

UC Berkeley

UC Berkeley Previously Published Works

Title

ViKiNG: Vision-Based Kilometer-Scale Navigation with Geographic Hints

Permalink

<https://escholarship.org/uc/item/19v5c3gr>

Authors

Shah, Dhruv
Levine, Sergey

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

ViKiNG: Vision-Based Kilometer-Scale Navigation with Geographic Hints

Dhruv Shah, Sergey Levine
UC Berkeley

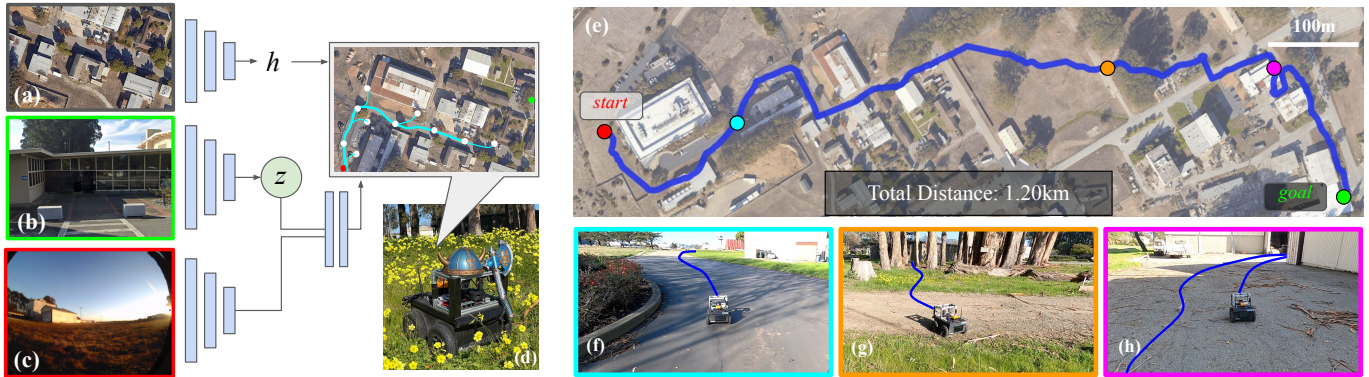


Fig. 1: Kilometer-scale autonomous navigation with ViKiNG: Our learning-based navigation system takes as input the current egocentric image (c), a photograph of the desired destination (b), and an overhead map (which may be a schematic or satellite image) (a) that provides a *hint* about the surrounding layout. The robot (d) uses learned models trained in *other* environments to infer a path to the goal (e), combining local traversability estimates with global heuristics derived from the map. This enables ViKiNG to navigate *previously unseen* environments (e), where a single traversal might involve following roads (f), off-road driving under a canopy (g), and backtracking from dead ends (h).

Abstract—Robotic navigation has been approached as a problem of 3D reconstruction and planning, as well as an end-to-end learning problem. However, long-range navigation requires both planning and reasoning about local traversability, as well as being able to utilize general knowledge about global geography, in the form of a roadmap, GPS, or other side information providing important cues. In this work, we propose an approach that integrates learning and planning, and can utilize side information such as schematic roadmaps, satellite maps and GPS coordinates as a planning heuristic, without relying on them being accurate. Our method, ViKiNG, incorporates a local traversability model, which looks at the robot’s current camera observation and a potential subgoal to infer how easily that subgoal can be reached, as well as a heuristic model, which looks at overhead maps for hints and attempts to evaluate the appropriateness of these subgoals in order to reach the goal. These models are used by a heuristic planner to identify the best waypoint in order to reach the final destination. Our method performs no explicit geometric reconstruction, utilizing only a topological representation of the environment. Despite having never seen trajectories longer than 80 meters in its training dataset, ViKiNG can leverage its image-based learned controller and goal-directed heuristic to navigate to goals up to 3 kilometers away in previously unseen environments, and exhibit complex behaviors such as probing potential paths and backtracking when they are found to be non-viable. ViKiNG is also robust to unreliable maps and GPS, since the low-level controller ultimately makes decisions based on egocentric image observations, using maps only as planning heuristics. For videos of our experiments, please check out our project page: sites.google.com/view/viking-release

I. INTRODUCTION

Robotic navigation has conventionally been approached as a geometric problem, where the robot constructs a 3D model of the environment and then plans a path through this

model. End-to-end learning-based methods offer an alternative approach, where the robot learns to correlate observations with traversability information directly from experience, without full geometric reconstruction [1–3]. This can be advantageous because, in many cases, geometry alone is neither necessary nor sufficient to traverse an environment, and a learning-based method can acquire patterns that are more directly indicative of traversability, for example by learning that tall grass is traversable [4] while seemingly traversable muddy soil should be avoided. More generally, such methods can learn about common patterns in their environment, such as that houses tend to be rectangular, or that fences tend to be straight. These patterns can lead to common-sense inferences about which path should be taken through an unknown environment even before that environment has been fully mapped out [5].

However, dispensing with geometry entirely may also be undesirable: the spatial organization of the world provides regularities that become important for a robot that needs to traverse large distances to reach its goal. In fact, when humans navigate new environments, they make use of both geographic knowledge, obtained from overhead maps or other cues, and learned patterns [6]. But in contrast to SLAM, humans don’t require maps or auxiliary signals to be very accurate: a person can navigate a neighborhood using a schematic that roughly indicates streets and houses, and reach a house marked on it. Humans do not try to accurately reconstruct geometric maps, but use approximate “mental maps” that relate landmarks to each other topologically [7]. Our goal is to devise learning-enabled methods that similarly make use of geographic *hints*, which could take the form of GPS, roadmaps, or satellite

imagery, without requiring these signals to be perfect.

We consider the problem of navigation from raw images in a *novel* environment, where the robot is tasked with reaching a user-designated goal, specified as an egocentric image, as shown in Figure 1. Note that the robot has *no prior experience* in the target environment. The robot has access to geographic side information in the form of a schematic roadmap or satellite imagery (see Figure 2), which may be outdated, noisy, and unreliable, and approximate GPS. This information, while not sufficient for navigation by itself, contains useful cues that can be used by the robot. The robot also has access to a large and diverse dataset of experience from *other* environments, which it can use to learn general navigational affordances. We posit that an effective way to build such a robotic system is to combine the strengths of machine learning with informed search, by incorporating the geographic hints into a learned heuristic for search. The robot uses approximate GPS coordinates and an overhead map as geographic side information to help solve the navigation task, but does not assume that this information is particularly accurate—resembling a person using a paper map, the robot uses the GPS localization and an overhead map as *hints* to aid in visual navigation. Note that while we do assume access to GPS, the measurements are only accurate up to 2-5 meters (4-10 \times the scale of the robot), and cannot be used for local control.

The primary contribution of this work is ViKiNG, an algorithm that combines elements of end-to-end learning-based control at the low level with a higher-level heuristic planning method that uses this image-based controller in combination with the geographic hints. The local image-based controller is trained on large amounts of prior data from *other* environments, and reasons about navigational affordances directly from images without any explicit geometric reconstruction. The planner selects candidate waypoints in order to reach a faraway goal, incorporating the geographic side information as a planning heuristic. Thus, when the hints are accurate, they help the robot navigate toward the goal, and when they are inaccurate, the robot can still rely on its image observations to search through the environment. We demonstrate ViKiNG on a mobile ground robot and evaluate its performance in a variety of open-world environments not seen in the training data, including suburban areas, nature parks, and a university campus. Our local controller is trained on 42 hours of navigational data, and we test our complete system in 10 different environments. Despite never seeing trajectories longer than 80 meters in its training data, ViKiNG can effectively use geographic side information in the form of overhead maps to reach user-specified goals in *previously unseen environments* over 2 kilometers away in under 25 minutes.

II. RELATED WORK

Robotic navigation has been studied from a number of different perspectives in different fields. Classically, it is often approached as a problem of geometric mapping or reconstruction followed by planning [8]. In unknown environments,

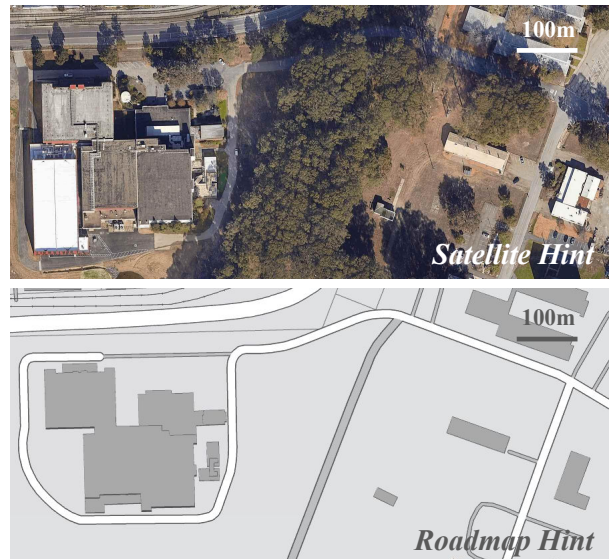


Fig. 2: Geographic hints used by ViKiNG. We evaluate our method with either satellite images or schematic roadmaps, though the approach could be used with any other information of this form, such as contour maps.

the mapping problem can be formulated in terms of information gain with local strategies [9–11], global strategies based on the frontier method [12–15], or by sampling “next-best views” [16–18], but such methods typically aim to map or reconstruct an entire environment, rather than achieve a single navigational goal. Active exploration methods have sought to modify this by jointly incentivizing an exploration objective along with reconstruction of the map [19, 20]. Both the goal-directed and mapping-focused methods aim to reconstruct the geometry of their environment, and do not directly benefit from training with prior data. Some approaches have sought to incorporate learning into mapping and reconstruction [21, 22], which benefits from prior data, but still aims at dense geometric reconstruction. Our approach uses a model that is trained with data from prior environments to predict *traversability* rather than geometry, and this model is then used in combination with geographic hints to plan a path to the goal.

In this respect, ViKiNG is also related to prior work on learning-based navigation, which is often formulated in terms of the “PointGoal” task [23]. Many such works rely on simulation and reinforcement learning, utilizing millions (or billions) of online trials to train a policy [24, 25]. In contrast, our method learns entirely from previously collected offline data, extrapolates to significantly longer paths than it is trained on, and does not require any simulation or online RL.

A number of prior learning-enabled methods also combine learned models with graph-based planning, using a topological graph to represent the environment [26–30]. These methods often assume access to data from the test environment to start with a viable graph, which may not be available in a new environment. Some works have studied this unseen setting by predicting explorable areas for semantically rich parts of the environment to accelerate visual exploration [31, 32]. While these methods can yield promising results in a variety of

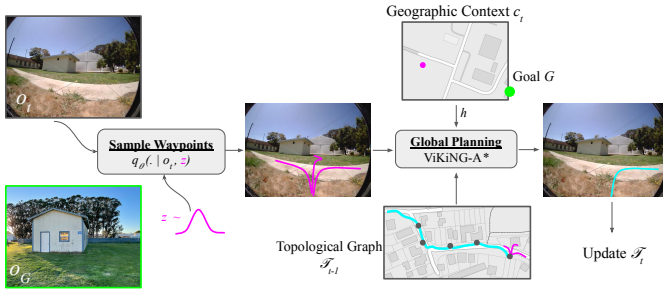


Fig. 3: An overview of our method. ViKiNG uses latent subgoals z proposed by a learned low-level controller, which operates on raw image observations o_t , for global planning on a topological graph \mathcal{T} to reach a distant goal o_G , indicated by a photograph and an approximate GPS location. A learned heuristic parses the overhead image c_t to bias this search towards the goal.

domains, they come at the cost of high sample complexity (over 10M samples) [23], making them difficult to use in the real world—the most performant algorithms take 10-20 minutes to find goals up to 50m away [33].

The closest prior work to ViKiNG is by Shah et al. (RECON) [33], which uses a learned representation over feasible subgoals to uniformly explore the environment. Like RECON, our method trains a local model that predicts temporal distances and actions for nearby subgoals, and then incorporates this model into a search procedure that incrementally constructs a topological graph in a novel environment. However, in contrast to RECON, which performs an uninformed search, ViKiNG incorporates geographic hints in the form of approximate GPS coordinates and overhead maps. This enables ViKiNG to reach faraway goals, up to $25\times$ further away than the furthest goal reported by RECON, and to reach goals up to $15\times$ faster than RECON when exploring a novel environment.

III. VISUAL NAVIGATION WITH GEOGRAPHIC HINTS

Our aim is to design a robotic system that learns to use first-person visual observations to reach user-specified landmarks, while also utilizing geographic *hints* in the form of approximate GPS coordinates and overhead maps. At the core of our approach is a deep neural network that takes in the robot’s current camera observation o_t , as well as an observation o_w of a potential subgoal w (we use “subgoal” and “waypoint” interchangeably), and predicts the time to reach w (or “temporal distance”), the best current action to do so, and the resulting spatial offset in terms of GPS coordinates. This model can also sample latent representations of potential *reachable* waypoints from the current observation o_t , which are used as candidate subgoals for planning. The model is trained on large amounts of data from a variety of training environments and, when the robot is placed in a *new* environment that it has not seen, it is used to incrementally construct a *topological* (non-geometric) graph to navigate to a distant user-specified goal. This goal is indicated by a photograph with an approximate GPS coordinate, and may be several kilometers away. The learned model alone is insufficient to navigate to such a distant goal in one shot, and therefore our planner uses a combination of the model’s predictions and geographic information to plan

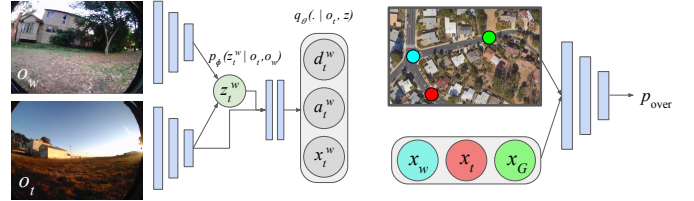


Fig. 4: The learned models used by ViKiNG. The latent goal model (left) takes in the current image o_t . It also takes in either a true waypoint image o_w , or samples a *latent* waypoint $z_t^w \sim r(z_t^w)$ from a prior distribution, and then predicts, its temporal distance from o_t (d_t^w), the action to reach it (a_t^w), and its approximate GPS offset (x_t^w). The heuristic model (right) takes in an overhead image c_t , the approximate GPS coordinates of the current location (x_t) and destination (x_G), and the coordinates of the waypoint inferred by the latent goal model (x_w), and predicts an approximate heuristic value of the waypoint w for reaching the final destination.

a sequence of subgoals that search for a path through the environment, incrementally constructing the graph.

This process corresponds to a kind of heuristic search, where the geographic side information provides a heuristic to bias the robot to explore towards the goal as it constructs the topological graph. The latent goal model is used to determine reachability in this topological graph, and the geographic heuristic is used to steer the graph by exploring the environment. In a novel environment, the robot must incrementally build this graph using *physical search*, by visiting new nodes and expanding its frontier. The decision about *where* to actually go is determined by the first-person images, and the geographic information is used only as a heuristic, allowing ViKiNG to remain robust to noisy or unreliable side information. We overview our method in Figure 3.

A. Low-level Control with a Latent Goal Model

Our low-level model maps the current image observation o_t and a waypoint observation o_w to: (1) the temporal distance d_t^w to reach w from o_t ; (2) the first action a_t^w that the robot must take now to reach w ; (3) a prediction of the (approximate) offset in GPS readings between o_t and w , x_t^w . (1) and (3) will be used by the higher-level planner, and (2) will be used to drive to w , if needed. We would also like this model to be able to *propose*, in a learned latent space, potential subgoals w that are reachable from o_t , and predict their corresponding values of d_t^w , a_t^w , and x_t^w .

We present the model in Figure 4, with precise architecture details in the supplementary materials. The model is trained by sampling pairs of time steps in the trajectories in the training set. For each pair, the earlier time step image becomes o_t , and the later image becomes o_w . The number of time steps between them provides the supervision for d_t^w , the action taken at the earlier time step supervises a_t^w , and the later GPS reading is transformed into the coordinate frame of the earlier time step to provide supervision for x_t^w . The model is trained via maximum likelihood. Note that by training the model on data in this way, we not only enable it to evaluate reachability of prospective waypoints, but also make it possible to inherit behaviors observed in the data. For example, in our

experiments, we will show that the model has a tendency to follow sidewalks and forest trails, a behavior it inherits from the portion of the dataset that is collected via teleoperation.

Besides predicting d_t^w , a_t^w , and x_t^w , our planner requires this model to be able to *sample* potential reachable waypoints from o_t (see Figure 3). We implement this via a variational information bottleneck (VIB) inside of the model that bottlenecks information from o_w . Thus, the model can either take as input a real image o_w of a prospective waypoint, or it can sample a *latent* waypoint $z_t^w \sim r(z_t^w)$ from a prior distribution. We train the model so that sampled latent waypoints correspond to feasible locations that the robot can reach from o_t without collision.

Training the latent goal model: The full model, illustrated in Figure 4, can be split into three parts: a waypoint encoder $p_\phi(z_t^w | o_w, o_t)$, a waypoint prior $r(z_t^w)$, and a predictor $q_\theta(\{a, d, x\}_t^w | z_t^w, o_t)$. The latent waypoint representation z_t^w can either be sampled from the prior (which is fixed to $r(z_t^w) \triangleq \mathcal{N}(0, I)$), or from the encoder $p_\phi(z_t^w | o_w, o_t)$ if a waypoint image o_w is provided. This latent waypoint is used together with o_t to predict all desired quantities according to $q_\theta(\{a, d, x\}_t^w | z_t^w, o_t)$. The training set consists of tuples $(o_t, o_w, \{a, d, x\}_t^w)$, but the model must be trained so that samples $z_t^w \sim r(z_t^w)$ also produce valid predictions. We accomplish this by means of the VIB [34], which regularizes the encoder $p_\phi(z_t^w | o_w, o_t)$ to produce distributions that are close to the prior $r(z_t^w)$ in terms of KL-divergence. We refer the reader to prior work for a derivation of the VIB [34], and present our training objective for p_ϕ and q_θ below:

$$\mathcal{L}_{\text{VIB}}(\theta, \phi) = -E_{\mathcal{D}}[\mathbb{E}_{p_\phi}[\log q_\theta(\{a, d, x\}_t^w | z_t^w, o_t)] + \beta \text{KL}(p_\phi(z_t^w | o_w, o_t) || r(z_t^w))] \quad (1)$$

The outer expectation over all tuples $(o_t, o_w, \{a, d, x\}_t^w) \in \mathcal{D}$ in the training distribution is estimating using the training set. The first term causes the model to accurately predict the desired information, while the second term forces the encoder to remain consistent with the prior, which makes the model suitable for sampling latent waypoints according to $z_t^w \sim r(z_t^w)$. As the encoder p_ϕ and decoder q_θ are conditioned on o_t , the representation z_t^w only encodes *relative* information about the subgoal from the context—this allows the model to represent *feasible* subgoals in new environments, and provides a compact representation that abstracts away irrelevant information, such as time of day or visual appearance. An analogous representation has been proposed in prior work [33], but did not predict spatial offsets and was used only for *uninformed* exploration without geographic hints.

B. Informed Search on a Topological Graph

The model described above can effectively reach nearby subgoals, for example those on which the robot has line of sight, but we wish to reach goals that are more than a kilometer away. To reach distant goals, we combine the model with a search procedure that incorporates geographic hints from satellite images or roadmaps. The system does not require this information to be accurate, instead using it as a planning

Algorithm 1 ViKiNG-A* for Physical Search

```

1: function ViKiNG-A*(start  $S$ , goal info  $o_G, x_G$ )
2:    $\Omega \leftarrow \{S\}$ 
3:   while  $\Omega$  not empty do
4:      $w_t \leftarrow \min(\Omega, f)$ 
5:     DriveTo( $w_t$ )  $\triangleright$  update visitations  $v$  on the way
6:     observe image  $o_t$ 
7:     add  $w_t$  to graph  $\mathcal{T}$   $\triangleright$  use  $q_{\theta, \phi}$  on  $o_t$  to get distances
8:     if close( $o_t, o_G$ ) finish  $\triangleright$  use  $q_{\theta, \phi}(\{a, d, x\}_t^w | o_t, o_G)$ 
9:     remove  $w_t$  from  $\Omega$ 
10:    sample waypoints  $w$  near  $w_t$  (Section III-A)
11:    for each  $w$  sampled near  $w_t$  do
12:      if not contains( $\Omega, w$ ) then add  $w$ 
13:      for each waypoint  $w \in \Omega$  do
14:         $f(w) = g(t, w) + d_{\text{Pr}[w]}^w + h(w) + v(\text{Pr}[w])$ 
15:    return failure

```

heuristic while still relying on egocentric camera images for control. Our high-level planner plans over a topological graph \mathcal{T} that it constructs incrementally using the low-level model in Section III-A as a local planner. We first describe a generic version of the algorithm for any heuristic, and then describe the data-driven heuristic function that we extract from the geographic hints via contrastive learning.

Challenges with physical search: Our “search” process involves the robot physically searching through the environment, and is not purely a computational process. In contrast to standard search algorithms (e.g., Dijkstra, A*, IDA*, D*, etc.), each “step” of our search involves the robot driving to a subgoal and updating the graph. Standard graph search algorithms assume (i) the ability to *visit* any arbitrary node, and (ii) access to a set of *neighbors* for every node and the corresponding “edge weight,” before visiting each neighbor. Physical search with a robot violates these assumptions, since robots cannot “teleport” and *visiting* a node incurs a driving cost. Furthermore, the real world does not provide “edge weights” and the robot needs to estimate the cost to reach an unvisited node *before* actually driving to it.

An algorithm for informed physical search: To solve these challenges, we design ViKiNG-A*, an A*-like search algorithm that uses our latent goal model and a learned heuristic to perform *physical search* in real-world environments. While ViKiNG-A* does prefer shorter paths, it does not aim to be optimal (in contrast to A*), only to reach the goal successfully. We will use a heuristic $h(w)$, fully described in the next section, which we assume provides a *comparative* evaluation of candidate waypoints in terms of their anticipated temporal distance to the destination. Algorithm 1 outlines ViKiNG-A*.

Like A*, ViKiNG-A* maintains a priority queue “open set” Ω of unexplored fringe nodes and a “current” node that represents the least-cost node in this set, which we refer to as w_t . It also maintains a graph with visited waypoints, \mathcal{T} , where nodes correspond to images seen at those nodes, and edges correspond to temporal distances estimated by the model

in Section III-A. At every iteration, the robot *drives* to the least-cost node in the open set (L5), using a procedure that we outline later. When it reaches w_t , it observes the image o_t using its camera (L6). This allows it to add o_t to the graph \mathcal{T} (L7), connecting it to other nodes by evaluating the distances using the model in Section III-A. The graph construction is analogous to prior work [29, 33]. If the robot is close to the final goal image o_G according to the model (L8), the search ends. o_t also allows it to sample nearby candidate waypoints using the model in Section III-A (L10): first sampling $z_t^w \sim r(z_t^w)$ from the prior, and then decoding distances d_t^w , a_t^w , and x_t^w , from which it can compute absolute locations as $x_w = x_t + x_t^w$. Each sampled waypoint is stored in the open set, and annotated with the current image o_t and d_t^w . We refer to w_t as the *parent* of w , and index it as $\text{Pr}[w]$. Note that we do *not* have access to the image o_w , as we have not visited the sampled waypoint w yet, and therefore we must store the current image o_t instead. This also means that we *cannot* connect these waypoints to the graph \mathcal{T} except through their parent. Next, we re-estimate the cost of each waypoint in the open set, including the newly added waypoints.

The cost for each waypoint $w \in \Omega$ from the current point w_t consists of four terms (L14): (1) $g(t, w)$, the cost to navigate to the *parent* of w , which is part of the graph \mathcal{T} ; this can be computed as a shortest path on the graph \mathcal{T} , and is zero for the current node. (2) $d_{\text{Pr}[w]}^w$, the distance from the parent of w to w itself. (3) $h(w)$, the heuristic cost estimate of reaching the final goal from w (see Section III-C). (4) $v(\text{Pr}[w])$, the visitation count of $\text{Pr}[w]$, computed as $CN(\text{Pr}[w])$, where C is a constant and $N(\text{Pr}[w])$ is a count of how many times the robot drove to $\text{Pr}[w]$ via the DriveTo subroutine; this acts as a *novelty bonus* to encourage the robot to explore novel states, a strategy widely used in RL [35, 36]. Summing these terms expresses a preferences for nodes that are fast to reach from w_t (1 + 2), get us closer to the goal (3), and have not been heavily explored before (4). At the next iteration (L4), the robot picks the lowest-cost waypoint and again drives to it.

To navigate to a selected waypoint w (DriveTo), the robot employs a procedure analogous to prior work on learning-based navigation with topological graphs [29, 33], planning the shortest path through \mathcal{T} , and selecting the next waypoint on this path. Once the waypoint w is selected, the model $q_{\theta, \phi}(\{a, d, x\}_t^w | o_t, o_w)$ is used to repeatedly choose the action a_t^w based on the current image o_t , until the distance d_t^w becomes small, indicating that the waypoint is reached and the robot can navigate to the next waypoint (in practice, it’s convenient to replan the path at this point, as is standard in MPC). Each time the DriveTo subroutine reaches a node, it also increments its count $N(w)$ which is used for the novelty bonus $v(w)$. The helper function *close* uses the model in Section III-A to check if the estimated temporal distance d_t^w is less than ϵ for two observations, and the *contains* operation on a set checks if a given node is *close* to any node inside the set. These modifications allow A*-like operations on the nodes of our graph, which are continuous variables.

C. Learning a Goal-Directed Heuristic for Search

We now describe how we extract a heuristic h from geographic side information. As a warmup, first consider the case where we only have the GPS coordinates for a waypoint (x_w) and final goal (x_G). We can use $\|x_G - x_w\|$ as a heuristic to bias the search to waypoints in the direction of the goal, and this heuristic can be readily obtained from the model in Section III-A. However, we would like to compute the heuristic function using some side information c_t , such as a roadmap or satellite image, that does not lie in a metric space. Thus, we need to *learn* the heuristic function from data. Since ViKiNG-A* does not aim to be optimal (only seeking a feasible path), we do not require the heuristic to be admissible.

We train the heuristic $h_{\text{over}}(x_w, x_G, x_t, c_t)$ to score the favorability of a sampled candidate waypoint w for reaching the goal G from current location x_t , given side information c_t . In our case, c_t is an overhead image that is roughly centered at the current location of the robot. Our heuristic is based on an estimator for the probability $p_{\text{over}}(w \rightarrow G | x_w, x_G, x_t, c_t)$ that a given waypoint w lies on a valid path to the goal G . We use the same training set as in Section III-A to learn a predictor for p_{over} . Given p_{over} , we can generate a heuristic $h_{\text{over}} := \lambda_{\text{over}}(1 - p_{\text{over}})$ to steer ViKiNG-A* towards the goal (Alg. 1 L13). Note that, since we evaluate the heuristic for sampled candidate waypoints, we do not have access to x_w , but we can predict it by using the model in Section III-A to infer the offset x_t^w using o_t and the sampled latent code, and then calculate x_w from x_t and x_t^w . Thus, the heuristic is technically a function of c_t , o_t , x_t , and x_G .

Our procedure for training $p_{\text{over}}(w \rightarrow G | x_w, x_G, x_t, c_t)$ is based on InfoNCE [37], a contrastive learning objective that can be seen as a binary classification problem between a set of *positives* and *negatives*. At each training iteration, we sample a random batch \mathcal{B} of sub-trajectories k from our training set, where x_S is the start of k and x_E is the end, and c_S is an overhead image centered at x_S . We sample a positive example by picking a random time step in this subtrajectory, and using its position x_{w^+} . The negatives x_{w^-} are locations of other randomly sampled time steps from other trajectories, comprising the set \mathcal{W}^- . In this way, we train a neural network model to represent $p_{\text{over}}(w \rightarrow G | x_w, x_G, x_t, c_t)$ (see Figure 4, right) via the InfoNCE objective:

$$\mathcal{L}_{\text{NCE}} = -\mathbb{E}_{\mathcal{B}} \left[\log \frac{p_{\text{over}}(w^+ \rightarrow E | x_{w^+}, x_E, x_S, c_S)}{\sum_{w^- \in \mathcal{W}^-} p_{\text{over}}(w^- \rightarrow E | x_{w^-}, x_E, x_S, c_S)} \right] \quad (2)$$

This heuristic can only reason about waypoints and goals at the scale of individual trajectories in the training set (up to 50m). For kilometer-scale navigation, the heuristic needs to make predictions for goals that are much further away, so we take inspiration from goal chaining in reinforcement learning [38] and combine overlapping trajectories in the training set (according to GPS positions) into larger trajectory groups. For a batch \mathcal{B} of trajectories, we combine two trajectories if they intersect in 2D space. The resulting macro-trajectories thus have multiple start and goal positions, and can extend

for several kilometers. We then sample the sub-trajectories for x_S , x_E , and x_{w^+} from these much longer macro-trajectories, giving us positive examples between very distant x_S , x_E pairs. This allows p_{over} to be trained on a vast pool of long-horizon goals and improves the reliability of the heuristic. We provide more details about this procedure in Appendix A.

IV. ViKiNG IN THE REAL WORLD

We now describe our experiments deploying ViKiNG in a variety of real-world outdoor environments for kilometer-scale navigation. Our experiments compare ViKiNG to other learning-based methods, evaluate its performance at different ranges, and study how it responds to degraded or erroneous geographic information.

A. Mobile Robot Platform

We implement ViKiNG on a Clearpath Jackal UGV platform (see Fig. 1). The default sensor suite consists of a 6-DoF IMU, a GPS unit for approximate global position estimates, and wheel encoders to estimate local odometry. Under open skies, the GPS unit is accurate up to 2-5 meters, which is 4-10 \times the size of the robot. In addition, we added a forward-facing 170 $^\circ$ field-of-view RGB camera. Compute is provided by an NVIDIA Jetson TX2 computer, and a cellular hotspot connection provides for monitoring and (if necessary) teleoperation. Our method uses only the monocular RGB images from the onboard camera, unfiltered measurements from onboard GPS, and overhead images (roadmap or satellite) queried at the current GPS location, without any other processing.

B. Offline Training Dataset

Our aim is to leverage data collected in a wide range of different environments to (i) enable the robot to learn navigational affordances that generalize to novel environments, and (ii) learn a global planning heuristic to steer physical search in novel environments. To create a diverse dataset capturing a wide range of navigation behavior, we use 30 hours of publicly available robot navigation data collected using an autonomous, randomized data collection procedure in office park style environments [33]. We augmented this dataset with another 12 hours of teleoperated data collected by driving on city sidewalks, hiking trails, and parks. Notably, ViKiNG never sees trajectories longer than 80 meters, but is able to leverage the learned heuristic (Section III-C) to reach goals over a kilometer away at over 80% of the average speed in the training set. The average trajectory length in the dataset is 45m, whereas our experiments evaluate runs in excess of 1km. The average velocity in the dataset is 1.68 m/s, and the average velocity the robot maintains in testing is 1.36 m/s. We provide more details about the dataset in Appendix B.

C. Kilometer-Scale Testing

For evaluation, we deploy ViKiNG in a variety of *previously unseen* open-world environments to demonstrate kilometer-scale navigation. Figure 5 shows the path taken by the robot in search for a user-specified goal image and location. ViKiNG

is able to utilize geographic hints, in the form of a roadmap or satellite image centered at its current position, to steer its search of the goal. In a university campus (Fig. 5(a, c)), we observe that the robot can identify large buildings along the way and plan around it, rather than following a greedy strategy. Since the training data often contains examples of the robot driving around buildings, ViKiNG is able to leverage this prior experience and generalize to novel buildings and environments. On city roads (Fig. 5(b)), the learned heuristic shows preference towards following the sidewalks, a characteristic of the training data in city environments. It is important to note that while the robot has seen some prior data on sidewalks and in suburban neighborhoods, it has never seen the specific areas (see Appendix B for further details). For videos of our experiments, please check out our project page.

These long-range experiments also exhibit successful backtracking behavior—when guided into a cul-de-sac by the planner, ViKiNG turns around and resumes its search for the goal from another node in the “openSet”, reaching the goal successfully (see Figure 1(h)). While the learned heuristic provides high-level guidance, the local control is done solely from first person images. This is illustrated in Figure 1(g), where the robot navigates through a forest, where the satellite image does not contain any useful information about navigating under a dense canopy. ViKiNG is able to successfully navigate through a patch of trees using the image-based model described in Section III-A. We can also provide ViKiNG with a set of goals to execute in a sequence to provide more guidance about the path (e.g., an inspection task with landmarks), as demonstrated in the next experiment.

A hiking ViKiNG: We deploy ViKiNG, with access to satellite images as hints, on a 2.7km hiking trail with a 70m elevation gain by providing a sequence of six checkpoint images and their corresponding GPS coordinates. Algorithmically, we run ViKiNG-A* on every goal (one at a time) while reusing the topological graph \mathcal{T} across goals. Figure 6 shows a top-down view of the path taken by the robot—ViKiNG is able to successfully combine the strengths of a learned controller for collision-free navigation with a learned heuristic that utilizes the satellite images to encourage on-trail navigation between checkpoints. Since the offline dataset contains examples of trail-following, the robot learns to stay on trails when possible. This behavior is emergent from the data—there is no other mechanism that encourages staying on the trails, and in several cases, a straight-line path between the goal waypoints would not stay on the trail (e.g., the first checkpoint in Figure 6).

Autonomous visual inspection: We further deploy ViKiNG in a suburban environment for the task of visual inspection specified by five images of interest. ViKiNG is able to successfully navigate to the landmarks by using satellite imagery, traveling a distance of 2.65km without any interventions. Figure 7 shows the specified images and a top-down view of the path taken by the robot on the trail.



Example Egocentric Observations Through Trajectory



Fig. 5: Examples of kilometer-scale goal-seeking in *previously unseen* environments using only egocentric images (right) and a schematic roadmap or satellite image as *hints* (left). ViKiNG can navigate in complex environments composed of roads, meadows, trees and buildings.

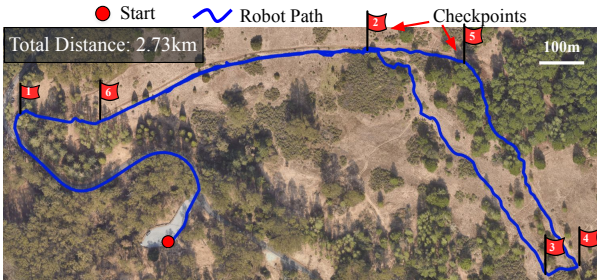


Fig. 6: ViKiNG can follow a sequence of goal checkpoints to perform search in complex environments, such as this 2.73km hiking trail.



Fig. 7: ViKiNG can utilize a satellite image to follow a sequence of visual landmarks (top) in complex suburban environments, such as this 2.65km loop stretching across buildings, meadows and roads.

D. Quantitative Evaluation and Comparisons

We compare ViKiNG to four prior approaches, each trained using the same offline data as our method. All methods have access to the egocentric images, GPS location, and satellite images, and control the robot via the same action space,

corresponding to linear and angular velocities.

Behavioral Cloning: A goal-conditioned behavioral cloning (BC) policy trained on the offline dataset that maps the three inputs to control actions [29, 39].

PPO: A policy gradient algorithm that maps the three inputs to control actions. This comparison is representative of state-of-the-art “PointGoal” navigation in simulation [25].

GCG: A model-based algorithm that uses a predictive model to plan a sequence of actions that reach the goal without causing collision [3]. We use GCG in the goal-directed mode with a GPS target, using the onboard camera and satellite images as input modalities.

RECON-H: A variant of RECON, which uses a latent goal model to represent reachable goals and plans over sampled subgoals to explore a novel environment [33]. We modify the algorithm to additionally accept the GPS and satellite images as additional inputs alongside the onboard camera image.

We evaluate the ability of ViKiNG to discover visually-indicated goals in 10 *unseen* environments of varying complexity. For each trial, we provide an RGB image of the desired target and its rough GPS location (accurate up to 5 meters). A trial is marked successful if the robot reaches the goal without requiring a human disengagement (due to a collision or getting stuck). We report the success rates of all methods in these environments in Table I and visualize overhead plots of the trajectories in one such environment in Figure 8.

ViKiNG outperforms all the prior methods, successfully navigating to goals that are over up to 500 meters away in our comparisons, including instances where no other method succeeds. RECON-H is the most performant of the other methods, successfully reaching most goals in the easier environments. Visualizing the robot trajectories (Fig. 8) reveals that RECON-H is unable to successfully utilize the geographic hints and explores greedily on encountering an obstacle. It also gets stuck and is unable to backtrack in 2/10 instances.

Method	Easy < 50m	Medium 50 – 150m	Hard 150 – 500m
Behavior Cloning	2/3	1/4	0/3
Offline PPO [40]	2/3	1/4	0/3
GCG [3]	3/3	2/4	0/3
RECON-H [33]	3/3	3/4	1/3
ViKiNG (Ours)	3/3	4/4	3/3

TABLE I: Comparison of goal-seeking performance against baselines. ViKiNG successfully reaches all goals. RECON-H and GCG succeed in simpler cases but are unable to utilize the hints effectively for distant goals. PPO and BC fail in all but the simplest cases.

Method	Avg. Displacement (m)	Avg. Velocity (m/s)
Behavior Cloning	19.5	0.35
Offline PPO [40]	47.2	0.85
GCG [3]	78.3	1.40
RECON-H [33]	188.3	0.41
ViKiNG (Ours)	250.0+	1.36

TABLE II: Average robot displacement and velocity before disengagement. ViKiNG successfully reaches all goals without requiring any disengagements. RECON-H also reaches some distant goals, but the low avg. velocity suggests that it takes an efficient path.

While GCG also performs well in simpler environments, it is limited by its planning horizon (up to 5 seconds) and gets stuck. PPO and BC both are both unable to learn from prior data and produce collisions with bushes and a parked car, respectively. In contrast, ViKiNG is able to effectively use the local controller to avoid the obstacles and reach the goal.

Analyzing the performance in the harder tasks with ranges of up to 500 meters (Table II), the average displacements and velocities before a user disengagement (due to collision or getting stuck) during these runs further confirm that ViKiNG is able to effectively use the geographic hints to steer the search without running into obstacles. While RECON-H manages to reach some faraway goals, it takes a greedy path to do so and is over $3\times$ slower than ViKiNG (see Fig. 8).

V. THE ROLE OF GEOGRAPHIC HINTS

In this section, we closely examine the role of geographic hints on the performance of ViKiNG by studying how it deals with a low-fidelity roadmap (versus a satellite image), and with incorrect hints and degraded geographic information. For the

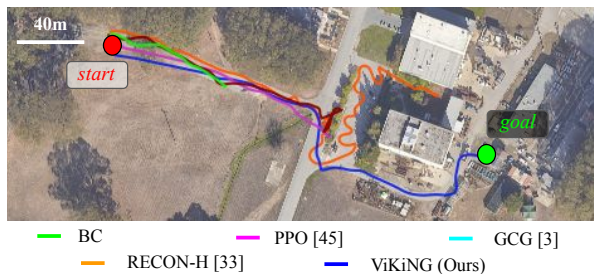


Fig. 8: Trajectories taken by the methods in a *previously unseen* environment. Only ViKiNG is able to effectively use the overhead images to reach the goal (270m away) successfully, following a smooth path around the building. RECON-H and GCG get stuck, while PPO and BC result in collisions.



Fig. 9: ViKiNG can use geographic hints in the form of a schematic roadmap or a satellite image. Providing roadmap hints encourages ViKiNG to follow marked roads (left); with satellite images, it is able to find a more direct path by cutting across a meadow (right).

experiment in Section V-A, we use models trained on the same dataset, but using schematic roadmaps as geographic hints. In Sections V-B and V-C, we use the same satellite image model from Section IV, with no additional retraining to accommodate missing or imperfect geographic information.

A. Comparing Different Types of Hints

To understand the nature of hints learned by the heuristic for different sources of geographic side information, we compare two separate versions of ViKiNG: one trained with schematic roadmaps as hints, and another trained with satellite images. Note that the method is identical in both cases, only the hint image in the data changes. For identical start-goal pairs, we observe that a model trained with roadmaps prefers following marked roads, whereas one trained with satellite images often cuts across patches of traversable terrain (e.g., grass meadows or trails) to take the quicker path, despite being trained on the same data. We hypothesize that this is due to the ability of the learned models to extract better correlations from the feature-rich satellite images, in contrast to the more abstract roadmap. Figure 9 shows a top-down view of the paths taken by the robot in the two cases in one such experiment.

B. Outdated Hints

To test the robustness of ViKiNG to outdated hints, we set up a goal-seeking experiment in one of the earlier environments and added a new obstacle—a large truck—blocking the path that ViKiNG took in the original trial. Since the satellite images are queried from a pre-recorded dataset, they do not reflect the addition of the truck, and hence continue to show a feasible path. We observe that the robot drives up to the truck and takes an alternate path to the goal, without colliding with it (see Figure 10). The lower-level latent goal model is robust to such obstacles and only proposes *valid* subgoal candidates that do not lead to collision; since the learned heuristic only evaluates valid subgoals, ViKiNG is robust to small discrepancies in the hints.

C. Incorrect Hints

Next, we set up a goal-seeking experiment in one of the easy environments with modified GPS measurements, so that



Fig. 10: On navigating with outdated hints, like the truck (top right) that is absent in the satellite image, ViKiNG uses its learned local controller to propose feasible subgoals that avoid obstacles and finds a new path (blue) to the goal that avoids the truck.

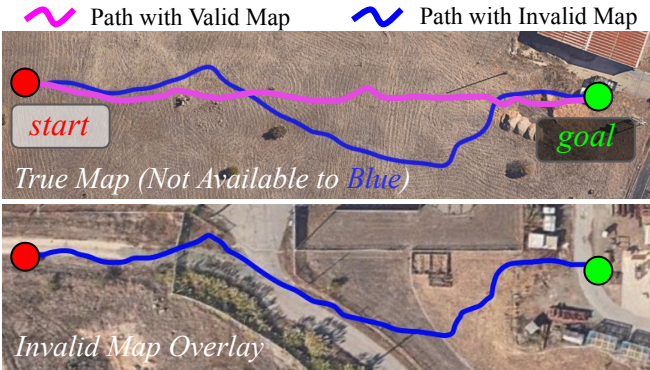


Fig. 11: On navigation with invalid hints, like the map at a different location, ViKiNG deviates from its original path (magenta) and reaches the goal by following the learned heuristic (blue).

the satellite images available to ViKiNG are offset by a $\sim 5\text{km}$ constant. As a result, this *hints* to the robot that there may be a road that it should follow, where in fact there isn't one (see Figure 11). We observe that the robot indeed deviates from its earlier path (with a valid map, the robot drives straight to the goal); upon overlaying this trajectory on the invalid map, we find that the learned heuristic indeed encourages the robot to follow the curvature of the road, but this path is still successful because it corresponds to open space.

D. A Disoriented ViKiNG

Finally, we analyze the effects of *disabling* the geographic hints and GPS localization on the goal-seeking performance of ViKiNG. Towards this, we run two variants of our algorithm: **No Overhead Image:** We provide the robot with GPS, but no satellite images. To accommodate this, we use a simple ℓ_2 heuristic $h_{\text{GPS}}(x_w, x_g, x_t) = \|x_g - x_w\|$.

No GPS: The robot does not have access to GPS or satellite images. To accommodate this, we remove the heuristic h from

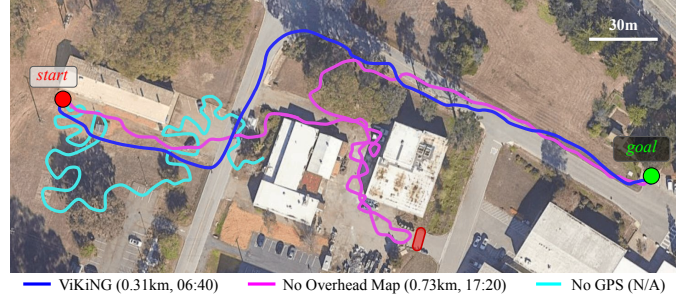


Fig. 12: Ablations of ViKiNG by withholding geographic hints. ViKiNG without overhead images (magenta) acts greedily, driving close to buildings, gets caught into a cul-de-sac and eventually reaches the goal $2.6\times$ slower than ViKiNG with access to satellite images (blue), which avoids the building cluster by following a smoother dirt path. Search without GPS (cyan) performs uninformed exploration and is unable to reach the goal in over 30 minutes.

ViKiNG-A*, making it an uninformed search algorithm.

Figure 12 summarizes the path taken by the robot, distance traversed, and time taken. When we disable the overhead hints and only use h_{GPS} , ViKiNG-A* can still reach the destination, but takes significantly longer to do so, initially exploring a dead-end path that it then has to back out of. That said, this experiment also illustrates the ability of ViKiNG-A* to handle less useful heuristics: while the path is significantly longer, the method is still able to eventually reach the destination, and in some sense the mistakes the method makes are to be expected of any system that has no prior map information. If we remove GPS as well, ViKiNG-A* corresponds to a Dijkstra-like uninformed search (resembling RECON [33]). In this case, the robot searches its environment without any guidance and is unable to reach the goal in over 30 minutes.

VI. DISCUSSION

We proposed a method for efficiently learning vision-based navigation in *previously unseen* environments at a kilometer-scale. Our key insight is that effectively leveraging a small amount of geographic knowledge in a learning-based framework can provide strong regularities that enable robots to navigate to distant goals. We find that incorporating geographic hints as goal-directed heuristics for planning enables emergent preferences such as following roads or hiking trails. Additionally, ViKiNG only uses the hints for biasing the high-level search; the learned control policy at the lower-level relies solely on egocentric image observations, and is thus robust to imperfect hints. While we only use overhead images in our experiments, an existing avenue for future work is to explore how such a system could use other information sources, including paper maps or textual instructions, which can be incorporated into our contrastive objective.

ACKNOWLEDGMENTS

This research was partially supported by DARPA Assured Autonomy, ARL DCIST CRA W911NF-17-2-0181, and DARPA RACER. The authors would like to thank Blazej Osinski, Brian Ichter and Katie Kang for useful discussions.

REFERENCES

- [1] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, “Target-driven visual navigation in indoor scenes using deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [2] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, “Deep-driving: Learning affordance for direct perception in autonomous driving,” in *IEEE International Conference on Computer Vision*, 2015.
- [3] G. Kahn, A. Villafior, B. Ding, P. Abbeel, and S. Levine, “Self-Supervised Deep RL with Generalized Computation Graphs for Robot Navigation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [4] G. Kahn, P. Abbeel, and S. Levine, “Badgr: An autonomous self-supervised learning-based navigation system,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, 2021.
- [5] M. Narasimhan, E. Wijmans, X. Chen, T. Darrell, D. Batra, D. Parikh, and A. Singh, “Seeing the un-scene: Learning amodal semantic maps for room navigation,” *CoRR*, 2020.
- [6] J. M. Wiener, S. J. Büchner, and C. Hölscher, “Taxonomy of human wayfinding tasks: A knowledge-based approach,” *Spatial Cognition & Computation*, 2009.
- [7] P. S. Foo, W. H. Warren, A. P. Duchon, and M. J. Tarr, “Do humans integrate routes into a cognitive map? map-versus landmark-based navigation of novel shortcuts.” *Journal of experimental psychology. Learning, memory, and cognition*, 2005.
- [8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, 2005.
- [9] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, “Information based adaptive robotic exploration,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 1, 2002, pp. 540–545 vol.1.
- [10] T. Kollar and N. Roy, “Efficient optimization of information-theoretic exploration in slam,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2008.
- [11] W. Tabib, K. Goel, J. Yao, M. Dabhi, C. Boirum, and N. Michael, “Real-time information-theoretic exploration with gaussian mixture model maps.” in *Robotics: Science and Systems*, 2019.
- [12] B. Yamauchi, “A frontier-based approach for autonomous exploration,” in *IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 1997.
- [13] D. Holz, N. Basilico, F. Amigoni, and S. Behnke, “A comparative evaluation of exploration strategies and heuristics to improve them,” 2011.
- [14] B. Charrow, S. Liu, V. Kumar, and N. Michael, “Information-theoretic mapping using cauchy-schwarz quadratic mutual information,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [15] B. Charrow, G. Kahn, S. Patil, S. Liu, K. Goldberg, P. Abbeel, N. Michael, and V. Kumar, “Information-theoretic planning with trajectory optimization for dense 3d mapping,” in *Robotics: Science and Systems (RSS)*, 2015.
- [16] C. Connolly, “The determination of next best views,” in *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, 1985.
- [17] C. Papachristos, S. Khattak, and K. Alexis, “Uncertainty-aware receding horizon exploration and mapping using aerial robots,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [18] M. Selin, M. Tiger, D. Duberg, F. Heintz, and P. Jensfelt, “Efficient autonomous exploration planning of large-scale 3-d environments,” *IEEE Robotics and Automation Letters*, 2019.
- [19] J. Delmerico, E. Mueggler, J. Nitsch, and D. Scaramuzza, “Active autonomous aerial exploration for ground robot path planning,” *IEEE Robotics and Automation Letters*, 2017.
- [20] I. Lluvia, E. Lazkano, and A. Ansuategi, “Active mapping and robot exploration: A survey,” *Sensors*, 2021.
- [21] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, “Semanticfusion: Dense 3d semantic mapping with convolutional neural networks,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [22] K. M. Jatavallabhula, G. Iyer, and L. Paull, “gradslam: Dense SLAM meets automatic differentiation,” *CoRR*, 2019.
- [23] Manolis Savva*, Abhishek Kadian*, Oleksandr Maksymets*, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, “Habitat: A Platform for Embodied AI Research,” in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [24] A. Kumar*, S. Gupta*, D. Fouhey, S. Levine, and J. Malik, “Visual Memory for Robust Path Following,” in *Neural Information Processing Systems (NeurIPS)*, 2018.
- [25] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, “DD-PPO: Learning Near-Perfect PointGoal Navigators from 2.5 Billion Frames,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [26] N. Savinov, A. Dosovitskiy, and V. Koltun, “Semi-Parametric Topological Memory for Navigation,” in *International Conference on Learning Representations*, 2018.
- [27] J. Bruce, N. Sunderhauf, P. Mirowski, R. Hadsell, and M. Milford, “Learning deployable navigation policies at kilometer scale from a single traversal,” in *Proceedings of The 2nd Conference on Robot Learning*, A. Billard, A. Dragan, J. Peters, and J. Morimoto, Eds., 2018.
- [28] A. Faust, K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson, “Prm-rl: Long-range robotic

navigation tasks by combining reinforcement learning and sampling-based planning,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5113–5120.

- [29] D. Shah, B. Eysenbach, G. Kahn, N. Rhinehart, and S. Levine, “ViNG: Learning Open-World Navigation with Visual Goals,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [30] X. Meng, N. Ratliff, Y. Xiang, and D. Fox, “Scaling Local Control to Large-Scale Topological Navigation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [31] D. Singh Chafflot, R. Salakhutdinov, A. Gupta, and S. Gupta, “Neural topological slam for visual navigation,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [32] D. S. Chafflot, H. Jiang, S. Gupta, and A. Gupta, “Semantic curiosity for active visual learning,” in *ECCV*, 2020.
- [33] D. Shah, B. Eysenbach, N. Rhinehart, and S. Levine, “Rapid exploration for open-world navigation with latent goal models,” in *5th Annual Conference on Robot Learning*, 2021.
- [34] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, “Deep variational information bottleneck,” *arXiv preprint arXiv:1612.00410*, 2016.
- [35] T. Lai and H. Robbins, “Asymptotically efficient adaptive allocation rules,” *Advances in Applied Mathematics*, 1985.
- [36] M. G. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, “Unifying count-based exploration and intrinsic motivation,” *arXiv preprint arXiv:1606.01868*, 2016.
- [37] A. van den Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” 2019.
- [38] Y. Chebotar, K. Hausman, Y. Lu, T. Xiao, D. Kalashnikov, J. Varley, A. Irpan, B. Eysenbach, R. Julian, C. Finn, and S. Levine, “Actionable models: Unsupervised offline reinforcement learning of robotic skills,” *arXiv preprint arXiv:2104.07749*, 2021.
- [39] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, “End-to-End Driving Via Conditional Imitation Learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [40] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” 2017.
- [41] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications,” 2017.
- [42] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.

A. Implementation Details

Layer	Input [Dimensions]	Output [Dimensions]	Layer Details
<i>Encoder</i> $p_\phi(z o_t, o_w) = \mathcal{N}(\cdot; \mu_p, \Sigma_p)$			
1	o_t, o_w [3, 160, 120]	I_t^w [6, 160, 120]	Concatenate along channel dimension.
2	I_t^w [6, 160, 120]	E_t^w [1024]	MobileNet Encoder [41]
3	E_t^w [1024]	μ_p [64], σ_p [64]	Fully-Connected Layer, exp activation of σ_p
4	σ_p [64]	Σ_p [64, 64]	torch.diag(σ_p)
<i>Decoder</i> $q_\theta(a, d, x o_t, z_t^w) = \mathcal{N}(\cdot; \mu_q, \Sigma_q)$			
1	o_t [3, 160, 120]	E_t [1024]	MobileNet Encoder [41]
2	E_t [1024], z_t^w [64]	$F = E_t \oplus z_t^w$ [1088]	Concatenate image and goal representation
3	F [1088]	μ_q [3], σ_q [3]	Fully-Connected Layer, exp activation of σ_q
4	σ_q [5]	Σ_q [5, 5]	torch.diag(σ_q)
5	μ_q [5]	\bar{a}_t^w [2], d_t^w [1], \bar{x}_t^w [2]	Split into actions, distances and offsets

TABLE III: Architectural details of the latent goal model (Section III-A)

1) *Latent Goal Model (Section III-A)*: Inputs to the encoder p_ϕ are pairs of observations of the environment—current and goal—represented by a stack of two RGB images obtained from the onboard camera at a resolution of 160×120 pixels. p_ϕ is implemented by a MobileNet encoder [41] followed by a fully-connected layer projecting the 1024-dimensional latents to a stochastic, context-conditioned representation z_t^w of the goal that uses 64-dimensions each to represent the mean and diagonal covariance of a Gaussian distribution. Inputs to the decoder q_θ are the context (current observation)—processed with another MobileNet—and z_t^w . We use the reparametrization trick [42] to sample from the latent and use the concatenated encodings to learn the optimal actions \bar{a}_t^w , temporal distances d_t^w and spatial offsets \bar{x}_t^w . Details of our network architecture are provided in Table III. During pretraining, we maximize \mathcal{L}_{VIB} (Eq. 1) with a batch size of 128 and perform gradient updates using the Adam optimizer with learning rate $\lambda = 10^{-4}$ until convergence.

2) *Learned Heuristic (Section III-C)*: Inputs to the encoder p_{over} are (i) satellite image c_S and (ii) the triplet of GPS locations $\{x_w, x_S, x_G\}$. p_{over} is implemented as a multi-input neural network with a MobileNet encoder [41] to featurize c_S , which is then concatenated with the location inputs. This is followed by a series of fully-connected layers [512, 128, 32, 1] down to a single cell to predict the binary classification scores. During pretraining, we minimize \mathcal{L}_{NCE} with a batch size of 256 and perform gradient updates using the Adam optimizer with learning rate $\lambda = 10^{-4}$ until convergence.

3) *Miscellaneous Hyperparameters*: We provide the hyperparameters associated with our algorithms in Table IV.

Hyperparameter	Value	Meaning
Δt	0.5	Time step of the robot (s)
ϵ	10	Threshold for close (Sec. III-B)
C	20	Scaling constant for v (Alg. 1 L14)
λ_{over}	200	Scaling constant for h_{over} (Sec. III-C)

TABLE IV: Hyperparameters used in our experiments.

B. Offline Trajectory Dataset

For the offline dataset discussed in Section IV-B, we use a combination of a 30 hours of autonomously collected data, and 12 hours of human teleoperated data. The complete dataset was collected by 3 independent sets of researchers over the course of 24 months in environments spanning multiple cities. We provide more information below.

1) *Autonomously Collected Data*: We use the published dataset by Shah et al. [33], that contains over 5000 self-supervised trajectories collected over 9 distinct real-world environments. These trajectories capture the interaction of the robot in diverse environments, including phenomena like collisions with obstacles and walls, getting stuck in the mud or pits, or flipping due to bumpy terrain.

During data collection, a robot is equipped with a 2D LIDAR sensor to detect collisions ahead of time and generate autonomous pseudo-labels for collision events. To ensure that the control policy achieves sufficient coverage of the environment while also ensuring that the action sequences executed by the robot are realistic, we use a time-correlated random walk to gather data.

	Training Dataset	ViKiNG Deployment
Avg. Length	45m	>1km
Avg. Velocity (m/s)	1.68	1.36

TABLE V: Trajectory statistics for offline training dataset and real-world deployment.

Environment Type	Amount of Data (hrs)
Paved Hiking Trails	01:45
City Sidewalks	02:15
Suburban Neighborhood Roads	01:30
Unpaved Grasslands	01:00
University/Office Campus	02:30
Miscellaneous	03:00
Total	12:00

TABLE VI: Approximate composition of various environment types in the teleoperated dataset.

2) *Human Teleoperated Data:* The above dataset contains extremely diverse dataset that is great for learning general notions of traversability and collision avoidance. However, the random nature of the dataset means that it does not contain any semantically interesting behavior that may be desired of a robotic system, such as following a sidewalk or through a patch of trees. To enhance the quality of learned behaviors, we augment this dataset with about 12 hours of human teleoperated data in semantically rich environments such as hiking trails, city sidewalks, parking lots and suburban neighborhoods. These environments represent realistic scenarios where such a robotic system would be deployed.

Table V summarizes key statistics of the trajectories, such as length and velocity. Table VI summarizes the various environments in which the dataset was collected, and their relative composition. Figure 13 visualizes the geographic locations of these data collection sites (location anonymized for the double-blind review process). We ensure no overlap between the training and test environments—success in these test environments requires *true* generalization to unseen environments.

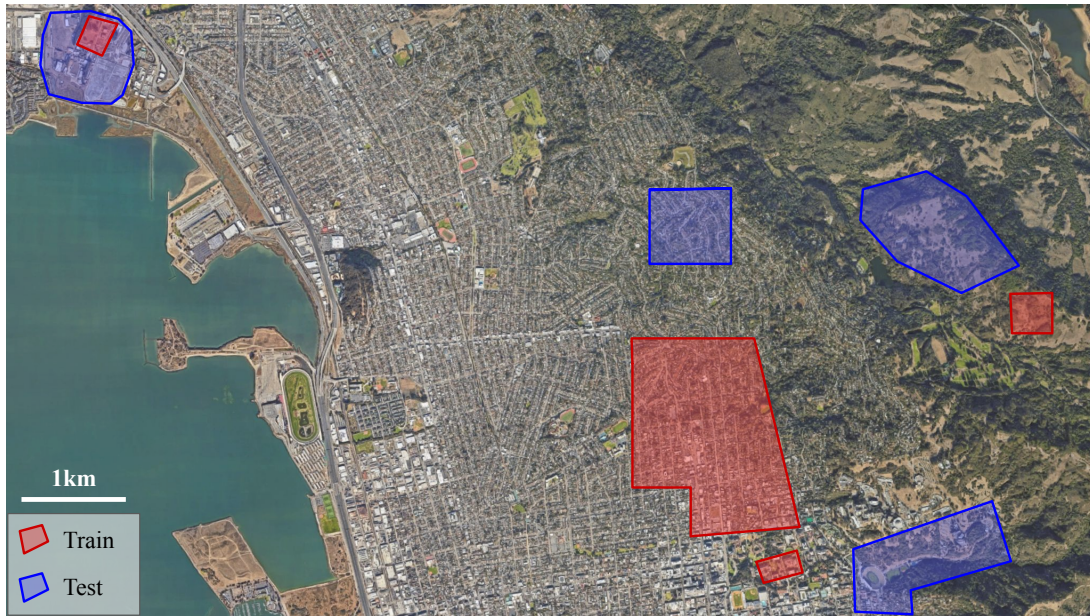


Fig. 13: Rough geographical locations of data collection by human teleoperation and testing (Section IV)

C. Project Page

We share experiment videos, including third-person perspectives of trajectories traversed by ViKiNG, on our project page: sites.google.com/view/viking-release.