



Sidekick Policy Learning for Active Visual Exploration

Santhosh K. Ramakrishnan¹(✉) and Kristen Grauman²

¹ The University of Texas at Austin, Austin, TX 78712, USA
srama@cs.utexas.edu

² Facebook AI Research, 300 W. Sixth Street, Austin, TX 78701, USA
grauman@fb.com

Abstract. We consider an *active visual exploration* scenario, where an agent must intelligently select its camera motions to efficiently reconstruct the full environment from only a limited set of narrow field-of-view glimpses. While the agent has full observability of the environment during training, it has only partial observability once deployed, being constrained by what portions it has seen and what camera motions are permissible. We introduce *sidekick policy learning* to capitalize on this imbalance of observability. The main idea is a preparatory learning phase that attempts simplified versions of the eventual exploration task, then guides the agent via reward shaping or initial policy supervision. To support interpretation of the resulting policies, we also develop a novel policy visualization technique. Results on active visual exploration tasks with 360° scenes and 3D objects show that sidekicks consistently improve performance and convergence rates over existing methods. Code, data and demos are available (Project website: <http://vision.cs.utexas.edu/projects/sidekicks/>).

Keywords: Visual exploration · Reinforcement learning

1 Introduction

Visual recognition has witnessed dramatic successes in recent years. Fueled by benchmarks composed of Web photos, the focus has been inferring semantic labels from *human-captured images*—whether classifying scenes, detecting objects, or recognizing activities [41, 51, 57]. By relying on human-taken images, the common assumption is that an intelligent agent will have already decided where and how to capture the input views. While sufficient for handling static

K. Grauman—On leave from University of Texas at Austin (grauman@cs.utexas.edu).

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01258-8_26) contains supplementary material, which is available to authorized users.




Fig. 1. Embodied agents that actively explore novel objects (left) or 360° environments (right) intelligently select camera motions to gain as much information as possible with very few glimpses. While they naturally face **limited observability** of the environment, during *learning* **fuller observability** may be available. We propose sidekicks to guide policy learning for active visual exploration.

repositories of photos (e.g., auto-tagging Web photos and videos), assuming informative observations glosses over a very real hurdle for embodied vision systems.

A resurgence of interest in perception tied to action takes aim at that hurdle. In particular, recent work explores agents that optimize their physical movements to achieve a specific perception goal, e.g., for active recognition [2, 28, 29, 31, 43], visual exploration [30], object manipulation [40, 46, 49], or navigation [2, 21, 70]. In any such setting, deep reinforcement learning (RL) is a promising approach. The goal is to learn a policy that dictates the best action for the given state, thereby integrating sequential control decisions with visual perception.

However, costly exploration stages and partial state observability are well-known impediments to RL. In particular, an active visual agent [21, 30, 70, 71] has to take a long series of actions purely based on the limited information available from its first person view. Due to poor action selection based on limited information, the most effective viewpoint trajectories are buried among many mediocre ones, impeding the agent’s exploration in complex state-action spaces.

We observe that agents lacking full observability when deployed may nonetheless possess full observability *during training*, in some cases. Overall, the imbalance occurs naturally when an agent is trained with a broader array of sensors than available at test-time, or trained free of the hard time pressures that limit test-time exploration. In particular, as we will examine in this work, once deployed, an active exploration agent can only move the camera to “look-around” nearby [30], yet if trained with omnidirectional panoramas, could access any possible viewpoint while learning. Similarly, an active object recognition system [2, 28, 29, 31, 65] can only see its previously selected views of the object; yet if trained with CAD models, it could observe all possible views while learning. Additionally, agents can have access to multiple sensors during training in simulation environments [10, 13, 48], yet operate on first-person observations during test-time. However, existing methods restrict the agent to the same partial observability during training [28–31, 65, 70].

We propose to leverage the imbalance of observability. To this end, we introduce *sidekick policy learning*. We use the name “sidekick” to signify how a sidekick to a hero (e.g., in a comic or movie) provides alternate points of view, knowledge, and skills that the hero does not have. In contrast to an *expert* [19,61], a sidekick *complements* the hero (agent), yet cannot solve the main task at hand.

We propose two sidekick variants. Both use access to the full state during a preparatory training period to facilitate the agent’s ultimate learning task. The first sidekick previews individual states, estimates their value, and shapes rewards to the agent for visiting valuable states during training. The second sidekick provides initial supervision via trajectory selections to accelerate the agent’s training, while gradually permitting the agent to act on its own. In both cases, the sidekicks learn to solve *simplified* versions of the main task with full observability, and use insights from those solutions to aid the training of the agent. At test time, the agent has to act without the sidekick.

We validate sidekick policy learning for active visual exploration [30]. The agent enters a novel environment and must select a sequence of camera motions to rapidly understand its entire surroundings. For example, an agent that has explored various grocery stores should enter a *new* one and, with a couple glimpses, (1) conjure a belief state for where different objects are located, then (2) direct its camera to flesh out the harder-to-predict objects and contexts. The task is like active recognition [2,29,31,65], except that the training signal is *pixel-wise* reconstruction error for the full environment rather than labeling error. Our sidekicks can look at any part of the environment in any sequence during training, whereas the actual agent is limited to physically feasible camera motions and sees only those views it has selected. On two standard datasets [65,66], we show how sidekicks accelerate training and promote better look around policies.

As a secondary contribution, we present a novel policy visualization technique. Our approach takes the learned policy as input, and displays a sequence of heatmaps showing regions of the environment most responsible for the agent’s selected actions. The resulting visualizations help illustrate how sidekick policy learning differs from traditional training.

2 Related Work

Active Vision and Attention: Linking intelligent control strategies to perception has early foundations in the field [1,5,6,63]. Recent work explores new strategies for active object recognition [2,28,29,31,65], object localization [9,20,71], and visual SLAM [32,58], in order to minimize the number of sampled views required to perform accurate recognition or reconstruction. Our work is complementary to any of the above: sidekick policy learning is a means to accelerate and improve active perception when observability is greater during training.

Models of saliency and attention allow a system to prioritize portions of its observation to reduce clutter or save computation [4,42,45,67,68]. However, unlike both our work and the active methods above, they assume full observability at test time, selecting among already-observed regions. Work in active

sensor placement aims to place sensors in an environment to maximize *coverage* [11, 36, 62]. We introduce a model for coverage in our policy learning solution (Sect. 3.3). However, rather than place and fix N static sensors, the visual exploration tasks entail selecting new observations dynamically and in sequence.

Supervised Learning with Observability Imbalance: Prior work in supervised learning investigates ways to leverage greater observability during training, despite more limited observability during test time. Methods for depth estimation [16, 22, 60] and/or semantic segmentation [25, 26, 56] use RGBD depth data, multiple views, and/or auxiliary annotations during training, then proceed with single image observations at test time. Similarly, self-supervised losses [27, 44] based on auxiliary prediction tasks at training time have been used to aid representation learning for control tasks. Knowledge distillation [24] lets a “teacher” network guide a “student” with the motivation of network compression. In learning with privileged information, an “expert” provides the student with training data having extra information (unavailable during testing) [37, 53, 61]. At a high level, all the above methods relate to ours in that a simpler learning task facilitates a harder one. However, in strong contrast, they tackle supervised classification/regression/representation learning, whereas our goal is to learn a *policy* for selecting actions. Accordingly, we develop a very different strategy—introducing rewards and trajectory suggestions—rather than auxiliary labels/modalities.

Guiding Policy Learning: There is a wide body of work aimed at addressing sparse rewards and partial observability. Several works explore *reward shaping* motivated by different factors. The intrinsic motivation literature develops parallel reward mechanisms, e.g., based on surprise [7, 47], to direct exploration. The TAMER framework [33–35] utilizes expert human rewards about the end-task. Potential-based reward shaping [23] incorporates expert knowledge grounded in potential functions to ensure policy invariance. Others convert control tasks into supervised measurement prediction task by defining goals and rewards as functions of measurements [12]. In contrast to all these approaches, our sidekicks exploit the observability difference between training and testing to transfer knowledge from a simpler version of the task. This external knowledge directly impacts the final policy learned by augmenting task related knowledge via reward shaping.

Behavior cloning provides expert-generated trajectories as supervised (state, action) pairs [8, 14, 17, 50]. *Offline planning*, e.g., with tree search, is another way to prepare good training episodes by investing substantial computation offline [3, 19, 54], but observability is assumed to be the same between training and testing. *Guided policy search* uses importance sampling to optimize trajectories within high-reward regions [39] and can utilize full observability [38], yet transfers from an expert in a purely supervised fashion. Our second sidekick also demonstrates good action sequences, but we specifically account for the observability imbalance by annealing supervision over time.

More closely related to our goal is the *asymmetric actor critic*, which leverages synthetic images to train a robot to pick/push an object [48]. Full state information from the graphics engine is exploited to better train the critic. While

this approach modifies the advantage expected for a state like our first sidekick, this is only done at the task level. Our sidekick injects a different perspective by solving simpler versions of the task, leading to better performance (Sect. 4.2).

Policy Visualization: Methods for post-hoc explanation of deep networks are gaining attention due to their complexity and limited interpretability. In supervised learning, heatmaps indicating regions of an image most responsible for a decision are generated via backprop of the gradient for a class label [15, 52, 55]. In reinforcement learning, policies for visual tasks (like Atari) are visualized using t-SNE maps [69] or heatmaps highlighting the parts of a *current* observation that are important for selecting an action [18]. We introduce a policy visualization method that reflects the influence of an agent’s *cumulative* observations on its action choices, and use it to illuminate the role of sidekicks.

3 Approach

Our goal is to learn a policy for controlling an agent’s camera motions such that it can explore novel environments and objects efficiently. Our key insight is to facilitate policy learning via sidekicks that exploit (1) full observability and (2) unlimited time steps to solve a simpler problem in a preparatory training phase.

We first formalize the problem setup in Sect. 3.1. After overviews observation completion as a means of active exploration in Sect. 3.2, we introduce our sidekick learning framework in Sect. 3.3. We tie together the observation completion and sidekick components with the overall learning objective in Sect. 3.4. Finally, we present our policy visualization technique in Sect. 3.5.

3.1 Problem Setup: Active Visual Exploration

The problem setting builds on the “learning to look around” challenge introduced in [30]. Formally, the task is as follows. The agent starts by looking at a novel environment (or object) X from some unknown viewpoint¹. It has a budget T of time to explore the environment. The learning objective is to minimize the error in the agent’s pixelwise reconstruction of the full—mostly unobserved—environment using only the sequence of views selected within that budget.

Following [30], we discretize the environment into a set of candidate viewpoints. In particular, the space of viewpoints is a viewgrid indexed by N elevations and M azimuths, denoted by $V(X) = \{x(X, \theta^{(i)}) | 1 \leq i \leq MN\}$, where $x(X, \theta^{(i)})$ is the 2D view of X from viewpoint $\theta^{(i)}$, which is comprised of two angles. More generally, $\theta^{(i)}$ could capture both camera angle and position; however, to best exploit existing datasets, we limit camera motions to rotations.

The agent expends the budget in discrete increments, called “glimpses”, by selecting $T - 1$ camera motions in sequence. At each time step, the agent gets

¹ For simplicity of presentation, we represent an *environment* as X where the agent explores a novel scene, looking outward in new viewing directions. However, experiments will also use X as an *object* where the agent moves around an object, looking inward at it from new viewing angles.

observation x_t from the current viewpoint. The agent makes an exploratory rotation (δ_t) based on its policy π . When the agent executes action $\delta_t \in \mathcal{A}$, the viewpoint changes according to $\theta_{t+1} = \theta_t + \delta_t$. For each camera motion δ_t executed by the agent, a reward r_t is provided by the environment (Sects. 3.3 and 3.4). Using the view x_t , the agent updates its internal representation of the environment, denoted $\hat{V}(X)$. Because camera motions are restricted to have proximity to the current camera angle (Sect. 4.1) and candidate viewpoints partially overlap, the discretization promotes efficiency without neglecting the physical realities of the problem (following [29–31, 43]).




Fig. 2. Active observation completion. The agent receives one view (shown in red), updates its belief and reconstructs the viewgrid at each time step. It executes an action (red arrows) according to its policy to obtain the next view. The active agent must rapidly refine its belief with well-chosen views. (Color figure online)

3.2 Recurrent Observation Completion Network

We start with the deep RL neural network architecture proposed in [30] to represent the agent’s recurrent observation completion. The process is deemed “completion” because the agent strives to hallucinate portions of the environment it has not yet seen. It consists of five modules: SENSE, FUSE, AGGREGATE, DECODE, and ACT with parameters W_s , W_f , W_r , W_d and W_a respectively.

- SENSE: Independently encodes the view (x_t) and proprioception (p_t) consisting of elevation at time t and relative motion from time $t - 1$ to t , and returns the encoded tuple $s_t = \text{SENSE}(x_t, p_t)$.
- FUSE: Consists of fully connected layers that jointly encode the tuple s_t and output a fused representation $f_t = \text{FUSE}(s_t)$.
- AGGREGATE: An LSTM that aggregates fused inputs over time to build the agent’s internal representation $a_t = \text{AGGREGATE}(f_1, f_2, \dots, f_t)$ of X .
- DECODE: A convolutional decoder which reconstructs the viewgrid $\hat{V}_t = \text{DECODE}(a_t)$ as a set of MN feature maps ($3MN$ for 3 channeled images) corresponding to each view of the viewgrid.
- ACT: Given the aggregated state a_t and proprioception p_t , the ACT module outputs a probability distribution $\pi(\delta|a_t)$ over the candidate camera motions $\delta \in \mathcal{A}$. An action sampled from this distribution $\delta_t = \text{ACT}(a_t, p_t)$ is executed.

At each time step, the agent receives and encodes a new view x_t , then updates its internal representation a_t by sensing, fusing, and aggregating. It decodes the

viewgrid \hat{V}_t and executes δ_t to change the viewpoint. It repeats the above steps until the time budget T is reached (see Fig. 2). See Supp. for implementation details and architecture diagram.

3.3 Sidekick Definitions

Sidekicks provide a preparatory learning phase that informs policy learning. Sidekicks have full observability during training; in particular, they can observe the results of arbitrary camera motions in arbitrary sequence. This is impossible for the actual look-around agent—who must enter novel environments and respect physical camera motion and budget constraints—but it *is* practical for the sidekick with fully observed training samples (e.g., a 360° panoramic image or 3D object model, cf. Sect. 4.1). Sidekicks are trained to solve a simpler problem with relevance to the ultimate look-around agent, serving to accelerate training and help the agent converge to better policies. In the following, we define two sidekick variants: a reward-based sidekick and a demonstration-based sidekick.

Reward-Based Sidekick. The reward-based sidekick aims to identify a set of K views $\{x(X, \theta_1), \dots, x(X, \theta_K)\}$ which can provide maximal information about the environment X . The sidekick is allowed to access X and select views without any restrictions. Hence, it addresses a simplified completion problem.

A candidate view is scored based on how informative it is, i.e., how well the entire environment can be reconstructed given only that view. We train a completion model (cf. Sect. 3.2) that can reconstruct $\hat{V}(X)$ from any single view (i.e., we set $T = 1$). Let $\hat{V}(X|y)$ denote the decoded reconstruction for X given only view y as input. The sidekick scores the information in observation $x(X, \theta)$ as:

$$\text{Info}(x(X, \theta), X) \propto^{-1} d\left(\hat{V}(X|x(X, \theta)), V(X)\right), \quad (1)$$

where d denotes the reconstruction error and $V(X)$ is the fully observed environment. We use a simple ℓ_2 loss on pixels for d to quantify information. Higher-level losses, e.g., for detected objects, could be employed when available. The scores are normalized to lie in $[0, 1]$ across the different views of X . The sidekick scores each candidate view. Then, in order to sharpen the effects of the scoring function and avoid favoring redundant observations, the sidekick selects the top K most informative views with greedy non-maximal suppression. It iteratively selects the view with the highest score and suppresses all views in the neighborhood of that view until K views are selected (see Supp. for details). This yields a map of favored views for each training environment. See Fig. 3, top row.

The sidekick conveys the results to the agent during policy learning in the form of an augmented reward (to be defined in Sect. 3.4). Thus, the reward-based sidekick previews observations and encourages the selection of those *individually* valuable for reconstruction. Note that while the sidekick indexes views in absolute angles, the agent will not; all its observations are relative to its initial (random) glimpse direction. This works because the sidekick becomes a part of

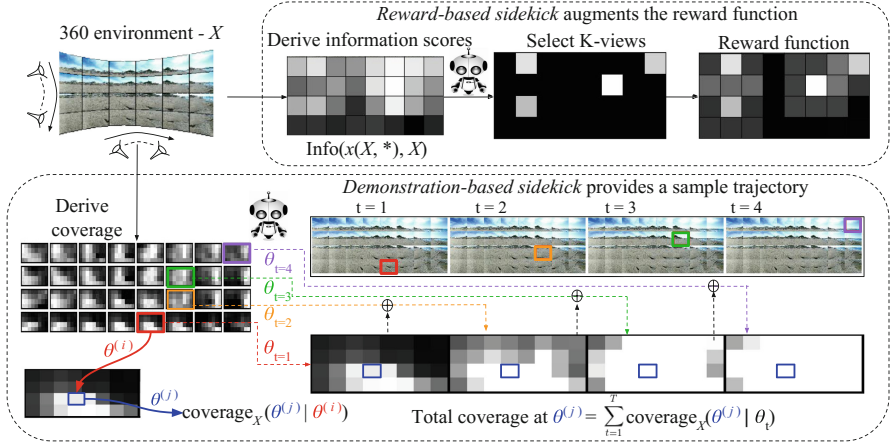


Fig. 3. Top left shows the 360° environment’s viewgrid, indexed by viewing elevation and azimuth. **Top: Reward sidekick** scores individual views based on how well they alone permit inference of the viewgrid X (Eq. 1). The grid of scores (center) is post-processed with non-max suppression to prioritize K non-redundant views (right), then is used to shape the agent’s rewards. **Bottom: Demonstration sidekick.** Left “grid-of-grids” displays example coverage score maps (Eq. 2) for all $\theta^{(i)}, \theta^{(j)}$ view pairs. The outer $N \times M$ grid considers each $\theta^{(i)}$, and each inner $N \times M$ grid considers each $\theta^{(j)}$ for the given $\theta^{(i)}$ (bottom left). A pixel in that grid is bright if coverage is high for $\theta^{(j)}$ given $\theta^{(i)}$, and dark otherwise. Each $\theta^{(i)}$ denotes an (elevation, azimuth) pair. While observed views and their neighbors are naturally recoverable (brighter), the sidekick uses broader environment context to also anticipate distant and/or different-looking parts of the environment, as seen by the non-uniform spread of scores in the left grid. Given the coverage function and a starting position, this sidekick selects actions to greedily optimize the coverage objective (Eq. 3). The bottom right strip shows the cumulative coverage maps as each of the $T = 4$ glimpses is selected.

the environment, i.e., it attaches rewards to the true views of the environment. In short, the reward-based sidekick shapes rewards based on its exploration with full observability.

Demonstration-Based Sidekick. Our second sidekick generates *trajectories* of informative views. Given a starting view in X , the demonstration sidekick selects a trajectory of T views that are deemed to be most informative about X . Unlike the reward-based sidekick above, this sidekick offers guidance with respect to a starting state, and it is subject to the same camera motion restrictions placed on the main agent. Such restrictions model how an agent cannot teleport its camera using one unit of effort.

To identify informative trajectories, we first define a scoring function that captures *coverage*. Coverage reflects how much information $x(X, \theta)$ contains about each view in X . The coverage score for view $\theta^{(j)}$ upon selecting view $\theta^{(i)}$ is:

$$\text{Coverage}_X \left(\theta^{(j)} | \theta^{(i)} \right) \propto^{-1} d \left(\hat{x}(X, \theta^{(j)}), x(X, \theta^{(j)}) \right), \quad (2)$$

where \hat{x} denotes an inferred view within $\hat{V}(X|x(X, \theta^{(i)}))$, as estimated using the same $T = 1$ completion network used by the reward-based sidekick. Coverage scores are normalized to lie in $[0, 1]$ for $1 \leq i, j \leq MN$.

$$\mathcal{C}(\Theta, X) = \sum_{j=1}^{MN} \sum_{\theta \in \Theta} \text{Coverage}_X(\theta^{(j)} | \theta), \quad (3)$$

The goal of the demonstration sidekick is to maximize the coverage objective (Eq. 3), where $\Theta = \{\theta_1, \dots, \theta_t\}$ denotes the sequence of selected views, and $\mathcal{C}(\Theta, X)$ saturates at 1. In other words, it seeks a sequence of reachable views such that *all* views are “explained” as well as possible. See Fig. 3, bottom panel.

The policy of the sidekick (π_s) is to greedily select actions based on the coverage objective. The objective encourages the sidekick to select views such that the overall information obtained about each view in X is maximized.

$$\pi_s(\Theta) = \arg \max_{\delta} \mathcal{C}(\Theta \cup \{\theta_t + \delta\}, X). \quad (4)$$

We use these sidekick-generated trajectories as supervision to the agent for a short preparatory period. The goal is to initialize the agent with useful insights learned by the sidekick to accelerate training of better policies. We achieve this through a hybrid training procedure that combines imitation and reinforcement. In particular, for the first t_{sup} time steps, we let the sidekick drive the action selection and train the policy based on a supervised objective. For steps t_{sup} to T , we let the agent’s policy drive the action selection and use REINFORCE [64] or Actor-Critic [59] to update the agent’s policy (see Sect. 4). We start with $t_{sup} = T$ and gradually reduce it to 0 in the preparatory sidekick phase (see Supp.). This step relates to behavior cloning [8, 14, 17], which formulates policy learning as supervised action classification given states. However, unlike typical behavior cloning, the sidekick is not an expert. It solves a simpler version of the task, then backs away as the agent takes over to train with partial observability.

3.4 Policy Learning with Sidekicks

Having defined the two sidekick variants, we now explain how they influence policy learning. The goal is to learn the policy $\pi(\delta|a_t)$ which returns a distribution over actions for the aggregated internal representation a_t at time t . Let $\mathcal{A} = \{\delta_i\}$ denote the set of camera motions available to the agent.

Our agent seeks the policy that minimizes reconstruction error for the environment given a budget of T camera motions (views). If we denote the set of weights of the network $[W_s, W_f, W_r, W_d, W_a]$ by W and W excluding W_a by $W_{/a}$ and W excluding W_d by $W_{/d}$, then the overall weight update is:

$$\Delta W = \frac{1}{n} \sum_{j=1}^n \lambda_r \Delta W_{/a}^{rec} + \lambda_p \Delta W_{/d}^{pol} \quad (5)$$

where n is the number of training samples, j indexes over the training samples, λ_r and λ_p are constants and $\Delta W_{/a}^{rec}$ and $\Delta W_{/d}^{pol}$ update all parameters except W_a and W_d , respectively. The pixel-wise MSE reconstruction loss (\mathcal{L}_t^{rec}) and corresponding weight update at time t are given in Eq. 6, where $\hat{x}_t(X, \theta^{(i)})$ denotes the reconstructed view at viewpoint $\theta^{(i)}$ and time t , and Δ_0 denotes the offset to account for the unknown starting azimuth (see [30]).

$$\begin{aligned} \mathcal{L}_{rec}^t(X) &= \sum_{i=1}^{MN} d\left(\hat{x}_t(X, \theta^{(i)} + \Delta_0), x(X, \theta^{(i)})\right), \\ \Delta W_{/a}^{rec} &= - \sum_{t=1}^T \nabla_{W_{/a}} \mathcal{L}_{rec}^t(X), \end{aligned} \quad (6)$$

The agent’s reward at time t (see Eq. 7) consists of the intrinsic reward from the sidekick $r_t^s = \text{Info}(x(X, \theta_t), X)$ (see Sect. 3.3) and the negated final reconstruction loss ($-\mathcal{L}_{rec}^T(X)$).

$$r_t = \begin{cases} r_t^s & 1 \leq t \leq T-2 \\ -\mathcal{L}_{rec}^T(X) + r_t^s & t = T-1 \end{cases} \quad (7)$$

The update from the policy (see Eq. 8) consists of the REINFORCE update, with a baseline b to reduce variance, and supervision from the demonstration sidekick (see Eq. 9). We consider both REINFORCE [64] and Actor-Critic [59] methods to update the ACT module. For the latter, the policy term additionally includes a loss to update a learned Value Network (see Supp.). For both, we include a standard entropy term to promote diversity in action selection and avoid converging too quickly to a suboptimal policy.

$$\Delta W_{/d}^{pol} = \sum_{t=1}^{T-1} \nabla_{W_{/d}} \log \pi(\delta_t | a_t) \left(\sum_{t'=t}^{T-1} r_{t'} - b(a_t) \right) + \Delta W_{/d}^{demo}, \quad (8)$$

The demonstration sidekick influences policy learning via a cross entropy loss between the sidekick’s policy π_s (cf. Sect. 3.3) and the agent’s policy π :

$$\Delta W_{/d}^{demo} = \sum_{t=1}^{T-1} \sum_{\delta \in \mathcal{A}} \nabla_{/d} (\pi_s(\delta | a_t) \log \pi(\delta | a_t)). \quad (9)$$

We pretrain the SENSE, FUSE, and DECODE modules with $T = 1$. The full network is then trained end-to-end (with SENSE and FUSE frozen). For training with sidekicks, the agent is augmented either with additional rewards from the reward sidekick (Eq. 7) or an additional supervised loss from the demonstration sidekick (Eq. 9). As we will show empirically, training with sidekicks helps overcome uncertainty due to partial observability and learn better policies.

3.5 Visualizing the Learned Motion Policies

Finally, we propose a visualization technique to qualitatively understand the policy that has been learned. The aggregated state a_t is used by the policy network to determine the action probabilities. To analyze which part of the agent’s belief (a_t) is important for the current selected action δ_t , we solve for the change in the aggregated state (Δa_t) which maximizes the change in the predicted action distribution ($\pi(\cdot|a_t)$):

$$\Delta a^* = \arg \max_{\Delta a_t} \sum_{\delta \in \mathcal{A}} (\pi(\delta|a_t) - \pi(\delta|a_t + \Delta a_t))^2 \quad (10)$$

s.t. $\|\Delta a_t\| \leq C\|a_t\|$

where C is a constant that limits the deviation in norm from the true belief. Equation 10 is maximized using gradient ascent (see Supp.). This change in belief is visualized in the viewgrid space by forward propagating through the DECODE module. The visualized heatmap intensities (H_t) are defined as follows:

$$H_t \propto \|\text{DECODE}(a_t + \Delta a^*) - \text{DECODE}(a_t)\|_2^2. \quad (11)$$

The heatmap indicates which parts of the agent’s belief *would have to change* to affect its action selection. The views with high intensity are those that affect the agent’s action selection the most.

4 Experiments

In Sects. 4.1 and 4.2, we describe our experimental setup and analyze the learning efficiency and test-time performance of different methods. In Sect. 4.3, we visualize learned policies and demonstrate the superiority of our policies over a baseline.

4.1 Experimental Setup

Datasets: We use two popular datasets to benchmark our models.

- **SUN360:** SUN360 [66] consists of high resolution spherical panoramas from multiple scene categories. We restrict our experiments to the 26 category subset used in [30, 66]. The viewgrid consists of 32×32 views captured across 4 elevations (-45° to 45°) and 8 azimuths (0° to 180°). At each step, the agent sees a 60° field-of-view. This dataset represents an agent looking out at a scene in a series of narrow field-of-view glimpses.
- **ModelNet Hard:** ModelNet [65] provides a collection of 3D CAD models for different categories of objects. ModelNet-40 and ModelNet-10 are provided subsets consisting of 40 and 10 object categories respectively, the latter being a subset of the former. We train on objects from the 30 categories not present in ModelNet-10 and test on objects from the unseen 10 categories. We increase

completion difficulty in “ModelNet Hard” by rendering with more challenging lighting conditions, textures and viewing angles than [30]; see Supp. It consists of 32×32 views sampled from 5 elevations and 9 azimuths. This dataset represents an agent looking in at a 3D object and moving it to a series of selected poses.

For both datasets, the candidate motions \mathcal{A} are restricted to a 3 elevations \times 5 azimuths neighborhood, representing the set of unit-cost actions. Neighborhood actions mimic real-world scenarios where the agent’s physical motions are constrained (i.e., no teleporting) and is consistent with recent active vision work [2, 28–30, 43]. The budget for number of steps is fixed to $T = 4$.

Baselines: We benchmark our methods against several baselines:

- **one-view:** the agent trained to reconstruct from one view ($T = 1$).
- **rnd-actions:** samples actions uniformly at random.
- **1t1a** [30]: our implementation of the “learning to look around” approach [30]. We verified our code reproduces results from [30].
- **rnd-rewards:** naive sidekick where rewards are assigned uniformly at random on the viewgrid.
- **asymm-ac** [48]: approach from [48] adapted for discrete actions. Critic sees the entire panorama/object and true camera poses (no experience replay).
- **demo-actions:** actions selected by demo-sidekick while training/testing.
- **expert-clone:** imitation from an expert policy that uses full observability (similar to critic in Fig. 2 of Supp.)

Evaluation: We evaluate reconstruction error averaged over uniformly sampled elevations, azimuths and all test samples (**avg**). To provide a worst case analysis, we also report an adversarial metric (**adv**), which evaluates each agent on its hardest starting positions in each test sample and averages over the test data.

4.2 Active Exploration Results

Table 1 shows the results on both datasets. For each metric, we report the mean error along with the percentage improvement over the **one-view** baseline. Our methods are abbreviated **ours(rew)** and **ours(demo)** referring to the use of our reward- and demonstration-based sidekicks, respectively. We denote the use of Actor-Critic instead of REINFORCE with **+ac**.

We observe that **ours(rew)** and **ours(demo)** with REINFORCE generally perform better than **1t1a** with REINFORCE [30]. In particular, **ours(rew)** performs significantly better than **1t1a** on both datasets on all metrics. **ours(demo)** performs better on SUN360, but is only slightly better on ModelNet Hard. Figure 4 shows the validation loss plots; using the sidekicks leads to significant improvement in the convergence rate over **1t1a**.

Figure 5 compares example decoded reconstructions. We stress that the vast majority of pixels are unobserved when decoding the belief state, i.e., only 4 views out of the entire viewing sphere are observed. Accordingly, they are blurry.

Table 1. Avg/Adv MSE errors $\times 1000$ (\downarrow lower is better) and corresponding improvements (%) over the **one-view** model (\uparrow higher is better), for the two datasets. The best and second best performing models are highlighted in green and blue respectively. Standard errors range from 0.2 to 0.3 on SUN360 and 0.1 to 0.2 on ModelNet Hard.

Method	SUN360				ModelNet Hard			
	Avg ($\times 1000$)		Adv ($\times 1000$)		Avg ($\times 1000$)		Adv ($\times 1000$)	
	Mean \downarrow	% \uparrow	Mean \downarrow	% \uparrow	Mean \downarrow	% \uparrow	Mean \downarrow	% \uparrow
one-view	38.31	-	55.12	-	9.63	-	17.10	-
rnd-actions	30.99	19.09	44.85	18.63	7.32	23.93	12.38	27.56
rnd-rewards	25.55	33.30	30.20	45.21	7.04	26.89	9.66	43.50
ltla [30]	24.94	34.89	31.86	42.19	6.30	34.57	8.78	48.65
asymm-ac [48]	23.74	38.01	29.92	45.72	6.24	35.20	8.55	50.00
expert-clone	23.98	37.38	28.50	48.28	6.41	33.44	8.52	50.13
ours(rew)	23.44	38.82	28.54	48.22	5.80	39.79	7.17	58.04
ours(demo)	24.24	36.73	29.01	47.36	6.32	34.37	8.64	49.47
ours(rew)+ac	23.36	39.01	28.26	48.72	5.75	40.26	7.10	58.44
ours(demo)+ac	24.05	37.22	28.52	48.26	6.13	36.31	8.26	51.64
demo-actions *	26.12	31.82	31.53	42.76	5.82	39.50	7.46	56.40

(* - requires full observability at test time)

Regardless, their differences indicate the differences in belief states between the two methods. A better policy more quickly fleshes out the general shape of the scene or object.

Next, we compare our model to **asymm-ac**, which is an alternate paradigm for exploiting full observability during training. First, we note that **asymm-ac** performs better than **ltla** across all datasets and metrics, making it a strong baseline. Comparing **asymm-ac** with **ours(rew)+ac** and **ours(demo)+ac**, we see our methods still perform considerably better on all metrics and datasets. As we show in the Supp, our methods also lead to faster convergence.

In order to contrast learning from sidekicks with learning from experts, we additionally compare our models to behavior cloning an expert that exploits full observability at training time. As shown in Table 1, **ours(rew)** outperforms **expert-clone** on both the datasets, validating the strength of our approach. It is particularly interesting because training an expert takes a lot longer ($17\times$) than training sidekicks (see Supp.). When compared with **demo-actions**, an ablated version of **ours(demo)** that requires full observability at *test time*, our performance is still significantly better on SUN360 and slightly better on ModelNet Hard. **ours(rew)** and **ours(demo)** also beat the remaining baselines by a significant margin. These results verify our hypothesis that sidekick policy learning can improve over strong baselines by exploiting full observability during training.




Fig. 4. Validation errors ($\times 1000$) vs. epochs on SUN360 (left) and ModelNet Hard (right). All models shown here use REINFORCE (see Supp. for more curves). Our approach accelerates convergence.

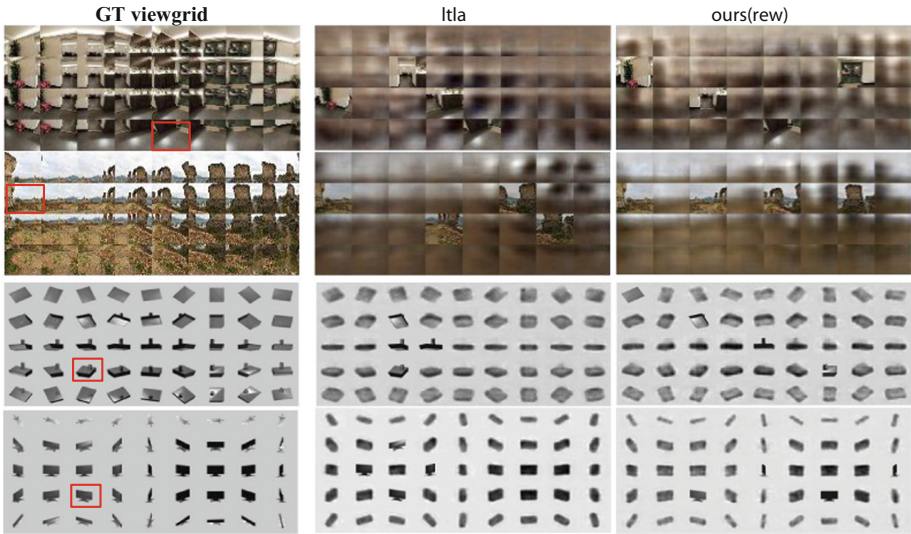


Fig. 5. Qualitative comparison of `ours(rew)` vs. `l1la` [30] on SUN360 (first 2 rows) and ModelNet Hard (last 2 rows). The first column shows the groundtruth viewgrid and a randomly selected starting point (marked in red). The 2nd and 3rd columns contain the decoded viewgrids from `l1la` and `ours(rew)` after $T = 4$ time steps. The reconstructions from `ours(rew)` are visibly better. For example, in the 3rd row, our model reconstructs the protrusion more clearly; in the 2nd row, our model reconstructs the sky and central hills more effectively. Best viewed on pdf with zoom. (Color figure online)

4.3 Policy Visualization

We present our policy visualizations for `l1la` and `ours(rew)` on SUN360 in Fig. 6; see Supp. for examples with `ours(demo)`. The heatmap from Eq. 10 is shown in pink and overlaid on the reconstructed viewgrids. For both models,




Fig. 6. Policy visualization: The viewgrid reconstructions of `ours(rew)` and `ltla` [30] are shown on two examples from SUN360. The first column shows the viewgrid with a randomly selected view (in red). Subsequent columns show the view received (in red), viewgrid reconstructed, action selected (red arrow), and the parts of the belief space our method deems responsible for the action selection (pink heatmap). Both the agents tend to move towards sparser regions of the heatmap, attempting to improve their beliefs about views that do not contribute to their action selection. `ours(rew)` improves its beliefs much more rapidly and as a result, performs more informed action selection. (Color figure online)

the policies tend to take actions that move them towards views which have low heatmap density, as witnessed by the arrows/actions pointing to lower density regions. Intuitively, the agents move towards the views that are not contributing effectively to their action selection to increase their understanding of the scene. It can be observed in many cases that `ours(rew)` model has a much denser heatmap across time when compared to `ltla`. Therefore, `ours(rew)` takes more views into account for selecting its actions earlier in the trajectory, suggesting that a better policy and history aggregation leads to more informed action selection.

5 Conclusion

We propose *sidekick policy learning*, a framework to leverage extra observability or fewer restrictions on an agent’s motion during training to learn better policies. We demonstrate the superiority of policies learned with sidekicks on two challenging datasets, improving over existing methods and accelerating training. Further, we utilize a novel policy visualization technique to illuminate the different reasoning behind policies trained with and without sidekicks. In future work, we plan to investigate the effectiveness of our framework on other active vision tasks such as recognition and navigation.

Acknowledgements. The authors thank Dinesh Jayaraman, Thomas Crosley, Yu-Chuan Su, and Ishan Durugkar for helpful discussions. This research is supported in part by DARPA Lifelong Learning Machines, a Sony Research Award, and an IBM Open Collaborative Research Award.

References

1. Aloimonos, J., Weiss, I., Bandyopadhyay, A.: Active vision. *Int. J. Comput. Vis.* **1**, 333–356 (1988)
2. Ammirato, P., Poirson, P., Park, E., Košecká, J., Berg, A.C.: A dataset for developing and benchmarking active vision. In: 2017 IEEE International Conference on Robotics and Automation (2017)
3. Anthony, T., Tian, Z., Barber, D.: Thinking fast and slow with deep learning and tree search. In: *Advances in Neural Information Processing Systems* (2017)
4. Ba, J., Mnih, V., Kavukcuoglu, K.: Multiple object recognition with visual attention. arXiv preprint [arXiv:1412.7755](https://arxiv.org/abs/1412.7755) (2014)
5. Bajcsy, R.: Active perception. In: *IEEE Proceedings*, vol. 76, no. 8, pp. 996–1006 (1988)
6. Ballard, D.H.: Animate vision. *Artif. Intell.* **48**(1), 57–86 (1991)
7. Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., Munos, R.: Unifying count-based exploration and intrinsic motivation. In: *Advances in Neural Information Processing Systems* (2016)
8. Bojarski, M., et al.: End to end learning for self-driving cars. arXiv preprint [arXiv:1604.07316](https://arxiv.org/abs/1604.07316) (2016)
9. Caicedo, J.C., Lazebnik, S.: Active object localization with deep reinforcement learning. In: 2015 IEEE International Conference on Computer Vision (2015)
10. Das, A., Datta, S., Gkioxari, G., Lee, S., Parikh, D., Batra, D.: Embodied Question Answering. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition (2018)
11. Dhillon, S.S., Chakrabarty, K.: Sensor placement for effective coverage and surveillance in distributed sensor networks. In: 2003 Wireless Communications and Networking. WCNC 2003. IEEE (2003)
12. Dosovitskiy, A., Koltun, V.: Learning to act by predicting the future. In: International Conference on Learning Representations (2017)
13. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: an open urban driving simulator. In: *Conference on Robot Learning* (2017)
14. Duan, Y., et al.: One-shot imitation learning. In: *Advances in Neural Information Processing Systems* (2017)
15. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: 2017 IEEE International Conference on Computer Vision (2017)
16. Garg, R., Vijay Kumar, B.G., Carneiro, G., Reid, I.: Unsupervised CNN for single view depth estimation: geometry to the rescue. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9912, pp. 740–756. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_45
17. Giusti, A., et al.: A machine learning approach to visual perception of forest trails for mobile robots. *IEEE Robot. Autom. Lett.* **1**, 661–667 (2016)
18. Greydanus, S., Koul, A., Dodge, J., Fern, A.: Visualizing and understanding atari agents. *CoRR* (2017)

19. Guo, X., Singh, S., Lee, H., Lewis, R.L., Wang, X.: Deep learning for real-time Atari game play using offline Monte-Carlo tree search planning. In: *Advances in Neural Information Processing Systems* (2014)
20. Gupta, S., Davidson, J., Levine, S., Sukthankar, R., Malik, J.: Cognitive mapping and planning for visual navigation. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition* (2017)
21. Gupta, S., Fouhey, D., Levine, S., Malik, J.: Unifying map and landmark based representations for visual navigation. *arXiv preprint [arXiv:1712.08125](https://arxiv.org/abs/1712.08125)* (2017)
22. Gupta, S., Hoffman, J., Malik, J.: Cross modal distillation for supervision transfer. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition* (2016)
23. Harutyunyan, A., Devlin, S., Vrancx, P., Nowe, A.: Expressing arbitrary reward functions as potential-based advice. In: *Twenty-Ninth AAAI Conference on Artificial Intelligence* (2015)
24. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arXiv preprint [arXiv:1503.02531](https://arxiv.org/abs/1503.02531)* (2015)
25. Hong, S., Noh, H., Han, B.: Decoupled deep neural network for semi-supervised semantic segmentation. In: *Advances in Neural Information Processing Systems* (2015)
26. Hong, S., Oh, J., Lee, H., Han, B.: Learning transferrable knowledge for semantic segmentation with deep convolutional neural network. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition* (2016)
27. Jaderberg, M., et al.: Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint [arXiv:1611.05397](https://arxiv.org/abs/1611.05397)* (2016)
28. Jayaraman, D., Grauman, K.: End-to-end policy learning for active visual categorization. *IEEE Trans. Pattern Anal. Mach. Intell.* (2018). <https://doi.org/10.1109/TPAMI.2018.2840991>
29. Jayaraman, D., Grauman, K.: Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016. LNCS, vol. 9909*, pp. 489–505. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46454-1_30
30. Jayaraman, D., Grauman, K.: Learning to look around: intelligently exploring unseen environments for unknown tasks. In: *2018 IEEE Conference on Computer Vision and Pattern Recognition* (2018)
31. Johns, E., Leutenegger, S., Davison, A.J.: Pairwise decomposition of image sequences for active multi-view recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition* (2016)
32. Kim, A., Eustice, R.M.: Perception-driven navigation: active visual slam for robotic area coverage. In: *2013 IEEE International Conference on Robotics and Automation* (2013)
33. Knox, W.B., Stone, P.: Interactively shaping agents via human reinforcement: the tamer framework. In: *Proceedings of the Fifth International Conference on Knowledge Capture* (2009)
34. Knox, W.B., Stone, P.: Combining manual feedback with subsequent MDP reward signals for reinforcement learning. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems* (2010)
35. Knox, W.B., Stone, P.: Reinforcement learning from simultaneous human and MDP reward. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems* (2012)
36. Krause, A., Guestrin, C.: Near-optimal observation selection using submodular functions. In: *AAAI* (2007)

37. Lapin, M., Hein, M., Schiele, B.: Learning using privileged information: SVM+ and weighted SVM. *Neural Netw.* **53**, 95–108 (2014)
38. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.* **17**, 1334–1373 (2016)
39. Levine, S., Koltun, V.: Guided policy search. In: *International Conference on Machine Learning* (2013)
40. Levine, S., Pastor, P., Krizhevsky, A., Quillen, D.: Learning hand-eye coordination for robotic grasping with large-scale data collection. In: Kulić, D., Nakamura, Y., Khatib, O., Venture, G. (eds.) *2016 International Symposium on Experimental Robotics* (2017)
41. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
42. Liu, T., et al.: Learning to detect a salient object. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**, 353–367 (2011)
43. Malmir, M., Sikka, K., Forster, D., Movellan, J.R., Cottrell, G.: Deep Q-learning for active recognition of germs: baseline performance on a standardized dataset for active learning. In: *British Machine Vision Conference* (2015)
44. Mirowski, P., et al.: Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673* (2016)
45. Mnih, V., Heess, N., Graves, A., et al.: Recurrent models of visual attention. In: *Advances in Neural Information Processing Systems* (2014)
46. Nair, A., et al.: Combining self-supervised learning and imitation for vision-based rope manipulation. In: *2017 IEEE International Conference on Robotics and Automation* (2017)
47. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: *International Conference on Machine Learning* (2017)
48. Pinto, L., Andrychowicz, M., Welinder, P., Zaremba, W., Abbeel, P.: Asymmetric actor critic for image-based robot learning. *Robot. Sci. Syst.* (2018)
49. Pinto, L., Gupta, A.: Supersizing self-supervision: learning to grasp from 50k tries and 700 robot hours. In: *2016 IEEE International Conference on Robotics and Automation* (2016)
50. Ross, S., Gordon, G., Bagnell, D.: A reduction of imitation learning and structured prediction to no-regret online learning. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (2011)
51. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015)
52. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization. In: *2017 IEEE International Conference on Computer Vision* (2017)
53. Sharmanska, V., Quadrianto, N., Lampert, C.H.: Learning to rank using privileged information. In: *2013 IEEE International Conference on Computer Vision*. IEEE (2013)
54. Silver, D., et al.: Mastering the game of go without human knowledge. *Nature* **550**, 354 (2017)
55. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013)

56. Song, S., Zeng, A., Chang, A.X., Savva, M., Savarese, S., Funkhouser, T.: Im2Pano3D: extrapolating 360 structure and semantics beyond the field of view. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition (2018)
57. Soomro, K., Zamir, A.R., Shah, M.: UCF101: a dataset of 101 human actions classes from videos in the wild. arXiv preprint [arXiv:1212.0402](https://arxiv.org/abs/1212.0402) (2012)
58. Spica, R., Giordano, P.R., Chaumette, F.: Active structure from motion: application to point, sphere, and cylinder. *IEEE Trans. Robot.* **30**, 1499–1513 (2014)
59. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. <https://mitpress.mit.edu/books/reinforcement-learning>
60. Tulsiani, S., Zhou, T., Efros, A.A., Malik, J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (2017)
61. Vapnik, V., Izmailov, R.: Learning with intelligent teacher. In: Gammerman, A., Luo, Z., Vega, J., Vovk, V. (eds.) COPA 2016. LNCS (LNAI), vol. 9653, pp. 3–19. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-33395-3_1
62. Wang, B.: Coverage problems in sensor networks: a survey. *ACM Comput. Surv.* **43**, 32 (2011)
63. Wilkes, D., Tsotsos, J.K.: Active object recognition. In: 1992 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (1992)
64. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. In: Sutton, R.S. (ed.) Reinforcement Learning, vol. 173, pp. 5–32. Springer, Boston (1992). https://doi.org/10.1007/978-1-4615-3618-5_2
65. Wu, Z., et al.: 3D shapeNets: a deep representation for volumetric shapes. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (2015)
66. Xiao, J., Ehinger, K.A., Oliva, A., Torralba, A.: Recognizing scene viewpoint using panoramic place representation. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (2012)
67. Xu, K., et al.: Show, attend and tell: neural image caption generation with visual attention. In: International Conference on Machine Learning (2015)
68. Yang, C., Zhang, L., Lu, H., Ruan, X., Yang, M.H.: Saliency detection via graph-based manifold ranking. In: 2013 IEEE Conference on Computer Vision and Pattern Recognition (2013)
69. Zahavy, T., Ben-Zrihem, N., Mannor, S.: Graying the black box: understanding DQNs. In: International Conference on Machine Learning (2016)
70. Zhu, Y., et al.: Visual semantic planning using deep successor representations. In: 2017 IEEE International Conference on Computer Vision (2017)
71. Zhu, Y., et al.: Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: 2017 IEEE International Conference on Robotics and Automation (2017)