# The Freiburg Groceries Dataset

Philipp Jund, Nichola Abdo, Andreas Eitel, Wolfram Burgard

arXiv:submit/1726170 [cs.CV] 17 Nov 2016

Abstract-With the increasing performance of machine learning techniques in the last few years, the computer vision and robotics communities have created a large number of datasets for benchmarking object recognition tasks. These datasets cover a large spectrum of natural images and object categories, making them not only useful as a testbed for comparing machine learning approaches, but also a great resource for bootstrapping different domain-specific perception and robotic systems. One such domain is domestic environments, where an autonomous robot has to recognize a large variety of everyday objects such as groceries. This is a challenging task due to the large variety of objects and products, and where there is great need for real-world training data that goes beyond product images available online. In this paper, we address this issue and present a dataset consisting of 5,000 images covering 25 different classes of groceries, with at least 97 images per class. We collected all images from real-world settings at different stores and apartments. In contrast to existing groceries datasets, our dataset includes a large variety of perspectives, lighting conditions, and degrees of clutter. Overall, our images contain thousands of different object instances. It is our hope that machine learning and robotics researchers find this dataset of use for training, testing, and bootstrapping their approaches. As a baseline classifier to facilitate comparison, we re-trained the CaffeNet architecture (an adaptation of the well-known AlexNet [20]) on our dataset and achieved a mean accuracy of 78.9%. We release this trained model along with the code and data splits we used in our experiments.

#### I. INTRODUCTION AND RELATED WORK

Object recognition is one of the most important and challenging problems in computer vision. The ability to classify objects plays a crucial role in scene understanding, and is a key requirement for autonomous robots operating in both indoor and outdoor environments. Recently, computer vision has witnessed significant progress, leading to impressive performance in various detection and recognition tasks [10, 14, 31]. On the one hand, this is partly due to the recent advancements in machine learning techniques such as deep learning, fueled by a great interest from the research community as well as a boost in hardware performance. On the other hand, publicly-available datasets have been a great resource for bootstrapping, testing, and comparing these techniques.

Examples of popular image datasets include ImageNet, CIFAR, COCO and PASCAL, covering a wide range of categories including people, animals, everyday objects, and much more [6, 8, 19, 22]. Other datasets are tailored towards specific domains such as house numbers extracted from Google Street View [24], face recognition [13], scene



Fig. 1: Examples of images from our dataset. Each image contains one or multiple instances of objects belonging to one of 25 classes of groceries. We collected these images in real-world settings at different stores and apartments. We considered a rich variety of perspectives, degree of clutter, and lighting conditions. This dataset highlights the challenge of recognizing objects in this domain due to the large variation in shape, color, and appearance of everyday products even within the the same class.

understanding and place recognition [28, 34], as well as object recognition, manipulation and pose estimation for robots [5, 11, 17, 27].

One of the challenging domains where object recognition plays a key role is service robotics. A robot operating in unstructured, domestic environments has to recognize everyday objects in order to successfully perform tasks like tidying up, fetching objects, or assisting elderly people. For example, a robot should be able to recognize grocery objects in order to fetch a can of soda or to predict the preferred shelf for storing a box of cereals [1, 30]. This is not only challenging due to the difficult lighting conditions and occlusions in real-world environments, but also due to the large number of everyday objects and products that a robot can encounter.

Typically, service robotic systems address the problem of object recognition for different tasks by relying on state-ofthe-art perception methods. Those methods leverage existing object models by extracting hand-designed visual and 3D descriptors in the environment [2, 12, 26] or by learning new feature representations from raw sensor data [4, 7]. Others rely on an ensemble of perception techniques and sources of information including text, inverse image search, cloud data, or images downloaded from online stores to categorize objects and reason about their relevance for different tasks [3, 16, 18, 25].

However, leveraging the full potential of machine learning approaches to address problems such as recognizing groceries and food items remains, to a large extent, unrealized. One of the main reasons for that is the lack of training data

All authors are with the University of Freiburg, 79110 Freiburg, Germany. This work has partly been supported by the German Research Foundation under research unit FOR 1513 (HYBRIS) and grant number EXC 1086.



Fig. 2: Number of images per class in our dataset.

for this domain. In this paper, we address this issue and present the Freiburg Groceries Dataset, a rich collection of 5000 images of grocery products (available in German stores) and covering 25 common categories. Our motivation for this is twofold: *i*) to help bootstrap perception systems tailored for domestic robots and assistive technologies, and to *ii*) provide a challenging benchmark for testing and comparing object recognition techniques.

While there exist several datasets containing groceries, they are typically limited with respect to the view points or variety of instances. For example, sources such as the OpenFoodFacts dataset or images available on the websites of grocery stores typically consider one or two views of each item [9]. Other datasets include multiple view points of each product but consider only a small number of objects. An example of this is the GroZi-120 dataset that contains 120 grocery products under perfect and real lighting conditions [23]. Another example is the RGB-D dataset covering 300 object instances in a controlled environment, of which only a few are grocery items [21]. The CMU dataset, introduced by Hsiao et al., considers multiple viewpoints of 10 different household objects [12]. Moreover, the BigBIRD dataset contains 3D models and images of 100 instances in a controlled environment [29].

In contrast to these datasets, the Freiburg Groceries Dataset considers challenging real-world scenes as depicted in Fig. 1. This includes difficult lighting conditions with reflections and shadows, as well as different degrees of clutter ranging from individual objects to packed shelves. Additionally, we consider a large number of instances that cover a rich variety of brands and package designs.

To demonstrate the applicability of existing machine learning techniques to tackle the challenging problem of recognizing everyday grocery items, we trained a convolutional neural network as a baseline classifier on five splits of our dataset, achieving a classification accuracy of 78.9%. Along with the dataset, we provide the code and data splits we used in these experiments. Finally, whereas each image in the main dataset contains objects belonging to one class, we include an additional set of 37 cluttered scenes, each containing several object classes. We constructed these scenes at our lab to emulate real-world storage shelves. We present qualitative examples in this paper that demonstrate using our classifier to recognize patches extracted from such images.

## II. THE FREIBURG GROCERIES DATASET

The main bulk of the Freiburg Groceries Dataset consists of 4947 images of 25 grocery classes, with 97 to 370 images per class. Fig. 2 shows an overview of the number of images per class. We considered common categories of groceries that exist in most homes such as pasta, spices, coffee, etc. We recorded this set of images, which we denote by  $\mathcal{D}_1$ , using four different smartphone cameras at various stores (as well as some apartments and offices) in Freiburg, Germany. The images vary in the degree of clutter and real-world lighting conditions, ranging from well-lit stores to kitchen cupboards. Each image in  $\mathcal{D}_1$  contains one or multiple instances of *one* of the 25 classes. Moreover, for each class, we considered a rich variety of brands, flavors, and packaging designs. We processed all images in  $\mathcal{D}_1$  by down-scaling them to a size of  $256 \times 256$  pixels. Due to the different aspect ratios of the cameras we used, we padded the images with gray borders as needed. Fig. 7 shows example images for each class.

Moreover, the Freiburg Groceries Dataset includes an additional, smaller set  $\mathcal{D}_2$  with 74 images of 37 cluttered scenes, each containing objects belonging to *multiple* classes. We constructed these scenes at our lab and recorded them using a Kinect v2 camera [32, 33]. For each scene, we provide data from two different camera perspectives, which includes a 1920×1080 RGB image, the corresponding depth image and a point cloud of the scene. We created these scenes to emulate real-world clutter and to provide a challenging benchmark with multiple object categories per image. We provide a "coarse" labeling of images in  $\mathcal{D}_2$  in terms of which classes exist in each scene.

We make the dataset available on this website: http://www2.informatik.uni-freiburg.de/ %7Eeitel/freiburg\_groceries\_dataset.html.



Fig. 3: Example test images for candy and pasta taken from the first split in our experiments. All images were correctly classified in this case. The classifier is able to handle large variations in color, shape, perspective, and degree of clutter.

There, we also include the trained classifier model we used in our experiments (see Sec. III). Additionally, we provide the code for reproducing our experimental results on github: https://github.com/PhilJd/freiburg\_groceries\_dataset.

## III. OBJECT RECOGNITION USING A CONVOLUTIONAL NEURAL NETWORK

To demonstrate the use of our dataset and provide a baseline for future comparison, we trained a deep neural network classifier using the images in  $\mathcal{D}_1$ . We adopted the CaffeNet architecture [15], a slightly altered version of the AlexNet [20]. We trained this model, which consists of five convolution layers and three fully connected layers, using the Caffe framework [15]. We initialized the weights of the model with those of the pre-trained CaffeNet, and fine-tuned the weights of the three fully-connected layers.

We partitioned the images into five equally-sized splits, with the images of each class uniformly distributed over the splits. We used each split as a test set once and trained on the remaining data. In each case, we balanced the training data across all classes by duplicating images from classes with fewer images. We trained all models for 10,000 iterations and always used the last model for evaluation.

We achieved a mean accuracy of 78.9% (with a standard deviation of 0.5%) over all splits. Fig. 3 shows examples of correctly classified images of different candy and pasta packages. The neural network is able to recognize the categories in these images despite large variations in appearance, perspectives, lighting conditions, and number of objects in each image. On the other hand, Fig. 4 shows examples of misclassified images. For example, we found products with white, plain packagings to be particularly challenging, often mis-classified as flour. Another source of difficulty is products with "misleading" designs, e.g., pictures of fruit (typically found on juice cartons) on cereal boxes.

Fig. 6 depicts the confusion matrix averaged over the five splits. The network performs particularly well for classes such as water, jam, and juice (88.1%-93.2%), while on the other hand it has difficulties correctly recognizing objects







Class: Beans Predicted: Sugar

Class: Chocolate Class: Cake Predicted: Flour Predicted: Chocolate







Class: Cereal Predicted: Juice Class: Candy Predicted: Coffee Class: Nuts Predicted: Beans

Fig. 4: Examples of misclassification that highlight some of the challenges in our dataset, e.g., cereal boxes with drawings of fruits on them are sometimes confused with juice.

from the class flour (59.9%). We provide the data splits and code needed to reproduce our results, along with the a Caffe model trained on all images of  $\mathcal{D}_1$ , on the web pages mentioned in Sec. II.

Finally, we also performed a qualitative test where we used a model trained on images in  $\mathcal{D}_1$  to classify patches extracted from images in  $\mathcal{D}_2$  (in which each image contains multiple object classes). Fig. 5 shows an example for classifying different manually-selected image patches. Despite a sensitivity to patch size, this shows the potential for using  $\mathcal{D}_1$ , which only includes one class per image, to recognize objects in cluttered scenes where this assumption does not hold. An extensive evaluation on such scenes is outside the scope of this paper, which we leave to future work.





(c)

Fig. 5: We used a classifier trained on  $\mathcal{D}_1$  to recognize manually-selected patches from dataset  $\mathcal{D}_2$ . Patches with a green border indicate a correct classification whereas those with a red border indicate a misclassification. We rescaled each patch before passing it through the network. (a) depicts the complete scene. (b) shows an example of the sensitivity of the network to changes in patch size. (c) shows classification results for some manually extracted patches.

### **IV. CONCLUSION**

In this paper, we introduced the Freiburg Groceries Dataset, a novel dataset targeted at the recognition of groceries. Our dataset includes ca 5000 labeled images, organized in 25 classes of products, which we recorded in several stores and apartments in Germany. Our images cover a wide range of real-world conditions including different viewpoints, lighting conditions, and degrees of clutter. Moreover, we provide images and point clouds for a set of 37 cluttered scenes, each consisting of objects from multiple classes. To facilitate comparison, we provide results averaged over five train/test splits using a standard deep network architecture, which achieved a mean accuracy of 78.9% over all classes. We believe that the Freiburg Groceries Dataset represents an interesting and challenging benchmark to evaluate stateof-the-art object recognition techniques such as deep neural networks. Moreover, we believe that this real-world training data is a valuable resource for accelerating a variety of service robot applications and assistive systems where the ability to recognize everyday objects plays a key role.

#### REFERENCES

[1] N. Abdo, C. Stachniss, L. Spinello, and W. Burgard. Organizing objects by predicting user preferences through collaborative filtering.

Int. J. of Robotics Research (IJRR), 35(13):1587-1608, 2016.

- [2] L. A. Alexandre. 3d descriptors for object and category recognition: a comparative evaluation. In Workshop on Color-Depth Camera Fusion in Robotics at Int. Conf. on Intelligent Robots and Systems (IROS), 2012.
- [3] M. Beetz, F. Bálint-Benczédi, N. Blodow, D. Nyga, T. Wiedemeyer, and Z.-C. Márton. Robosherlock: unstructured information processing for robot perception. In *Int. Conf. on Robotics & Automation (ICRA)*, 2015.
- [4] L. Bo, X. Ren, and D. Fox. Unsupervised Feature Learning for RGB-D Based Object Recognition. In Int. Symposium on Experimental Robotics (ISER), 2012.
- [5] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols. *arXiv preprint* arXiv:1502.03143, 2015.
- [6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [7] A. Eitel, J. T. Springenberg, L. Spinello, M. Riedmiller, and W. Burgard. Multimodal deep learning for robust rgb-d object recognition. In Int. Conf. on Intelligent Robots and Systems (IROS), 2015.
- [8] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *Int. J. of Computer Vision (IJCV)*, 88(2):303–338, 2010.
- [9] S. Gigandet. Open food facts. http://en.wiki. openfoodfacts.org/. Accessed: 2016-11-04.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Int. Conf. on Computer Vision (ICCV)*, 2015.
- [11] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Asian Conf.* on Computer Vision, 2012.
- [12] E. Hsiao, A. Collet, and M. Hebert. Making specific features less discriminative to improve point-based 3d object recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [13] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, 2007.
- [14] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167, 2015.
- [15] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093, 2014.
- [16] P. Kaiser, M. Lewis, R. Petrick, T. Asfour, and M. Steedman. Extracting common sense knowledge from text for robot planning. In *Int. Conf. on Robotics & Automation (ICRA)*, 2014.
- [17] A. Kasper, Z. Xue, and R. Dillmann. The kit object models database: An object model database for object recognition, localization and manipulation in service robotics. *Int. J. of Robotics Research (IJRR)*, 31(8):927–934, 2012.
- [18] B. Kehoe, A. Matsukawa, S. Candido, J. Kuffner, and K. Goldberg. Cloud-based robot grasping with the google object recognition engine. In *Int. Conf. on Robotics & Automation (ICRA)*, 2013.
- [19] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in Neural Information Processing Systems (NIPS), 2012.
- [21] K. Lai, L. Bo, X. Ren, and D. Fox. A large-scale hierarchical multiview rgb-d object dataset. In *Int. Conf. on Robotics & Automation* (ICRA), 2011.
- [22] 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 Conf. on Computer Vision (ECCV)*, 2014.
- [23] M. Merler, C. Galleguillos, and S. Belongie. Recognizing groceries in situ using in vitro training data. In *IEEE Conf. on Computer Vision* and Pattern Recognition (CVPR), 2007.
- [24] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

- [25] D. Pangercic, V. Haltakov, and M. Beetz. Fast and robust object detection in household environments using vocabulary trees with sift descriptors. In Workshop on Active Semantic Perception and Object Search in the Real World at the Int. Conf. on Intelligent Robots and Systems (IROS), 2011.
- [26] K. Pauwels and D. Kragic. Simtrack: A simulation-based framework for scalable real-time object pose detection and tracking. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2015.
- [27] C. Rennie, R. Shome, K. E. Bekris, and A. F. De Souza. A dataset for improved rgbd-based object detection and pose estimation for warehouse pick-and-place. *IEEE Robotics and Automation Letters*, 1(2):1179–1185, 2016.
- [28] N. Silberman and R. Fergus. Indoor scene segmentation using a structured light sensor. In Workshop on 3D Representation and Recognition at the Int. Conf. on Computer Vision (ICCV), 2011.
- [29] A. Singh, J. Sha, K. S. Narayan, T. Achim, and P. Abbeel. Bigbird: A large-scale 3d database of object instances. In *Int. Conf. on Robotics* & Automation (ICRA), 2014.
- [30] S. S. Srinivasa, D. Ferguson, C. J. Helfrich, D. Berenson, A. Collet,

R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. V. Weghe. Herb: a home exploring robotic butler. *Autonomous Robots*, 28(1): 5–20, 2010.

- [31] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [32] T. Wiedemeyer. IAI Kinect2, 2014 2015. URL https:// github.com/code-iai/iai\_kinect2.
- [33] L. Xiang, F. Echtler, C. Kerl, T. Wiedemeyer, Lars, hanyazou, R. Gordon, F. Facioni, laborer2008, R. Wareham, M. Goldhoorn, alberth, gaborpapp, S. Fuchs, jmtatsch, J. Blake, Federico, H. Jungkurth, Y. Mingze, vinouz, D. Coleman, B. Burns, R. Rawat, S. Mokhov, P. Reynolds, P. Viau, M. Fraissinet-Tachet, Ludique, J. Billingham, and Alistair. libfreenect2: Release 0.2, 2016. URL https://doi. org/10.5281/zenodo.50641.
- [34] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In Advances in Neural Information Processing Systems (NIPS), 2014.

	beans	cake	candy	cereal	chips	chocolate	coffee	corn	flour	honey	jam	juice	milk	nuts	oil	pasta	rice	soda	spices	sugar	tea	tomato sauce	tuna	vinegar	water
beans	78.1	0.5	0.5	1.9	0.0	1.5	0.0	3.1	0.0	1.8	2.8	0.8	0.8	0.8	0.5	0.8	0.8	0.0	1.0	0.8	0.8	1.1	0.0	0.0	0.8
cake	0.0	82.5	0.6	5.2	0.9	2.9	1.1	0.6	0.0	0.0	0.0	0.0	0.4	0.4	0.0	1.9	0.0	0.0	0.0	0.0	1.8	0.6	0.6	0.0	0.0
candy	0.0	0.5	82.9	1.1	1.7	1.9	0.8	0.0	0.2	0.2	0.5	1.0	1.1	1.9	0.0	0.2	0.2	0.0	0.0	0.2	1.9	1.6	0.6	0.0	0.8
cereal	0.3	4.8	2.9	78.2	0.3	1.0	0.6	0.3	0.7	0.0	0.0	1.1	0.6	1.8	0.4	1.3	0.7	0.0	1.0	0.6	1.4	0.4	0.3	0.0	0.4
chips	2.2	1.6	7.1	1.1	70.1	1.6	1.1	0.5	0.5	0.0	0.0	6.0	0.0	4.4	0.0	1.6	0.0	0.0	0.0	0.5	0.5	0.5	0.0	0.0	0.0
chocolate	0.0	3.1	2.0	4.0	1.0	70.4	4.7	0.0	1.3	0.9	0.3	0.6	0.9	3.3	0.0	0.6	0.0	0.7	0.0	0.2	4.8	0.3	0.0	0.0	0.0
coffee	0.0	0.3	0.3	1.3	1.7	5.1	71.9	0.0	0.7	1.0	3.9	1.3	1.0	0.9	1.5	0.3	0.0	0.6	2.4	0.3	2.2	0.3	0.9	0.0	1.3
corn	0.9	0.0	0.0	0.0	0.0	0.0	0.9	88.3	0.0	1.9	0.0	3.8	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.9	1.0	0.0	0.0	1.0
flour	2.0	1.7	3.2	5.2	0.0	0.8	0.9	0.8	59.9	0.8	0.0	1.7	3.8	0.8	0.8	6.4	0.8	0.0	0.9	2.5	5.4	0.0	0.0	0.0	0.8
honey	0.0	0.4	0.4	0.4	1.1	1.0	2.3	0.0	0.0	73.6	9.2	2.1	0.0	1.1	0.5	0.0	0.0	1.0	0.4	0.0	2.5	1.1	0.0	1.0	1.0
jam	0.0	0.0	0.4	0.0	0.3	0.0	2.5	0.0	0.0	0.8	91.4	0.0	0.0	0.4	0.0	0.0	0.0	0.4	0.7	0.0	0.7	1.5	0.4	0.0	0.0
juice	0.0	0.3	1.0	0.0	0.0	0.7	0.6	0.0	0.0	0.0	0.6	88.1	0.4	0.0	0.6	0.0	0.0	0.6	0.6	0.0	1.2	2.0	0.0	2.0	0.7
milk	0.0	0.0	0.0	1.0	0.0	0.5	3.3	0.0	0.0	0.0	0.0	5.7	81.7	0.0	0.0	0.0	0.7	0.6	0.5	1.7	2.6	1.2	0.0	0.0	0.0
nuts	2.2	1.2	2.9	4.1	4.4	8.2	1.0	0.5	0.0	1.2	1.0	0.4	0.0	67.7	0.0	1.2	0.4	0.0	0.0	0.0	0.4	0.9	0.0	0.0	1.0
oil	0.6	0.0	0.0	0.5	0.8	0.0	0.0	0.0	0.6	0.0	1.3	4.1	0.6	0.0	78.0	0.0	0.0	1.6	0.5	0.0	1.4	0.0	0.0	9.3	0.0
pasta	0.5	0.5	0.5	4.0	1.8	3.4	3.1	0.0	1.6	0.0	1.2	0.5	0.0	0.0	1.1	76.6	1.1	0.0	1.0	0.0	2.3	0.0	0.0	0.0	0.0
rice	0.0	1.1	1.0	5.0	1.9	0.7	2.3	0.7	1.7	0.0	0.0	0.5	1.6	0.6	0.0	3.4	69.3	0.0	2.6	0.0	4.5	0.5	0.6	0.6	0.5
soda	0.0	0.0	0.0	0.0	1.1	0.9	1.5	0.0	0.0	0.4	1.6	8.5	0.0	0.0	0.8	0.0	0.0	71.1	0.0	0.0	0.0	2.3	0.0	3.2	8.2
spices	0.4	0.0	0.4	1.0	0.5	0.6	1.3	0.0	0.4	0.5	2.4	2.1	1.8	0.4	0.6	0.0	0.0	0.9	80.1	1.3	0.8	1.0	0.9	0.8	1.0
sugar	0.7	0.7	1.3	0.0	0.6	0.0	0.7	0.0	2.3	0.0	0.7	0.7	2.8	2.0	0.0	0.0	1.4	0.0	1.1	82.0	1.1	1.1	0.0	0.0	0.0
tea	0.0	0.0	2.4	2.5	0.3	3.2	2.9	0.0	0.0	0.0	0.3	2.1	0.7	0.3	1.0	0.8	0.0	0.3	0.3	0.3	81.0	0.6	0.3	0.0	0.0
tomato sauce	1.1	0.0	0.6	0.6	0.5	1.0	0.5	0.0	0.5	0.5	3.5	1.8	1.1	0.6	0.0	1.2	0.0	0.5	1.8	0.0	0.0	82.9	0.6	0.0	0.0
tuna	0.0	0.0	0.0	0.9	0.0	3.1	1.6	1.1	0.0	0.7	0.0	0.0	0.7	1.1	0.0	0.7	0.0	0.0	0.0	0.0	3.5	0.0	85.3	0.0	0.8
vinegar	0.0	0.0	0.5	0.0	0.0	0.0	1.4	0.0	0.0	0.7	1.7	9.7	0.5	0.0	6.3	0.0	0.0	1.0	1.8	0.0	0.7	0.0	0.7	73.6	0.7
water	0.0	0.0	0.7	0.4	0.3	0.4	0.3	0.0	0.0	0.0	0.3	0.7	0.0	0.4	0.6	0.0	0.0	1.0	0.0	0.0	0.0	0.8	0.0	0.3	93.2

Fig. 6: The confusion matrix averaged over the five test splits. We achieve a mean accuracy of 78.9% over all classes.



Fig. 7: Example images for each class.