

# A Dataset for Developing and Benchmarking Active Vision

Phil Ammirato<sup>1</sup>, Patrick Poirson<sup>1</sup>, Eunbyung Park<sup>1</sup>, Jana Košecká<sup>2</sup>, Alexander C. Berg<sup>1</sup>

**Abstract**—We present a new public dataset with a focus on simulating robotic vision tasks in everyday indoor environments using real imagery. The dataset includes 20,000+ RGB-D images and 50,000+ 2D bounding boxes of object instances densely captured in 9 unique scenes. We train a fast object category detector for instance detection on our data. Using the dataset we show that, although increasingly accurate and fast, the state of the art for object detection is still severely impacted by object scale, occlusion, and viewing direction all of which matter for robotics applications. We next validate the dataset for simulating active vision, and use the dataset to develop and evaluate a deep-network-based system for next best move prediction for object classification using reinforcement learning. Our dataset is available for download at [cs.unc.edu/~ammirato/active\\_vision\\_dataset\\_website/](http://cs.unc.edu/~ammirato/active_vision_dataset_website/).

## I. INTRODUCTION

The ability to recognize objects is a core functionality for robots operating in everyday human environments. While there has been amazing recent progress in computer vision on object classification and detection, especially with deep models, these lines of work do not address some of the core needs of vision for robotics. Partly this is due to biases in the imagery considered and the fact that these recognition challenges are performed in isolation for each image. In robotic applications, the biases are different and recognition is performed over multiple images, often with active control of the sensing platform (active vision). This paper attempts to address part of this disconnect by introducing a new approach to studying active vision for robotics by collecting very dense imagery of scenes in order to allow simulating a robot moving through an environment by sampling appropriate imagery.

The goals are two-fold, to *provide a research and development resource for computer vision* without requiring access to robots for experiments, and to *provide a way to benchmark and compare different approaches to active vision* without the difficulty and expense of evaluating the algorithms on the same physical robotics testbed. We begin by collecting a large dataset of dense RGB-D imagery of common everyday rooms: kitchens, living rooms, dining rooms, offices, etc. This imagery is registered and used to form a 3D reconstruction of each scene. This reconstruction is used to simplify labeling of objects in the collection in 3D as opposed to individually in the thousands of images of those objects. The geometric relationship between images is

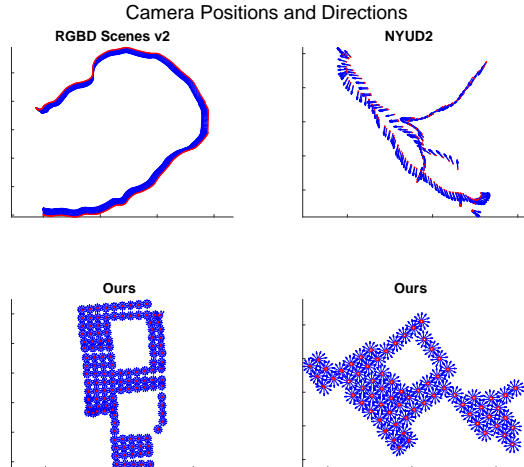


Fig. 1. Visualization of camera locations (red) and viewing directions (blue) from our collections (bottom) and previous datasets (top). We collect densely sampled RGB-D images of scenes for use in training and benchmarking active vision systems. The dense sampling allows “virtually” moving a camera through a scene. While other datasets do sample multiple images per scene, they often sample either from just a few positions or along only a few paths through the environment [1], [2]. Note that the physical scale is different in each plot.

also used to define connectivity for determining what image would be seen next when moving in a given direction from a given camera position (e.g. what would I see if I turned right? went backwards?).

Given this labeled data we adapt a state-of-the-art fast object category detector [3] based on deep convolutional networks to the task of recognizing specific object instances in the dataset. While most deep-learning approaches have focused on category detection, instance detection can be practically useful for robotics. This distinction between recognizing a category of object, such as chair, versus a specific object, such as a particular 8.4oz Red Bull can is important. Our results show that the category detection framework can be adapted to instance detection well, with some caveats.

Where the detection framework has difficulty is in the range of scales, viewing directions, and occlusions present in everyday scenes (e.g. our data) that is different from the biases present in Internet collected datasets. While the detector performs well for large frontal views of objects its performance falls for other views. This is quantified in Sec. IV-A. This view-dependent variation in recognition performance motivates active-vision for object recognition, controlling the sensing platform to acquire imagery that

\*Supported by NSF NRI 1527208,1526367 & NSF 1452851 & 1446631

<sup>1</sup>University of North Carolina at Chapel Hill, Computer Science [ammirato, poirson, eunbyung, aberg]@cs.unc.edu

<sup>2</sup>George Mason University, Computer Science kosecka@gmu.edu

improves recognition accuracy.

Our high-level goals are based on using the pre-collected dense imagery to develop and test active-vision algorithms. To validate this approach we begin by demonstrating that the imagery is sampled densely enough. In particular we care that the results and accuracy of recognition algorithms on samples of the densely collected imagery are close to the results that would be achieved if the robot moved continuously through the environment. This is explored in Sec. IV-C.

Given this validation, we proceed to use the densely sampled dataset to train and evaluate a deep-network for determining the next best move to improve object classification. The recognition component for this is pre-trained with external data and then a combined network that performs recognition and selects a direction to move in to improve accuracy is trained on a subset of the densely sampled data using reinforcement learning. To illustrate one way to use the dataset, we employ multiple train/test splits to determine the expected increase in accuracy with multiple moves using our next best move network. See Sec. IV-D.

The collected dataset and labels are available at [http://cs.unc.edu/~ammirato/active\\_vision\\_dataset\\_website/](http://cs.unc.edu/~ammirato/active_vision_dataset_website/), as well as a small toolbox for visualizations and loading. We hope to also provide the functionality to allow groups to submit algorithms for evaluation on completely private test data in the future. Before collection of imagery, release forms were signed and collected allowing free and legal access to the collected data.

## II. RELATED WORK

This paper proposes an approach to collecting and using datasets to train and benchmark object detection and recognition, especially for *active recognition*. We briefly discuss some of the most related work in each area.

The datasets that have been a driving force in pushing the deep learning revolution in object recognition, Pascal VOC [4], the ImageNet Challenge [5], and MS COCO [6] are all collected from web images (usually from Flickr) using web search based on keywords. These image collections introduce biases from the human photographer, the human tagging, and the web search engine. As a result objects are usually of medium to large size in images and are usually frontal views with small amounts of occlusion. In addition these datasets focus on object category recognition. The state of the art for object classification and recognition in these datasets is based on either object proposals and feature pooling following [7] with advanced deep networks [8], [9] or on fully convolutional networks implementing a modern take on sliding windows [3], [10], [11] that provide frame-rate or faster performance on high-end hardware for some reduction in accuracy.

Instance recognition (as opposed to object category recognition) has generally been approached using local features or template matching techniques. A recent relevant example using these types of models is [12] that trains on objects

in a room and is tested on the same objects in the room after rearrangement. In our experiments we are interested in generalization to new environments in order to avoid training in each new room. More recently, [13] shows how deep-learning for comparing instances can be applied to instance classification and outperform classic matching methods. For our data, we are also interested in instance detection, including localization in a large image. We use the system from [3] to build a much faster detector for object instances than would be possible with explicit matching.

There are many RGB-D datasets available today, but none with a focus on simulating robot motion through an environment. [14] gives a list of a various RGB-D datasets, some focus on single objects [15], [16], in what we call “table-top” style data. This type of data, especially the data in BigBIRD [15], is similar to what manufacturers may provide for robots in the future. While not capturing real-world scenes, the number of views and detail for each instance in this data can provide valuable training data for instance recognition systems. We include over 30 object instances similar to those in the BigBIRD dataset in our scenes.

Scene dataset, [2], [17],[1], and [18] do explore environments more than “table-top” data but do not have a dense set of views to simulate robot motion. These data-sets often have only one or two paths through the scene. An actual robot in the real-world has many choices of where to move, and the controller has to be able to pick a good path. See Figure 1 for a comparison of the available paths through scenes in previous datasets and our data.

Active vision has a long history in robotics. Early work largely centered around view selection [19], [20]. Others [21], [22], [23] have worked on the problem from a more theoretical perspective, but under many simplified settings for possible motions, or assumptions about known object models. In recent years, next best view prediction has been one of the more popular active vision problems. However, most of these approaches use CAD models of the objects of interest [24], [25], [26], with some small sets of real-world images [27]. CAD models produce encouraging results, but leave out some real-world challenges in perception.

[27] gives a system for object detection, pose estimation, and next best view prediction. They are able to test their detection and pose estimation system on existing real image datasets, but need to collect their own data to test their active vision framework. They collect a small scale dataset of only “table-top” style scenes with about 30-60 images each. This shows the need for a dataset for active vision, while also showing how difficult it can be to collect such data at a large scale.

## III. DATA COLLECTION

Our dataset covers a variety of scenes from office buildings and homes, often capturing more than one room. For example a kitchen, living room, and dining room may all be present in one scene. We capture a total of 9 unique scenes, but have a total of 17 scans since some scenes are scanned twice. Each scene has from 696-2,412 images, for a total of 20,916




Fig. 2. Four dense reconstructions of scenes from our collected data. We label objects in 3D using the dense reconstructions then project to each camera image to obtain 2D bounding boxes. (Reconstruction tool from [28].)

images and 54,247 bounding boxes. We use the Kinect v2 sensor and code from [29] for collection.

As stated, we aim to be able to simulate robotic motion through each scene with our scans. At first it may seem the best way to do this is to capture video as the camera moves around the scene. However, in order to get more than one view at any given point the camera must be rotated at that point. It is not possible to visit the infinite number of points in each scene, so a discrete set of points must be chosen. In a video, even if a consistent frame rate and rotation speed are maintained, there will be images in between the points of rotation that still represent only a single view of a position in the scene. This is unnatural for movement. Imagine a robot arriving at a location and being unable to turn in place.

We choose to have the camera visit a set of discrete points throughout the scene in order to provide some consistency among the images and camera positions. A video could still be collected at each point of rotation, but this would increase the dataset size unnecessarily. We choose to sample every 30 degrees at each point of rotation, providing substantial overlap between images while keeping the number of images in each scene manageable.

The set of points our robot visits in each scene is essentially a rectangular grid over the scene. We make our points 30 centimeters apart, and justify this in later experiments. Our scenes have between 58-201 points, which allow many choices of how to move.

Two scans of a scene will have different placements of objects. Only objects that would be naturally moved in daily life are relocated. For example chairs, books, and BigBIRD objects may be moved, but sofas and refrigerators will stay put. There are two advantages to scanning each scene twice. First, we are able to get more data from each scene, which is important given the limited availability of scenes. Second, we can test a system that learns about objects and a scene from an initial scan, and then is tested on the same scene with moved or new objects, e.g. [12].




Fig. 3. A comparison between an initial depth image (left) and the improved depth image (right). The improved depth images allow us to better handle occlusion when projecting point cloud labels from the dense reconstruction to bounding boxes in the RGB images.

### A. Labels

We aim to collect 2D bounding boxes of our 33 common instances across all scenes. In addition, we need to provide movement pointers from each image to allow movement through the scene. We provide pointers for rotation clockwise and counter clockwise, as well as translation forward, backward, left, and right.

For each scan of each scene, we create a sparse reconstruction of the scene using the RGB structure from motion tool COLMAP from Schönberger et al [30], [31]. From the reconstruction we get the camera position and orientation for each image. We don't use depth information for the reconstruction because our sampling is so dense that we are rarely testing the limits of the RGB system. See Figure 1 for example reconstructed camera positions.

Using the camera positions and orientations we are able to calculate the movement pointers that allow navigation through each scene using natural robotic movements.

To label every object instance in each scan, we feed the output of COLMAP into the dense reconstruction system CMVS/PMVS [28], [32]. This gives us a denser point cloud of the scene that makes it easy for humans to recognize objects. We then extract the point cloud of each instance from this dense reconstruction, and are able to get 2D bounding boxes in every image by projecting the point clouds for each object into each image. See Figure 2. Given that most of our scans include multiple rooms and lots of clutter, we must account for occlusion or the point clouds will project through walls and occluding objects and give low quality 2D bounding boxes. We are able to use the Kinect depth maps with the reconstructed point clouds and camera poses to account for some occlusion, but not all. Some occlusion is missed by the raw depth maps because they are sometimes noisy, giving wrong or no values for reflectivity shiny surfaces, and are not at the same resolution as the RGB images.

To improve a given depth map  $D$ , we build a dense reconstruction by back projecting the depth maps of cameras that see similar areas of the scene. This solves the resolution problem, as the other depth maps cover the areas missed by  $D$ . We are also able to fill in many of the missing or wrong values on specular surfaces by taking advantage of the fact that these values are either zero, or much greater than the true depth. Each depth image has a slightly different

view of the specular surface, and so has various correct and incorrect values on that surface. By projecting the point clouds of many depth images into  $D$  and keeping the smallest value for each pixel, we are able to remove most of the wrong values that are too large, and fill in a lot of the missing values. As a last step we perform some simple interpolation to attempt to fill in any holes of missing values that are left. See Figure 3 for a comparison of original to improved depth maps.

Though the improved depths are much better they are still not perfect. There is also noise in the dense reconstruction and noise in the labeled point clouds. Knowing this, we inspect every bounding box ourselves to make sure it contains the correct object, and is not of poor quality (too large or small for the object). We have labeled our scans for BigBIRD objects, yielding an average of over 3000 2D bounding boxes per scan. We provide some measure of difficulty for each bounding box based on its size, leaving adding a measure of occlusion for future work. For our experiments we only consider boxes with a size of at least 50x30 pixels.

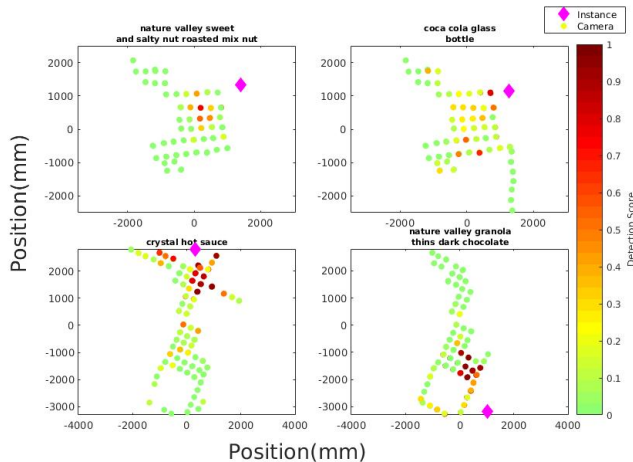


Fig. 4. Detection scores for four different instances in various scenes. Dots are camera position, color indicates score. Only cameras that see the instance (purple diamond) are shown. Notice certain viewpoints consistently yield higher scores. It would be advantageous for a robot to move from green views to red ones.

#### IV. EXPERIMENTS

We aim to show four things: a baseline for instance detection on our data, why it is important to design systems specifically for robot motion, how our dataset can be used to simulate motion, and a system demonstrating an active vision task on our dataset.

##### A. Instance Detection

We use a state-of-the-art class level object detector as a baseline for instance detection on our dataset. We choose the Single Shot Detection (SSD) network from [3] because it offers both real time detection performance (72 FPS) while maintaining a high-level of accuracy. This is exciting for robotics applications for which real time performance is

Instance	Split 1	Split 2	Split 3
Boxes > 100x75	.39	.55	.53
Boxes > 50x30	.26	.41	.42

TABLE I

MAP DETECTION RESULTS. SINCE SMALL BOXES ARE CHALLENGING FOR DETECTION SYSTEMS TO REPRODUCE, WE TRAIN/TEST OUR DETECTOR FIRST USING ONLY BOXES OF SIZE AT LEAST 100x75, AND THEN RE-TRAIN/TEST ON ALL BOXES AT LEAST 50x30.

crucial. The SSD network consists of a base network, in our case VGG [33], with additional feature maps added on top of the base network through a series of 1x1 and 3x3 convolutions.

We separate our dataset into three training and testing splits. Each split consists of eleven scans from seven scenes as training and three scans from two scenes for testing. Since small objects present a particularly difficult challenge for our detector, we first only consider boxes of size at least 100x75 pixels for training and testing. We then include all boxes of size at least 50x30, adding more training data but also a more difficult test scenario.

We use 500x500 images for training SSD. We train the network using an initial learning rate of 0.001 and train the network for 20,000 iterations with a stepsize of 6,000. We choose to use the same hyperparameter settings across all splits of the data. The Mean Average Precision results for each split are shown in Table I. From this table we can see that the network’s performance can vary depending upon the training and testing split used. In the next section we explore how the detection performance is affected by numerous factors in our dataset.

##### B. Qualitative Results

As our data has a wide variety of views of each object, varying pose and scale, we wanted to see how the detector fared with respect to different views. Figure 4 shows how detection score changed when camera position changed relative to an object instance. We can see there is a clear pattern showing the detector is more reliable in some camera positions than in others. Figure 6 shows how occlusion and object pose can greatly impact the detector even though there are training examples for both cases. We observed similar performance for many objects in all of our test scenes. This behavior motivates an active system that can move from a position with poor detection outputs to one with improved performance.

##### C. Ability to Simulate Motion

There are many parts of a robotic system that may be impacted by movement, but we are focused on the vision system, in particular object recognition. To find an appropriate sampling resolution for object recognition, we see how a vision system’s output changes as a function of camera movement. We need to find a sampling resolution that can simulate motion but is also practical for data collection purposes.



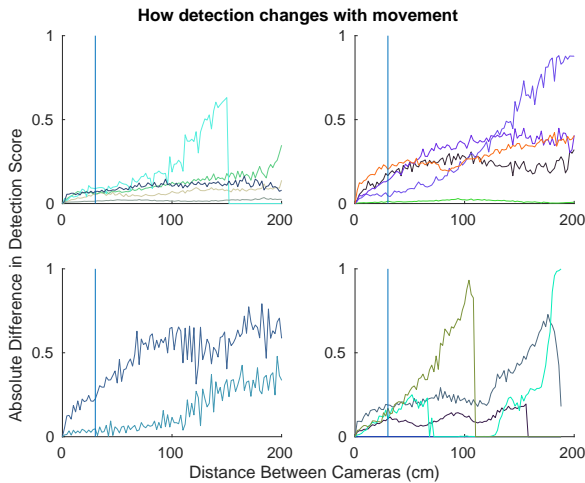


Fig. 5. How sensitive our detection system is to change in camera position. As the distance between two images of an instance increases(x-axis), the change in detection score(y-axis) tends to increase. Each line represents one instance. The vertical blue line shows our chosen sampling resolution of 30cm.

We first drive our robot around some scenes, capturing video as if the robot is naturally moving through the environment. We then label all BigBIRD instances in the videos, and run our instance detector on each image. For each video, we calculate the difference in detection score for each instance in all pairs of images. For example, we take the fourth and tenth frame and plot the difference in score for an instance against the distance the camera moved between frames. We plot the results from four videos in Figure 5.

For all instances that were detected in at least one image (score greater than 0), even the smallest movement of the camera results in some change in detection score. As the distance between cameras increases, there is a greater change in detection score. We considered the trade-off of having lower variation in our vision system against practicality of data collection. The vertical blue line in each plot in Figure 5 shows our chosen resolution, 30 cm. We found that for most instances, the change in score at 30cm is not much different than the changes at smaller resolutions like 10 or 20cm.

#### D. Active Vision

In this section we propose a baseline for an active instance classification task on our dataset. We envision a scenario where a robotic system is given an area of interest, and the system must classify the object instance at that location. We assume that given an initial area, localizing the same area in subsequent images is straight forward. Based on these assumptions, we propose the following problem setting. As input our agent receives an initial image with a bounding box for the target object. The agent can then choose an action at each timestep and will receive a new image and bounding box corresponding with the action. The goal is for the agent to learn an action policy which will increase the accuracy of the instance classifier.

A straightforward way of training an active vision system

for object recognition would be to train the system to acquire new views of an object when there is occlusion. However, it is not easy to label and quantify the level of occlusion of a target object. Furthermore, even if these labels were readily available our intuition about which views are difficult for a classifier would not necessarily be correct. For example, a classifier may be able to easily recognize some heavily occluded objects by only looking at some small discriminative part of the object. In addition, our dataset contains numerous factors which make the classification task difficult in addition to occlusions, such as varying object scale and lighting conditions. Therefore, we choose to use classification score as the training signal for our active vision system. A new view of an instance can increase both the confidence and accuracy of our classifier. This leads our model to learn a policy which attempts to move the agent to views that improve recognition performance.

As a feature extractor, we used the first 9 convolutional layers of ResNet-18 models [9], which recently showed compelling results on the 1000 way imagenet classification task. We used pre-trained models written in the torch framework [34]. The weights for the network are fixed for all experiments although our overall system is end-to-end trainable. The instance classifier and action network share the feature extractor. See Figure 8.

We first train an instance classifier for BigBIRD [15] instances, which appear in our dataset. One natural choice might be to train the classifier and action network simultaneously on our dataset. However, deep neural networks can easily achieve almost 100% classification accuracy on our training dataset. This type of over-fitting would prevent our action network from learning a meaningful policy, and does not perform well on the test set.

Thus, we use images from the BigBIRD [15] dataset for training our instance classifier. Even though the BigBIRD dataset provides many viewpoints of an instance, it can't be directly used for training since it consists of objects against a plain white background. We instead use the provided object masks to crop the object and overlay it on a random background sampled from SUN397 dataset[35], [36]. In order to prevent our network from overfitting, we aggressively applied various data augmentations. These included randomly cropping part of the image, performing color jittering, and sampling different lightening. Additionally, since our dataset consists of many small object instances, we randomly scaled the object by a factor ranging from 0.02 to 1.

Our baseline action network is inspired by a recent active vision approach [37], [26]. We use the REINFORCE algorithm to train a network to predict an action at each time step. At each time step our action network receives as input an image and a bounding box for the current position. Our network then outputs a score for each action: forward, backward, left, right, clockwise rotation, and counter-clockwise rotation. We fix the maximum number of timesteps during training to be  $T = 5$  steps or until the classifier achieves more than 0.9 confidence score. If the instance classifier correctly classifies the instance at the final timestep or reaches a 0.9 score at any




Fig. 6. Example of how movement affects detection output for a single instance. The proposed box with highest score  $> .1$  for the crystal hot sauce bottle instance is shown in each image. Object instance and scene correspond to the bottom left plot in Figure 4

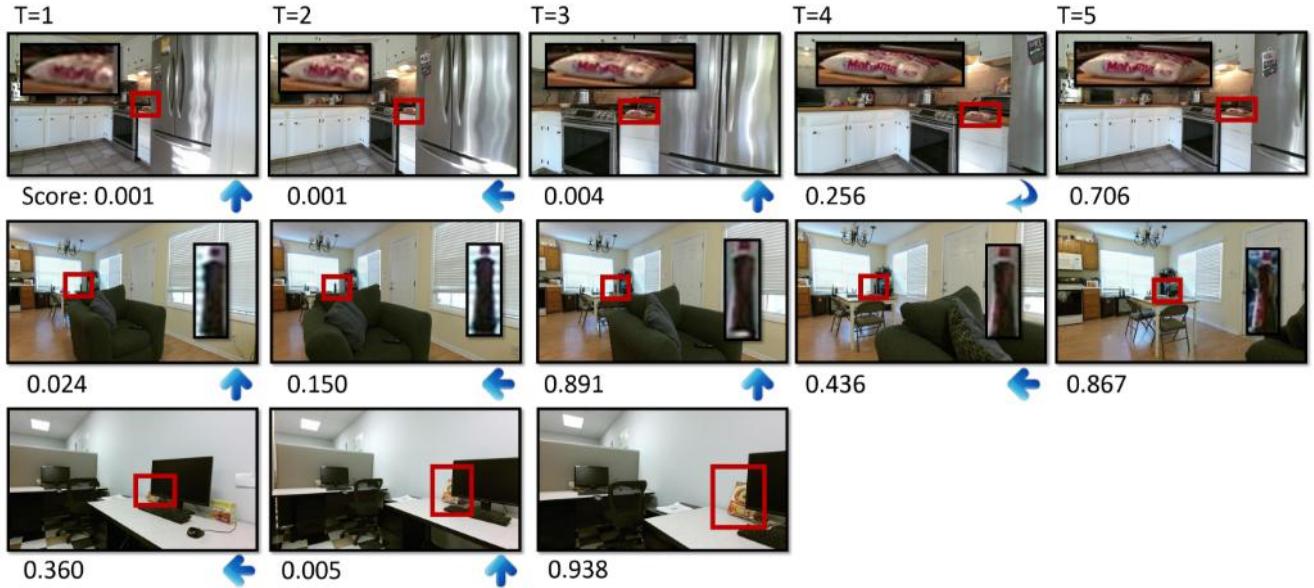


Fig. 7. Example paths taken by our active vision system. The arrow indicates the action chosen by the action network.

timestep, we consider the actions taken by the action network as correct. We then give the network a positive reward signal to adjust the weights of the action network to encourage the chosen moves.

More formally, we want to maximize the expected reward with respect to the policy distribution represented by our action network.

$$J(\theta) = \mathbb{E}_{p(a_{1:T}|\phi(I_{1:T}), bb_{1:T}; \theta)}[R] \quad (1)$$

Where  $\phi(I_{1:T})$  are the CNN features for the images,  $bb_{1:T}$  are the bounding boxes of target objects. If the classification is correct  $R$  is the score of the classifier, otherwise  $R = 0$ . For simplicity, we assumed the policy distributions to be independent at each timestep,  $p(a_{1:T}|\phi(I_{1:T}), bb_{1:T}; \theta) = \prod_t^T p(a_t|\phi(I_{t-1}), bb_{t-1}; \theta)$ . In order to compute gradients

with respect to the parameters of our action network, we use the REINFORCE algorithm, which is sample approximation to the gradient introduced by [38] and recently popularized by [39].

$$\nabla_{\theta} J \approx \frac{1}{M} \sum_{i=1}^M \sum_{t=1}^T \nabla_{\theta} \log p(a_t^i | \phi(I_{t-1}^i), bb_{t-1}^i; \theta) R^i \quad (2)$$

We evaluate our action network by comparing the accuracy of our classifier at different timesteps. The action network is used to choose an action at each image location at each time step, moving to a new image location for the next timestep. We consider how the classification accuracy changes as the maximum timestep,  $T$ , increases.

Since many of the instances in our dataset are small and far away in the image a natural baseline policy is one that always

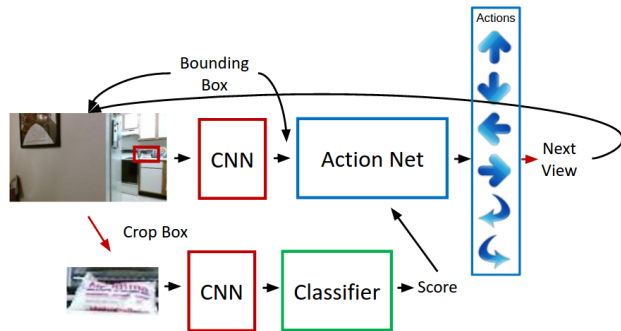


Fig. 8. Overall architecture of our active recognition system. It consists of three components. A CNN for extracting image features from the entire image given the current view, an instance classifier for classifying the cropped object, and an action network for selecting the next action in order to improve classification.

chooses the move forward action. We additionally compare against a policy of choosing a random action. Figure 9 shows how our system is able to greatly improve classification accuracy by moving to new image locations. We are also able to outperform the two obvious baselines. Figure 7 shows some qualitative examples of our system moving through a scene.

One potential improvement to our active classification model is a method for aggregating the views at each time step in order to choose the next action and perform multi-view classification[25], [26]. We also would like to explore the recurrent models that could consider history of actions taken. Additionally, the active vision task difficulty can be further increased by not providing the bounding box. This would require a policy that considers several hypothesis of both the location and class of the object. We expect that our dataset will provide a challenging test bed for further active vision research.

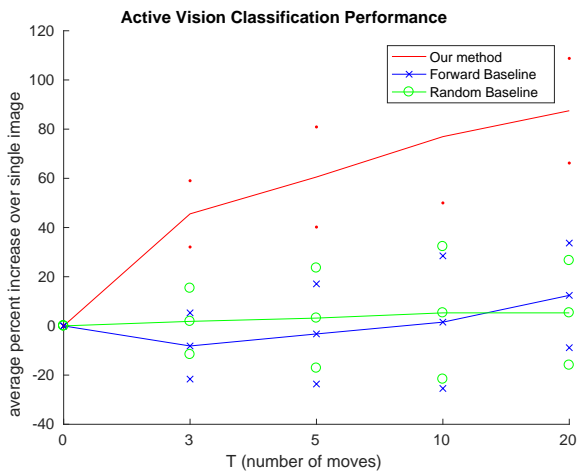


Fig. 9. The relative improvement in classification accuracy for different active vision policies. As the system makes more virtual moves through the scene(T increases), our method is able to move to a position that increases classification performance. Making random moves, or just moving forward, does not improve performance much.

Number of Moves	0	3	5	10	20
Method					
Split 1					
Ours	.30	.43	.45	.49	.51
Random	.30	.26	.28	.28	.33
Forward	.30	.29	.29	.29	.29
Split 2					
Ours	.25	.40	.46	.52	.53
Random	.25	.24	.26	.29	.33
Forward	.25	.29	.30	.31	.31
Split 3					
Ours	.42	.56	.62	.67	.73
Random	.42	.38	.40	.42	.46
Forward	.42	.39	.39	.40	.40

TABLE II

ACTIVE VISION RESULTS FOR DIFFERENT SPLITS. COLUMNS REPRESENT NUMBER OF MOVES. NUMBERS ARE ACCURACY OF THE CLASSIFIER, AVERAGED ACROSS ALL INSTANCES IN ALL TEST SCENES. THE GOAL OF OUR SYSTEM IS TO MOVE IN THE SCENE TO INCREASE CLASSIFICATION ACCURACY FOR A PARTICULAR INSTANCE.

## V. CONCLUSIONS

We introduce a new labeled dataset for developing and benchmarking object recognition methods in challenging indoor environments and active vision strategies for these tasks. We establish a baseline for object instance detection and show that the data is suitable for training a modern deep-learning-based system for next best view selection, using reinforcement learning, something that usually requires using a robot in the loop or synthetic computer graphics models. Using our densely sampled RGB-D imagery allows systems to see and be evaluated on real-world visual perception challenges which include large variations in scale and viewpoint as well as real imaging conditions that may not be present in CG. We validate experimentally that current state-of-the-art detection systems benefit from active vision on this real-world data. The dataset and toolbox for processing are now public.

## REFERENCES

- [1] K. Lai, B. Liefeng, and D. Fox, "Unsupervised feature learning for 3d scene labeling," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [2] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *European Conference on Computer Vision (ECCV)*, 2012.
- [3] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," *European Conference on Computer Vision (ECCV)*, 2016.
- [4] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [6] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conference on Computer Vision (ECCV)*. Springer, 2014, pp. 740–755.
- [7] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *International journal of computer vision*, vol. 104, no. 2, pp. 154–171, 2013.

- [8] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 580–587.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [10] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," *arXiv preprint arXiv:1506.02640*, 2015.
- [11] P. Poirson, P. Ammirato, C.-Y. Fu, J. Liu, Wei Koeck, and A. C. Berg, "Fast single shot detection and pose estimation," in *International Conference on 3DVision(3DV)*, 2016.
- [12] S. Song, L. Zhang, and J. Xiao, "Robot in a room: Toward perfect object recognition in closed environments," *CoRR*, vol. abs/1507.02703, 2015. [Online]. Available: <http://arxiv.org/abs/1507.02703>
- [13] D. "Held, S. Thrun, and S. Savarese, ""robust single-view instance recognition"," in "International Conference on Robotics and Automation (ICRA)", "2016".
- [14] M. Firman, "RGBD Datasets: Past, Present and Future," in *CVPR Workshop on Large Scale 3D Data: Acquisition, Modelling and Analysis*, 2016.
- [15] A. Singh, J. Sha, K. S. Narayan, T. Achim, and P. Abbeel, "Bigbird: A large-scale 3d database of object instances," *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [16] K. Lai, L. Bo, X. Ren, and D. Fox, "A large-scale hierarchical multi-view rgb-d object dataset," in *International Conference on Robotics and Automation (ICRA)*, 2011.
- [17] S. S., S. Lichtenberg, and J. Xiao, "Sun rgb-d: A rgb-d scene understanding benchmark suite," in *Proceedings of 28th IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [18] G. Georgakis, M. A. Reza, A. Mousavian, P.-H. Le, and J. Kosecka, "Multiview rgb-d dataset for object instance detection," in *International Conference on 3DVision(3DV)*, 2016.
- [19] R. Bajcsy, "Active perception," *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966–1005, 1988.
- [20] Z. Jia, A. Saxena, and T. Chen, "Robotic object detection: Learning to improve the classifiers using sparse graphs for path planning," in *Proceedings-International Joint Conference on Artificial Intelligence (IJCAI)*, vol. 22, no. 3, 2011, p. 2072.
- [21] V. Karasev, A. Chiuso, and S. Soatto, "Controlled recognition bounds for visual learning and exploration," in *Neural Information Processing Systems (NIPS)*, 2012, pp. 2915–2923.
- [22] N. Atanasov, B. Sankaran, J. Le Ny, T. Koletschka, G. J. Pappas, and K. Daniilidis, "Hypothesis testing framework for active object detection," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2013, pp. 4216–4222.
- [23] J. Velez, G. Hemann, A. S. Huang, I. Posner, and N. Roy, "Modelling observation correlations for active exploration and robust object detection," *Journal of Artificial Intelligence Research*, 2012.
- [24] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1912–1920.
- [25] H. Su, S. Maji, E. Kalogerakis, and E. G. Learned-Miller, "Multi-view convolutional neural networks for 3d shape recognition," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [26] S. L. E. Johns and A. Davison, "Pairwise decomposition of image sequences for active multi-view recognition," in *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [27] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6d object pose and predicting next-best-view in the crowd," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [28] Y. Furukawa and J. Ponce, "Accurate, dense, and robust multi-view stereopsis," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 32, no. 8, pp. 1362–1376, 2010.
- [29] T. Wiedemeyer, "IAI Kinect2," [https://github.com/code-iai/iai\\_kinect2](https://github.com/code-iai/iai_kinect2), Institute for Artificial Intelligence, University Bremen, 2014 – 2015, accessed June 12, 2015.
- [30] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [31] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, "Pixel-wise view selection for unstructured multi-view stereo," in *European Conference on Computer Vision (ECCV)*, 2016.
- [32] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, "Towards internet-scale multi-view stereo," in *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [33] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [34] R. Collobert, K. Kavukcuoglu, and C. Farabet, "Torch7: A matlab-like environment for machine learning," in *BigLearn, NIPS Workshop*, 2011.
- [35] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [36] H. Su, C. R. Qi, Y. Li, and L. J. Guibas, "Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views," in *IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [37] D. Jayaraman and K. Grauman, "Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion," in *European Conference on Computer Vision (ECCV)*, 2016.
- [38] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *MACHINE LEARNING*, vol. 8, no. 3, pp. 229–256, 1992.
- [39] V. Mnih, N. Heess, A. Graves, and K. kavukcuoglu, "Recurrent models of visual attention," in *Neural Information Processing Systems (NIPS)*, 2014.