	Box	Car	Human	Robot	Book	Airplane	Bus	Motorbike
Box	**0.958**	0	0.017	0.025	0	0	0	0
Car	0.010	**0.927**	0	0.021	0	0	0	0.042
Human	0.080	0.024	**0.820**	0.060	0.016	0	0	0
Robot	0.027	0	0.042	**0.899**	0.027	0	0	0.005
Book	0.016	0	0	0.042	**0.942**	0	0	0
Airplane	0.029	0.051	0	0.023	0.009	**0.888**	0	0
Bus	0	0.072	0	0	0	0	**0.856**	0.072
Motorbike	0	0.073	0	0.010	0.016	0	0.062	**0.839**

Table 1: Object recognition accuracy averaged over different models (i.e., subcategories) in each object category.

models (GMMs) in normalized HSV color space. The object model includes the GMMs, and their relative positions and sizes.

The learned object models are used for object recognition in novel scenes, irrespective of whether the objects are stationary or moving. For any test image, robots perform object recognition by identifying ROI(s) and computing the probability of occurrence of each learned object in the ROI(s). For image sequences with moving objects, ROIs are identified using the same approach used to identify ROIs during learning. For individual snapshots of objects, the iterated conditional modes (ICM) energy minimization algorithm is used to iteratively select candidate ROIs. For any candidate ROI and each learned object model, the robot computes the probability of occurrence of the corresponding object in the ROI based on each component of the object model. Probabilistic generative models that capture the conditional relationships between components of the learned model are then used to merge these probabilities and compute the net probability of occurrence of the corresponding object in the ROI. Robots are thus able to exploit the complementary properties of different visual cues for reliable and efficient object recognition in different scenes.

2. EXPERIMENTAL RESULTS

The test platform was a wheeled robot equipped with cameras that provide 640×480 images. Data from range finders were used to learn the domain map. Although the robot has Wi-Fi capability, all experiments were performed using an on-board 2GHz processor and 1GB RAM. Experimental trials evaluated the robot's ability to learn models of interesting objects from a small set of images using appearance-based and contextual visual cues, using these models to reliably and efficiently recognize objects in novel scenes.

Robots learned 30 different object models over eight different object categories: human, box, airplane, book, car, motorbike, bus and humanoid robot. Since it is a challenge to obtain an image dataset of objects with well-defined motion, experiments were conducted over ≈ 1400 images, ≈ 700 of which were captured by the robot in indoor and outdoor environments. To establish applicability to different domains, images of airplanes, motorbikes and buses (and some cars) were chosen from the *Pascal VOC2006* benchmark dataset to obtain $\approx$ 700 images. To make learning challenging, each object model was learned autonomously using $\approx 3 - 5$ images, with $\approx$ 150 images used for learning all object models; the remaining images were used for evaluation. The images used for learning and recognition were chosen randomly in repeated trials. The robot is able to process $3 - 5$ frames/second to identify moving objects, learn models and recognize objects in novel scenes, while performing other tasks such as path planning for navigation.

Table 1 shows the object recognition accuracy for different object categories, averaged over different subcategories in each category. Accurate recognition requires an object in a test image to be matched with the correct subcategory. The robot exploits complementary properties of visual cues to reliably learn object models and recognize objects in novel scenes—the combination of cues

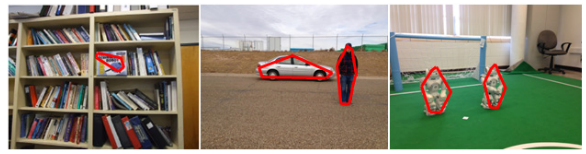

Figure 2: Illustrative examples of using the proposed algorithm to recognize one or more objects in test images.

provides higher accuracy than any single component. Most classification errors occur when an insufficient number of (test) image features are matched with the learned object models as a result of motion blur or a large difference in scale or viewpoint. Incremental revision of the learned object models helps eliminate many of these errors. Figure 2 shows examples of object recognition in test images—robots can reliably and efficiently recognize objects in cluttered backgrounds, and recognize multiple objects or multiple instances of the same object.

3. CONCLUSIONS

This paper described an approach that enables mobile robots to use the complementary properties of appearance-based and contextual visual cues to identify interesting objects and learn representative models from a small set of images. These object models are used for reliable and efficient object recognition in novel scenes. Future work will consider additional visual cues and fully integrate learning with planning and collaboration [5].

ACKNOWLEDGMENT

This work was supported in part by the ONR Science of Autonomy award N00014-09-1-0658.

4. REFERENCES

[1] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59(2):167–181, 2004.

[2] Xiang Li and Mohan Sridharan. Autonomous Learning of Object Models on a Mobile Robot using Visual Cues. In *International Conference on Robotics and Automation*, 2011.

[3] K. Mikolajczyk and C. Schmid. Scale and Affine Invariant Interest Point Detectors. *International Journal of Computer Vision*, 60(1):63–86, 2004.

[4] J. Porway and S. C. Zhu. C4: Computing Multiple Solutions in Graphical Models by Cluster Sampling. *Pattern Analysis and Machine Intelligence*, 33(9):1713–1727, 2011.

[5] Mohan Sridharan. Integrating Visual Learning and Hierarchical Planning for Autonomy in Human-Robot Collaboration. In *AAAI Spring Symposium Series, Designing Intelligent Robots: Reintegrating AI II*, March 2013.