

DIVERSE, DIFFICULT, AND ODD INSTANCES (D2O): A NEW TEST SET FOR OBJECT CLASSIFICATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Test sets are an integral part of evaluating models and gauging progress in object recognition, and more broadly in computer vision and AI. Existing test sets for object recognition, however, suffer from shortcomings such as bias towards the ImageNet characteristics and idiosyncrasies (*e.g.* ImageNet-V2), being limited to certain types of stimuli (*e.g.* indoor scenes in ObjectNet), and underestimating the model performance (*e.g.* ImageNet-A). To mitigate these problems, here we introduce a new test set, called D2O, which is sufficiently different from existing test sets. Images are diverse, unmodified, and representative of real-world scenarios and cause state-of-the-art models to misclassify them with high confidence. To emphasize generalization, our dataset by design does not come paired with a training set. It contains 8,060 images spread across 36 categories, out of which 29 appear in ImageNet. The best Top-1 accuracy on our dataset is around 60% which is much lower than 91% best Top-1 accuracy on ImageNet. We find that popular vision APIs perform very poorly in detecting objects over D2O categories such as “faces”, “cars”, and “cats”. Our dataset also comes with a “miscellaneous” category, over which we test the image tagging models. Overall, our investigations demonstrate that the D2O test set contain a mix of images with varied levels of difficulty and is predictive of the average-case performance of models. It can challenge object recognition models for years to come and can spur more research in this fundamental area. Data and code are publicly available at [Masked].

1 INTRODUCTION

The object recognition problem remains in an unclear state. Despite compelling performance of state-of-the-art object recognition methods, several questions such as out of distribution generalization (Recht et al., 2019; Barbu et al., 2019; Shankar et al., 2020; Taori et al., 2020; Koh et al., 2020), “superhuman performance” (He et al., 2015; Geirhos et al., 2018), adversarial vulnerability (Goodfellow et al., 2014), and invariance to image transformations and distortions (Hendrycks & Dietterich, 2019) still persist. Raw performance on test sets has been the main indicator of the progress and the major feedback about the state of the field. Few test sets have been proposed for evaluating object recognition models. Some follow the footsteps of ImageNet (Recht et al., 2019). Some filter images based on failures of models (Hendrycks et al., 2021). Researchers have also used controlled settings to collect data (Barbu et al., 2019; Borji et al., 2016). While being invaluable, these datasets suffer from few shortcomings. For example, datasets that only include examples for which the best models fail give the worst case scenario accuracy. While being useful, they underestimate model performance. Datasets that are biased towards certain environments (*e.g.* indoor scenes in ObjectNet (Barbu et al., 2019)), may not capture the full spectrum of visual stimuli. Most of the new datasets for object recognition have been centered on ImageNet (*e.g.* ImageNet-V2, ImageNet-A, ImageNet-O) and thus may have inherited its biases. This may in turn give us a biased assessment of visual recognition capability of models. We argue that a good test set should strike the right balance between sample difficulty and diversity and should reflect the average-case performance of models. We also believe that new test sets that are sufficiently different from existing ones can provide new insights into the object recognition problem. To this end, we include images that contain rich semantics and require cognitive processing (*e.g.* artistic scenes). Existing datasets lack enough samples of such cases.

Here, we emphasize image diversity and difficulty over scale. While scaling up test sets has a clear advantage (*e.g.* covering rare cases), it comes with some shortcomings. It is hard to ensure privacy,




Figure 1: Sample images from D2O dataset. Categories in order are: acorn, basketball, and cat (left column), banana, car, and clock (right column).

security, quality, and spurious correlations¹ in datasets containing millions of images. These problems are easier to tackle in small scale expert-made datasets. Nonetheless, both small and large datasets are needed and are complementary.

Our dataset includes 8,060 images across 36 categories (Fig. 1). Images are carefully collected, verified, and labeled. We do not limit ourselves to object recognition models proposed in academia, and also consider prominent vision APIs in industry. This allows us to test models over a wider range of categories than those available in ImageNet and obtain a broader sense of image understanding by models. State-of-the-art models show a 30% absolute drop in Top-1 acc on D2O test set compared to the best ImageNet accuracy (around 20% drop using Top-5 acc). Further, over categories for which we know humans are very good at (*e.g.* faces, cars), current APIs fail drastically.

D2O test set is intentionally not paired with a training set. It comes with a license that disallows researchers to update the parameters of any model on it. This helps avoid over-fitting on the dataset. Additionally, to mitigate the danger of leaking our data to other datasets, we mark every image by a one pixel green border which must be removed on the fly before using.

2 OBJECT RECOGNITION TEST SETS

A plethora of datasets have been proposed for image classification². Here, we are concerned with datasets that focus on core object recognition. ImageNet (Deng et al., 2009) is one of the most used datasets in computer vision and deep learning. It contains 1000 classes of common objects, with more than a million training images. Its test set contains 50,000 images. ImageNet test examples tend to be simple (by today’s standards), clear, and close-up images of objects. As such, they may not represent harder images encountered in the real world. Further, ImageNet annotations are limited to a single label per image. To remedy the problems with ImageNet, new test sets have been proposed, which have been instrumental in gauging performance of models and measuring the gap between models and humans. The major ones are reviewed below. In addition to these, several other datasets such as CIFAR-100 (Krizhevsky et al., 2009), SUN (Xiao et al., 2010), Places (Zhou et al., 2017), ImageNet-Sketch (Wang et al., 2019), and iLab20M (Borji et al., 2016) have also been introduced.

ImageNet-V2. Recht et al. (2019) built this test by closely following the ImageNet creation process. They reported a performance gap of about 11% (Top-1 acc.) between the performance of the best models on this dataset and their performance on the original test set. Engstrom et al. (2020) estimated that the accuracy drop from ImageNet to ImageNet-V2 is less than 3.6%. Some other works have also evaluated and analyzed models on this dataset (Shankar et al., 2020; Taori et al., 2020).

ImageNet-A, ImageNet-O, ImageNet-C, and ImageNet-P. These datasets are built to measure the robustness of image classifiers against out of distribution examples and image distortions (Hendrycks et al., 2021; Hendrycks & Dietterich, 2019; Hendrycks et al., 2020). Specifically, ImageNet-A dataset contains images for which a pre-trained ResNet-50 model fails to predict the correct label (Fig. 2). It has 7,500 images scrapped from iNaturalist, Flickr, and DuckDuckGo websites. Following the approach in Hendrycks et al. (2021), researchers have gathered “natural adversarial examples” for other problems such as object detection (Lau et al., 2021) and microscopy analysis (Pedraza et al., 2022). In contrast to these datasets which benchmark the worst case performance of models, here we are interested in the average case performance. To this end, instead of filtering images to fool a

¹Deep models take advantage of correlations between testing and training sets (*a.k.a* “spurious cues” or “shortcuts”). These correlations are easily accessible to models but are not to humans (Geirhos et al., 2020).

²<https://paperswithcode.com/datasets?task=image-classification&page=2>

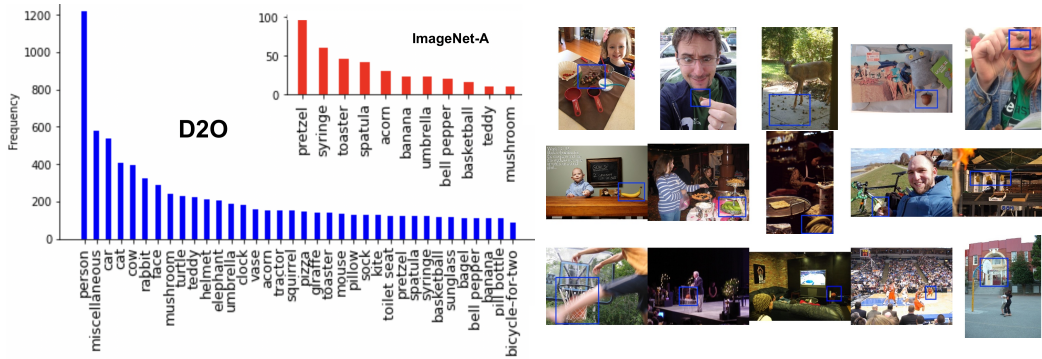


Figure 2: Left: Distribution of number of objects per category over our D2O dataset as well as the 11 classes from ImageNet-A that are in common with D2O (inset panel). This subset of Imagenet-A has a total of 393 images. In total, our dataset contains 8060 images over 36 categories. Right: Sample images from acorn, banana, and basketball categories of Imagenet-A. The blue box demonstrates the cropped region used to build the Isolated ImageNet-A dataset. See also Appendix B

classifier, we include a mix of easy and hard examples to get a better sense of accuracy. Notice that a model that can only solve a very hard test set is not guaranteed to also solve an easy one.

Reassessed Labels (ReaL). Beyer et al. (2020) collected new human annotations over the ImageNet validation set and used them to reassess the accuracy of ImageNet classifiers. They showed that model gains are substantially smaller than those reported using the original ImageNet labels. Further, they found that ReaL labels eliminate more than half of the ImageNet labeling mistakes. This implies that they provide a superior estimate of the model accuracy.

ObjectNet. To remove the biases of the ImageNet, Barbu et al. (2019) introduced the ObjectNet dataset. Images are pictured by Mechanical Turk workers using a mobile app in a variety of backgrounds, rotations, and imaging viewpoints. ObjectNet contains 50,000 images across 313 categories, out of which 113 are in common with ImageNet categories. Astonishingly, Barbu *et al.* found that the state-of-the-art object recognition models perform drastically lower on ObjectNet compared to their performance on ImageNet (about 40-45% drop). Later on, Borji (2021) revisited the Barbu *et al.* ’s results and found that applying deep models to the isolated objects, rather than the entire scene as is done in the original paper, leads to 20-30% performance improvement.

3 D2O TEST SET

We followed two approaches to collect the data. In the first one, we used publicly-available and free-to-distribute sources. We crawled images from the Flickr and Google image search engine using different search queries. The queries contained terms specifying countries, locations, materials, different times (*e.g.* 80s), odd appearances (*e.g.* odd helmet), etc. We also included images from various categorized panels in the search results (*e.g.* drawing, sketch, clip art, icon, neon, clay, etc.). In the second approach, we used image generation tools such as DALL-E 2 (Ramesh et al., 2022), Midjourney, and Stable Diffusion (Rombach et al., 2021) to generate some image or search for some images generated these tools. We only selected the images that had good quality. Some sample generated images are shown in Appendix I. We did our best to ensure that images do contain sensitive material, have poor resolution, or violate copyright law³. The gathered encompass a wide variety of visual concepts over both RGB images, paintings, drawings, cartoons, and clip arts. To reduce ambiguity, most of the images contain one main object. Categories were selected based on the following two criteria: b) they should be possible to collect a variety of instances for them with different levels of difficulty, and c) one would consider model errors on them egregious (*i.e.* confusing a cat with a dog is more troublesome than confusing a beaver with a marmot). During data collection, we emphasized choosing the odd items.

Three annotators, who were also computer vision researchers, were presented with an image as well as its corresponding label. They were asked to verify the label by checking the correct or the incorrect box. We found that the three annotators agreed with the correct label over all images.

³We choose images that were public domain, did not have copyright, or were released by the government.

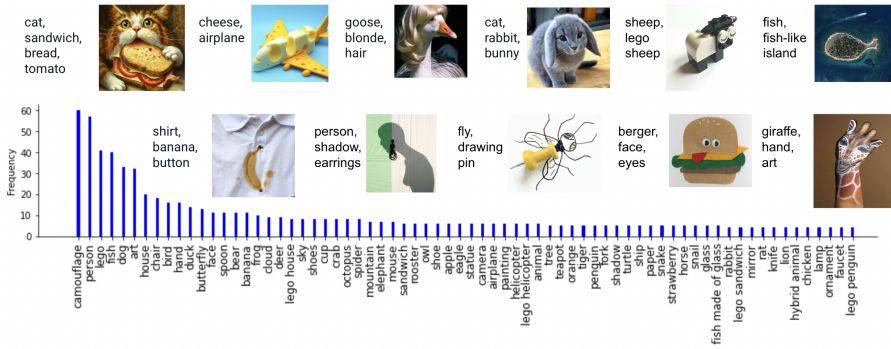


Figure 3: Sample images from the D2O miscellaneous category along with their tags. Bar plot shows the frequencies of the top 70 most frequent tags. See Appendices C and D for more stats.

D2O class	ImageNet class
clock	digital clock, wall clock
elephant	Indian elephant, African elephant
helmet	crash helmet, football helmet, gas mask, respirator, gas helmet
rabbit	wood rabbit, cottontail, cottontail rabbit, Angora, Angora rabbit
squirrel	fox squirrel, eastern fox squirrel, Sciurus niger
sun glass	sun glass, sunglasses, dark glasses, shades
turtle	loggerhead, loggerhead turtle, Caretta caretta, leatherback turtle, leatherback, leathery turtle, Dermochelys coriacea, mud turtle, box turtle, box tortoise

Table 1: The mapping between categories in our dataset and ImageNet categories.

We did not incorporate any bias towards gender, age, or race during data collection, and tried to be as inclusive as possible. Most of the categories are about objects. Few classes such as bicycle-built-for-two, face, helmet, person, sunglass, and umbrella contain humans and faces. We include and balance the number of images containing different ages and genders. The age groups are (child, 22), (teenager, 30), (adult, 51), and (elderly, 43). The gender groups include (woman, 66) and (man, 80). Notice that these issues are more important to address over large training sets. This is because sometimes models trained on such datasets are directly deployed in the real-world.

Our dataset contains 8,060 images spread across 36 categories, out of which 29 are in common with ImageNet. Six categories including {car, cat, cow, face, giraffe, person} do not appear in ImageNet, and are included mainly because they are very common and easily recognizable by humans. Seven of the categories correspond to multiple ImageNet categories, as shown in the mapping in Table 1. For a certain class, if a model predicted any of its corresponding ImageNet classes, we considered the prediction a hit. Sample images from our dataset are shown in Fig. 1. Distribution of object frequencies is shown in the left panel of Fig. 2. The most frequent class is the person followed by car and cat classes. Interestingly, there is a large variation of person in images as this topic has fascinated many artists over time (e.g. person made of wire, clay, matches).

The miscellaneous category. This category includes images that do not simply fall under a specific category and may cover multiple concepts (e.g. hybrid animals or strange objects). Thus, this category is suitable for testing image tagging algorithms. It contains 576 images covering a wide variety of concepts and topics including hybrid animals, hybrid objects, art, illusions, camouflage objects, out of context objects, shadow, animals, fruits, drawings, paintings, objects made from different materials (e.g. glass, metal, clay, cloud, tattoos, or Lego), impersonating objects, strange daily objects, and objects from odd viewpoints. Sample images alongside their tags are shown in Fig. 3. This figure also presents the tag frequencies. The top 10 most frequent tags include (camouflage, 60), (person, 57), (Lego, 41), (fish, 40), (dog, 33), (art, 32), (house, 20), (chair, 18), (bird, 16), (hand, 16), and (duck, 14).

To put our dataset in perspective with other datasets and for cross-dataset comparison, we also evaluate models over 11 classes of the ImageNet-A dataset that also exist in our dataset. Sample images from three of these categories are shown in the right panel of Fig. 2.

The D2O dataset is substantially different from ImageNet and ImageNet-A validation sets measured in terms of the Fréchet Inception Distance (FID) (Heusel et al., 2017). The FID between D2O and these sets in order are 45.2 and 51.3 indicating a large distribution shift. To put these numbers in




Figure 4: Per category performance of the models on our dataset (left) and on ImageNet-A dataset (right), averaged over 10 models. The dashed lines show the average performance over categories. See Appendix C for performance of individual models.




Figure 5: Per category performance of the best model (`resnext101_32x8d_ws`).

perspective, the FID between ImageNet’s validation and the test set is approximately 0.99. Notice that the lower the FID, the more similar the two distributions.

4 RESULTS AND ANALYSES

4.1 GENERIC OBJECT RECOGNITION

We tested 10 state-of-the-art object recognition models⁴, pre-trained on ImageNet, on our dataset. These models have been published over the past several years and have been immensely successful over the ImageNet benchmarks. They include AlexNet (Krizhevsky et al., 2012), MobileNetV2 (Sandler et al., 2018), GoogleNet (Szegedy et al., 2015), DenseNet (Huang et al., 2017), ResNext (Xie et al., 2017), ResNet101 and ResNet152 (He et al., 2016), Inception_V3 (Szegedy et al., 2016), Deit (Touvron et al., 2021), and ResNext_WSL (Mahajan et al., 2018). Details on accuracy computation are given in Appendix A.

Models are trained on ImageNet and tested only on the classes shared with ImageNet. Performance per D20 category, averaged over models, is shown in Fig. 4. The average performance, over all models and categories, is around 30% using Top-1 acc and around 50% using Top-5 acc. The corresponding numbers over the ImageNet-A dataset are about 5% and 15%, respectively. Therefore, ImageNet-A images are on average harder than D20 images for models, perhaps because they contain a lot of clutter. Prior research has shown that clutter and crowding severely hinder deep models (e.g. Volokitin et al. (2017)). It is not clear which object is the main one in most of the ImageNet-A images (See Fig. 2). To pinpoint whether and how much clutter contributes to low performance on this dataset, we manually cropped the object of interest in images (the blue bounding boxes in Fig. 2). The cropped objects have low resolution, but they are still recognizable by humans. Results on this dataset, called ImageNet-A-Isolated, will be discussed in the following.

Among the models, `resnext101_32x8d_ws`⁵ ranks the best over both datasets, as shown in Fig. 5. It achieves around 60% Top-1 accuracy, which is much higher than its Top-1 acc over the

⁴Models are available in PyTorch hub: <https://pytorch.org/hub/>. We used a 12 GB NVIDIA Tesla K80 GPU to conduct the experiments.

⁵This model scores 85.4% (97.6% top-5) over ImageNet-1k validation set (single-crop).

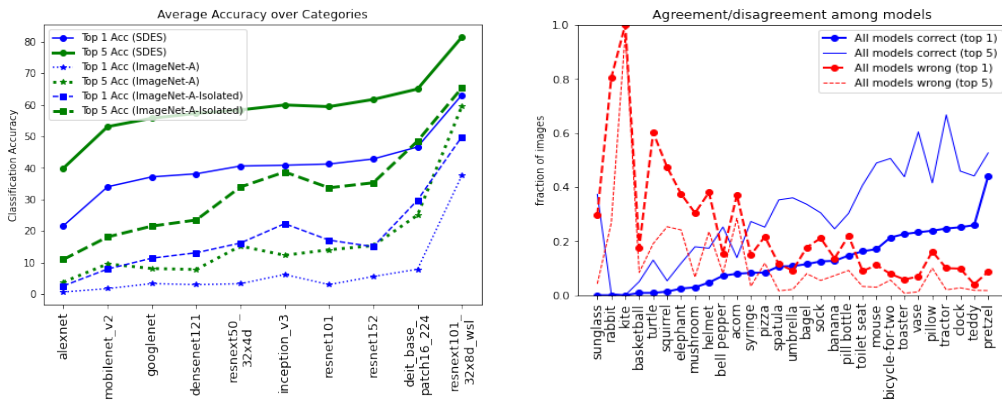


Figure 6: Left: Performance of models averaged over categories (29 categories of D2O and 11 categories of ImageNet-A). Right: Fraction of images over which all models fail or they all succeed.



Figure 7: Sample failures of resnext101_32x8d_ws model over the kite (top) and other categories. Model predictions are overlaid on images. The confidence for the top-1 prediction is also shown. Please see also Appendix F.

ImageNet-A dataset (~40%). The Top-5 acc of this model on our dataset is about 80% compared to its 60% over ImageNet-A. The success of this model can be attributed to the fact that it is trained to predict hashtags on billions of social media images in a weakly supervised manner. The best performance on our dataset is much lower than the best performance on the ImageNet validation set. The best Top-1 and Top-5 performance over the latter are about 91% and 99%, respectively⁶. We find that better performance on ImageNet translates to better performance on D2O.

Performance per model, averaged over D2O categories, is shown in the left panel of Fig. 6. We observe a big difference between accuracy of resnext101_32x8d_ws model vs. other models. This model is about 20% better than the second best model deit_base_patch16_224, using Top-1 accuracy. The former model is based on weakly supervised learning whereas the latter is a transformer-based model. See Appendix for performance of individual models.

As shown in Fig. 6, models perform much better over the Isolated ImageNet-A dataset than the original ImageNet-A, even though the former has low resolution images due to region cropping. This supports our argument that lower performance on ImageNet-A dataset is partly due to its scenes being cluttered. All categories enjoy an improvement in accuracy (See Appendix C).

We also test the SwinTransformer (Liu et al., 2021)⁷ on D2O. This model has improved the state-of-the-art over several computer vision problems including object detection, semantic segmentation, and action recognition. It scores 58.2% (Top-1) and 76.1% (Top-5). It performs much better than ResNet, Inception and Deit models, but is slightly below the resnext101_32x8d_ws model.

⁶<https://paperswithcode.com/sota/image-classification-on-imagenet>

⁷<https://github.com/microsoft/Swin-Transformer>




Figure 8: Sample face images along with predictions of OpenCV (Red), Microsoft API (Green), MEGVII Face++ API (Blue), and Google MediaPipe (Black) face detectors.

Illustrative Failure Modes. According to Fig. 4, the top five most difficult D2O categories for all models in order are *kite*, *rabbit*, *squirrel*, *turtle*, and *mushroom* which happen to be the most difficult categories for the best model as well (Fig. 5). The kite class is often confused with parachute, balloon, and umbrella classes. Sample failure cases from the categories along with the predictions of the `resnext101_32x8d_ws` model are shown in Fig. 7. Models often fail on drawings, unusual objects, or images where the object of interest is not unique. We also computed the fraction of images, per category, over which all models succeed, or they all fail, as shown in the right panel of Fig. 6. For some categories models consistently fail (e.g. kite, rabbit, turtle, squirrel), while for some others they all do very well (e.g. toaster, tractor, pretzel). When all models succeed, they are correct at best over 30% of the images (toaster). This result indicates that models share similar weaknesses and strengths.

4.2 PERFORMANCE OF VISION APIS

We also tested several APIs from **Microsoft**⁸, **Google**⁹, and **MEGVII**¹⁰ over the D2O categories that do not exist in ImageNet: {*face*, *person*, *car*, *cat*, *cow*, *giraffe*}. These APIs are very popular and highly accurate. The goal here is to see how models behave beyond ImageNet.

Face Detection. D2O face category has 289 images and includes a lot of odd and difficult faces, some are shown in Fig. 8. We are mainly interested in finding whether a model is close enough in detecting faces. To this end, we refrain from using mAP to evaluate the APIs and use the accuracy score, which is easier to understand and interpret. An image is considered as a hit if the API is able to generate a bounding box with IOU greater than or equal to 0.5 with a face in the image. Otherwise, the image is considered a mistake. We also manually verified all the predicted boxes. Our evaluation is an overestimation of performance rather than being a strict benchmark. Over the images for which the APIs failed, most often the predicted boxes did not overlap with any face in the image. In the majority of the mistakes, though, the face was missed.

Even with the above relaxed evaluation, APIs did not do well. Microsoft Azure face detection API achieves 45.3% accuracy in detecting D2O faces. The MEGVII Face++ API achieves 23.9% accuracy, slightly above the 23.2% by OpenCV face detector. Google MediaPipe face detector achieves 50.5% accuracy. Sample face images and predictions of the APIs are shown in Fig. 8.

Person Detection. MEGVII API person detector obtains about 16% accuracy over the 1,217 person images in the person category. OpenCV person detector achieves 5% accuracy. Microsoft object detection API achieves 27.4% accuracy. If this API predicted a correct bounding box with any of the following classes {*person*, *snowman*, *bronze sculpture*, *sculpture*, *doll*}, we counted it as a hit. Evaluation is done the same way as in face detection.

Cat Detection. Over the cat category (407 images), Microsoft object detection API, predicted 95 images as cat ($95/407=0.23$), 9 images as Persian cat ($9/407=0.022$), 26 images as animal ($6/407 =$

⁸<https://azure.microsoft.com/en-us/services/cognitive-services/>

⁹<https://google.github.io/mediapipe/>

¹⁰<https://www.faceplusplus.com/face-detection/>




Figure 9: Sample images from the person, cat, cow, giraffe, and car categories (row-wise) along with predictions of Microsoft object detector. For the person category, the blue and green boxes represent Microsoft and MEGVII person detectors, respectively.

0.064), and 95 images as mammal (95/407= 0.23). Considering all of these images as hits, this API achieves 54.8% accuracy.

Cow Detection. Over the cat category (407 images), Microsoft object detection API, predicted 95 images as cat (95/407=0.23), 9 images as Persian cat (9/407=0.022), 26 images as animal (6/407 = 0.064), and 95 images as mammal (95/407= 0.23). Considering all of these images as hits, this API achieves 54.8% accuracy.

Giraffe Detection. Over the giraffe category (138 images), Microsoft object detection API predicted 30 images as giraffe (30/138=0.217), 15 images as animal (15/138=0.108), and 46 images as mammal (46/138=0.333). Considering all of these images as hits, this API achieves 65.9% accuracy.

Car Detection. Microsoft object detection API achieves 25.8% accuracy (139 out of 539) on this category. An image was considered a hit if the API predicted a correct bounding box with any of these labels {car, land vehicle, all terrain vehicle, taxi, vehicle, race car, limousine, Van, station wagon}.

On the one hand, our investigation shows that APIs perform very poorly, even considering overestimated accuracy, over categories that are very easy for humans. On the other hand, it reveals that our test set is challenging for a large array of models trained on a variety of datasets. Sample images from the above categories and predictions of the APIs on them are shown in Fig. 9.

4.3 TAGGING RESULTS

We used the Microsoft tagging API to annotate the 576 images in the miscellaneous category. For 46.4% of the images, there is a common tag between predicted tags and ground truth tags (calculated for each image and then averaged over images). For the remaining 53.6% of images, there is no overlap between the two sets. The fractional overlap between predicted and GT tags per image, computed as the number of common tags over the number of GT tags, is 21.8%. The Google Vision API¹¹ performed slightly worse than the Microsoft API. Results are shown in Table 2 and Fig. 10.

API	overlap (%)	no overlap (%)	no prediction	fractional overlap
MSFT Tagger	46.4	53.6	0	21.8
" Detector	8.5	36.3	55.2	17.1
" Recognizer	0	91.3	8.7	0
" Captioning	30.9	69.1	0	13.8
Google Tagger	39.9	60.1	0	17.1

Table 2: Image tagging performance of APIs over the miscellaneous category of our dataset. “No prediction” column shows the percent of images for which there was no prediction.

¹¹<https://cloud.google.com/vision>

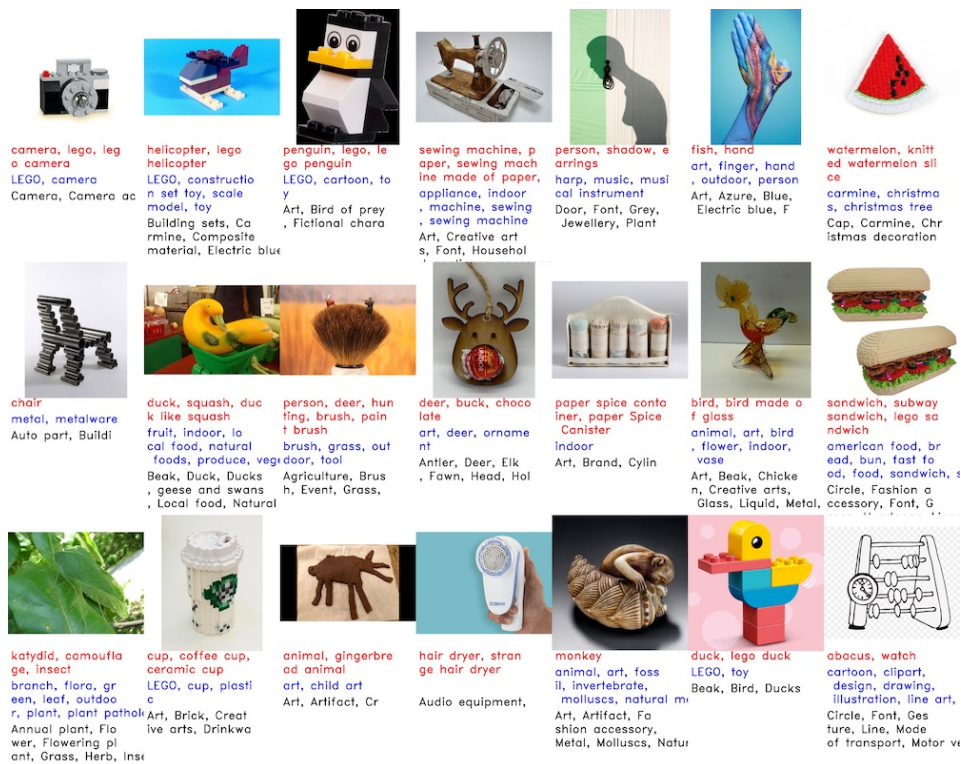


Figure 10: Sample images from the miscellaneous category along with ground truth tags (red) and predictions of the Microsoft (blue), and Google (black) tagging APIs.

We also used Microsoft detection, recognition, and captioning APIs as image taggers by considering their generated labels or words as tags. Using the detection API, for 8.5% of the images, there was at least one tag in common between the detected label set and the ground truth tags. For 36.3% of the images, there was no overlap at all. For the remaining 55.2% of images, the API did not detect anything. The fractional overlap between the predicted and ground-truth tags is 17.1%. Using the recognition API, 91.3% of images had no overlap and the remaining 8.7% had no prediction at all. Using the Microsoft captioning API, for 30.9% of the images there was an overlap between predicted and ground-truth tags (69.1% had no overlap). The fractional overlap between the two sets is 13.8%. Overall, the tagging APIs perform better in tagging images than APIs that are not tailored for this task, but all of them still perform poorly in tagging images. Please see Table 2.

5 DISCUSSION AND CONCLUSION

We introduced a new test for object recognition and evaluated several models and APIs on it. We find that, despite years of research and significant progress in this field, there is still a large gap in performance of models over our dataset vs. ImageNet.

Some datasets rooted in ImageNet are biased towards borrowing its images and classes, or its data collection strategy. Here, we intended to deviate from these biases. For example, unlike ImageNet-A, our dataset is model independent. ImageNet-A contains images from ImageNet for which models fail. Our dataset also includes categories, such as cows or cats, that perhaps everyone can easily recognize. These categories are missing in the ImageNet-based datasets. Further, our work encourages researchers and small teams to build carefully-curated, small-scale and versatile test sets frugally. Presently, the mindset is that datasets can only be collected by large institutions since data collection and annotation is difficult and expensive.

We share the dataset and code to facilitate future work, and will organize an annual challenge and associated workshop to discuss state-of-the-art methods and best practices. To stop or limit the misuse of our D2O by bad actors, we have made a dataset request form¹². We review the requests that we receive and allow access for a legitimate use. The dataset we share contains images and questions in a zip file. The package also contains the detailed documentation with all relevant metadata specified to users. Our dataset is licensed under Creative Commons Attribution 4.0 (Appendix H).

¹²<https://bit.ly/3JxIKh3>

REFERENCES

- Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In *Advances in Neural Information Processing Systems*, pp. 9448–9458, 2019.
- Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020.
- Ali Borji. Contemplating real-world object classification. *arXiv preprint arXiv:2103.05137*, 2021.
- Ali Borji, Saeed Izadi, and Laurent Itti. ilab-20m: A large-scale controlled object dataset to investigate deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2221–2230, 2016.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Jacob Steinhardt, and Aleksander Madry. Identifying statistical bias in dataset replication. In *International Conference on Machine Learning*, pp. 2922–2932. PMLR, 2020.
- Robert Geirhos, Carlos RM Temme, Jonas Rauber, Heiko H Schütt, Matthias Bethge, and Felix A Wichmann. Generalisation in humans and deep neural networks. In *Advances in neural information processing systems*, pp. 7538–7550, 2018.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11): 665–673, 2020.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. *arXiv preprint arXiv:2006.16241*, 2020.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15262–15271, 2021.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Sara Beery, et al. Wilds: A benchmark of in-the-wild distribution shifts. *arXiv preprint arXiv:2012.07421*, 2020.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

- Felix Lau, Nishant Subramani, Sasha Harrison, Aerin Kim, Elliot Branson, and Rosanne Liu. Natural adversarial objects. *arXiv preprint arXiv:2111.04204*, 2021.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 181–196, 2018.
- Anibal Pedraza, Oscar Deniz, and Gloria Bueno. Really natural adversarial examples. *International Journal of Machine Learning and Cybernetics*, 13(4):1065–1077, 2022.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*, 2019.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4510–4520, 2018.
- Vaishaal Shankar, Rebecca Roelofs, Horia Mania, Alex Fang, Benjamin Recht, and Ludwig Schmidt. Evaluating machine accuracy on imagenet. In *International Conference on Machine Learning (ICML)*, 2020.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *arXiv preprint arXiv:2007.00644*, 2020.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pp. 10347–10357. PMLR, 2021.
- Anna Volokitin, Gemma Roig, and Tomaso A Poggio. Do deep neural networks suffer from crowding? In *Advances in Neural Information Processing Systems*, pp. 5628–5638, 2017.
- Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 3485–3492. IEEE, 2010.
- Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6): 1452–1464, 2017.

A MEASURING ACCURACY

Since some of our classes cover multiple ImageNet classes¹³, we had to make some adjustments for computing accuracy. For example, ImageNet has three types of clocks including ‘digital clock’, ‘wall clock’, and ‘analog clock’. Here, we only have the ‘clock’ class, containing mostly analog clocks. We chose to give the benefit of the doubt to models. A prediction is correct if the ground-truth label is in the set of the words predicted by the model. In the mentioned scenario, if a model predicts ‘wall clock’, then a hit is counted. If the model predicts ‘wall’ or anything else, then the prediction would be considered a mistake. The same is true for the top-5 accuracy computation. For example, if the top five model predictions are ‘bib’, ‘necklace’, ‘toilet seat’, ‘pick’, ‘wall clock’, then the prediction is counted as a hit. In practice, first all words in the predicted labels are extracted, and then the prediction is counted as a hit if the ground-truth is in this set. In case of ground-truth having two words (*e.g.* ‘toilet seat’), then it should happen in the set of words exactly as it is.

Notice that this way of accuracy measurement gives an overestimation of the model performance, but it is good enough for our purposes here. Even with this overestimation, as we will show, models still perform poorly.

¹³<https://deeplearning.cms.waikato.ac.nz/user-guide/class-maps/IMAGENET/>

B SAMPLES IMAGES FROM THE IMAGENET-A AND ISOLATED IMAGENET-A DATASETS

The rows correspond to acorn, banana, basketball, spatula, teddy, and toaster categories. The blue bounding box shows the region that we cropped to construct the Isolated ImageNet-A dataset.



Figure 11: Some sample images from the ImagenetA dataset. The blue box shows the isolated region.

C CLASSIFICATION ACCURACY OF INDIVIDUAL MODELS



Figure 12: Performance of models over D20 dataset.




Figure 13: Performance of models over ImageNet-A dataset.




Figure 14: Performance of models over Isolated ImageNet-A dataset.

D FREQUENCY OF GROUND-TRUTH TAGS AND MODEL PREDICTED TAGS




Figure 15: Frequency of ground-truth tags and model predicted tags. From top to bottom: gt tags and predicted tags by a) Microsoft API, b) Google API, c) MSFT object detector, d) MSFT object classifier, and e) MSFT image captioning.

E SAMPLE FAILURE CASES OF BEST MODEL (RESNEXT101_32x8D_WS)
OVER D2O DATASET



Figure 16: Predictions of the resnext101_32x8d_wsl model over spatula category.




Figure 17: Predictions of the resnext101_32x8d_wsl model over helmet category.




Figure 18: Predictions of the resnext101_32x8d_wsl model over sunglasses category.

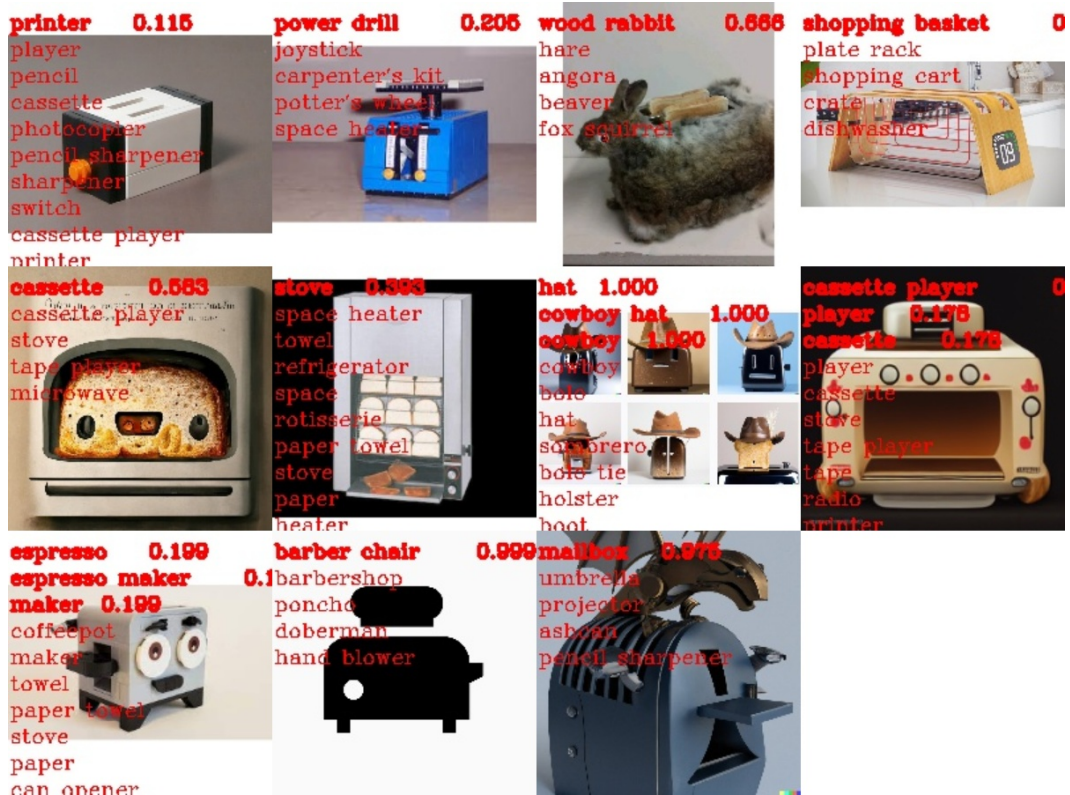


Figure 19: Predictions of the resnext101_32x8d_wsl model over toaster category.

F SAMPLE FAILURE CASES OF BEST MODEL (RESNEXT101_32x8D_WS)
OVER IMAGENET-A DATASET

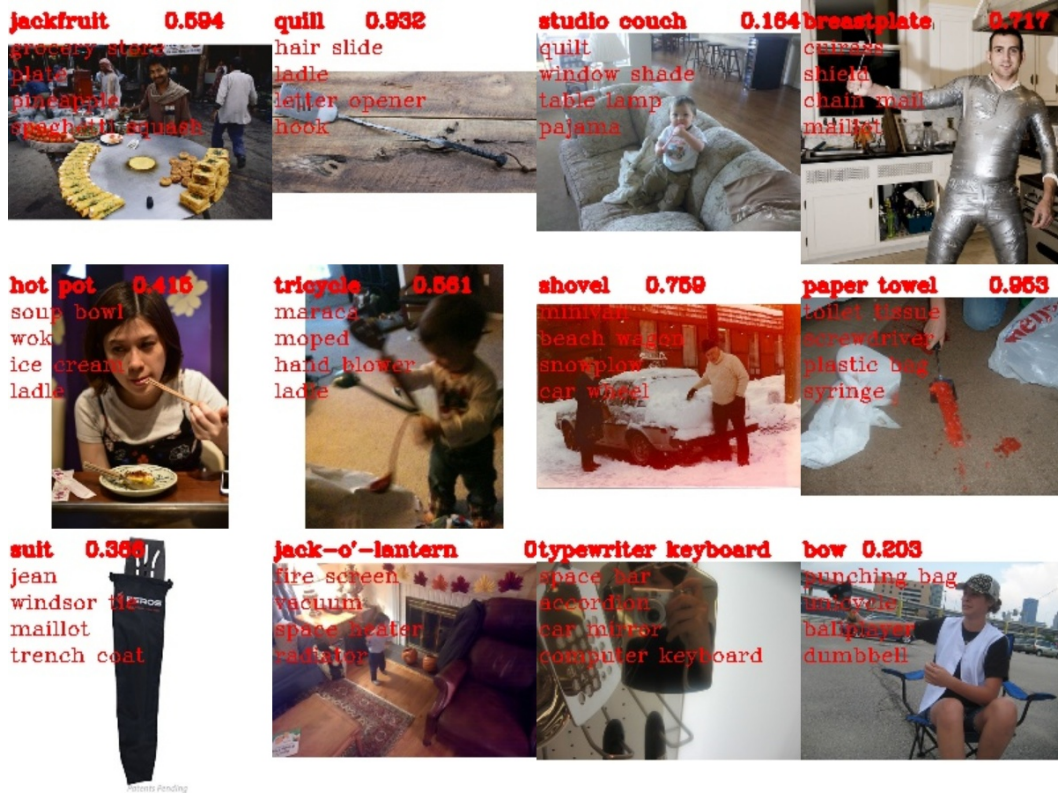


Figure 20: Predictions of the resnext101_32x8d_wsl model over spatula category.



Figure 21: Predictions of the resnext101_32x8d_wsl model over toaster category.



Figure 22: Predictions of the resnext101_32x8d_wsl model over bell pepper category.



Figure 23: Predictions of the resnext101_32x8d_wsl model over mushroom category.




Figure 24: Predictions of the resnext101_32x8d_wsl model over acorn category.



Figure 25: Predictions of the resnext101_32x8d_wsl model over basketball category.

G SAMPLE FAILURE CASES OF BEST MODEL (RESNEXT101_32x8D_WS)
OVER SISOATED IMAGENET-A DATASET




Figure 26: Predictions of the resnext101_32x8d_wsl model over spatula category.



Figure 27: Predictions of the resnext101_32x8d_wsl model over toaster category.



Figure 28: Predictions of the resnext101_32x8d_wsl model over bell pepper category.




Figure 29: Predictions of the resnext101_32x8d_wsl model over mushroom category.



Figure 30: Predictions of the resnext101_32x8d_wsl model over acorn category.



Figure 31: Predictions of the resnext101_32x8d_wsl model over basketball category.

H DATASET LICENSE

D2O dataset is free to use only for research and academic purposes (not commercial). It is licensed under Creative Commons Attribution 4.0 with three additional clauses:

1. D2O may never be used to tune the parameters of any model.
2. The images containing people should not to be posted anywhere unless the people in the images are appropriately de-identified. Even in this case, written agreement from dataset creators is required. This is to check whether all the clauses are properly followed.

To stop or limit the misuse of our D2O by bad actors, we have made a dataset request form¹⁴. We review the requests that we receive and allow access for a legitimate use. The dataset we share contains images and questions in a zip file. The package also contains the detailed documentation with all relevant metadata specified to users.

¹⁴<https://bit.ly/3bDY0MS>

I SAMPLE GENERATED IMAGES




Figure 32: Sample generated images by Midjourney from Pillow, Person, Pizza, Banana, Helmet, Toaster, Mushroom, and Rabbit classes.