

# YouTube-BoundingBoxes: A Large High-Precision Human-Annotated Data Set for Object Detection in Video

Esteban Real  
Google Brain  
ereal@google.com

Jonathon Shlens  
Google Brain  
shlens@google.com

Stefano Mazzocchi  
Google Research  
stefanom@google.com

Xin Pan  
Google Brain  
xpan@google.com

Vincent Vanhoucke  
Google Brain  
vanhoucke@google.com

## Abstract

*We introduce a new large-scale data set of video URLs with densely-sampled object bounding box annotations called YouTube-BoundingBoxes (YT-BB). The data set consists of approximately 380,000 video segments about 19s long, automatically selected to feature objects in natural settings without editing or post-processing, with a recording quality often akin to that of a hand-held cell phone camera. The objects represent a subset of the COCO [32] label set. All video segments were human-annotated with high-precision classification labels and bounding boxes at 1 frame per second. The use of a cascade of increasingly precise human annotations ensures a label accuracy above 95% for every class and tight bounding boxes. Finally, we train and evaluate well-known deep network architectures and report baseline figures for per-frame classification and localization to provide a point of comparison for future work. We also demonstrate how the temporal contiguity of video can potentially be used to improve such inferences. The data set can be found at <https://research.google.com/youtube-bb>. We hope the availability of such large curated corpus will spur new advances in video object detection and tracking.*

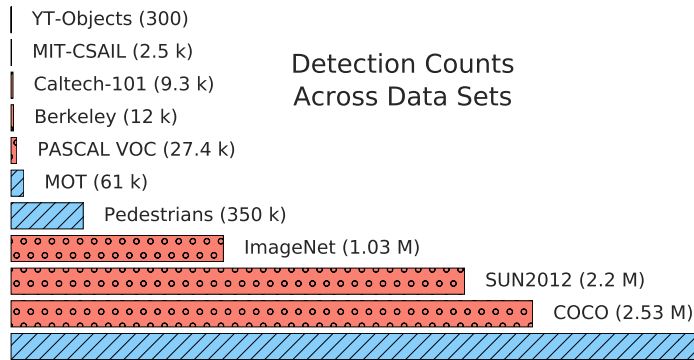
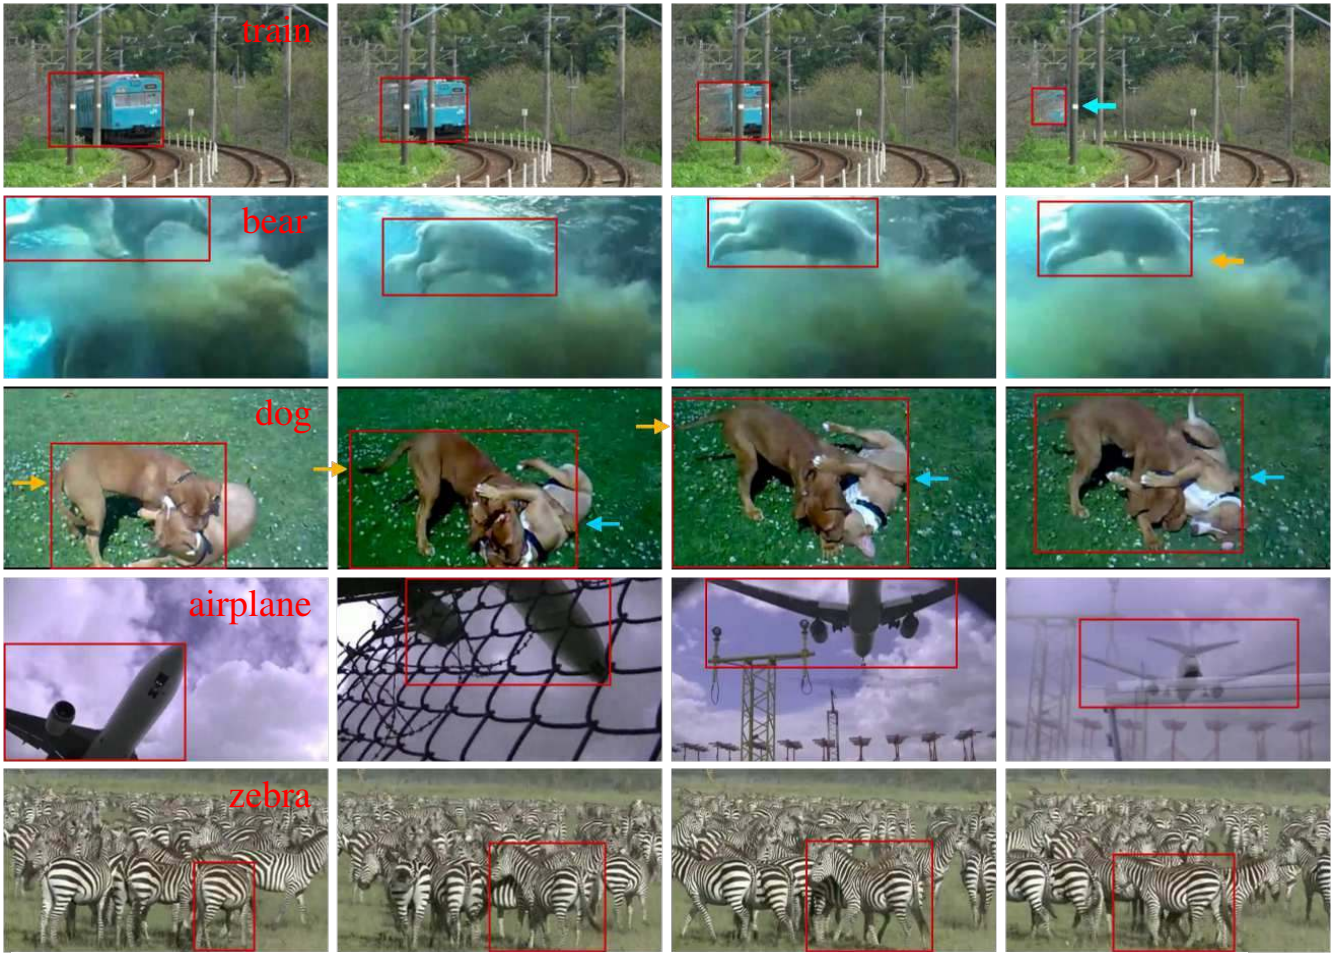
## 1. Introduction

The exceptional pace of progress in recent years on the tasks of object recognition and detection in still images was enabled by the creation of large-scale, publicly available data sets [9, 13, 14, 19, 32, 33, 34, 42, 52, 57]. These data sets established challenging benchmarks to evaluate new methods for visual object recognition that have substantially improved the state-of-the-art across a broad range of computer vision tasks [20, 28, 46, 48, 49].

Open academic challenges paired with open-source recipes have further accelerated the development of the field [6, 43, 51]. Most notably, systems that perform well on image recognition and object detection may be applied to other computer vision problems in which minimal training data is available [24]. Such systems have also become part of larger machine learning pipelines that stretch beyond visual recognition (e.g. multi-modal learning [36, 53]).

The increased speed and memory of modern computing architectures places the research community in a position to aim for comparable results in video, a natural goal for machine perception. The quest for large video data sets, however, has been more elusive. One challenge is that the online corpus of videos is *weakly labeled*, i.e. the label information is very noisy [50]. Sifting through a large sample may therefore require considerable human involvement.

Exacerbating this problem is the recognition that large data sets are necessary to prevent over-fitting of cutting-edge models (e.g. [35, 45]). Although the temporal dimension provides vastly more data, much of the information is redundant due to correlations of pixels across frames. Thus, increasing the data set size is not merely about gathering more sequential frames from a small number of videos. Instead, we need a large, diverse sample of videos. Attaining it requires paying special attention to how the videos are mined. Some of the larger existing vision data sets rely indirectly on aggregate measurements of human preference [18]. Consequently, those data sets favor aesthetically pleasing viewpoints of labeled objects. This leads to object recognition systems that are precise but may lack variety in terms of realistic lighting conditions, occlusions or the non-canonical viewpoints often observed in real life. Video may be less prone to some of these biases (especially the viewpoint bias), but a random YouTube sample would still suffer drastically from them. Mining videos with diversity in mind, on the other hand, can address this problem explic-



YT-BB Counts and Statistics

	Class.	Boxes	Videos	Motion
person	1.8 M	1.3 M	68 k	0.122
dog	560 k	240 k	10 k	0.165
train	340 k	240 k	8.9 k	0.072
		...		
NONE	26 k	-	-	-
TOTAL	9.5 M	5.6 M	240 k	-

Figure 1. IMAGES: Detection examples. Each row shows frames from one video segment. A frame containing an object whose identity can be deduced from other frames is boxed too, as in the last frame of the *train example* (blue arrow). Note how only visible parts are included in the box: the orange arrow in the *bear example* points to the hidden head. The *dog example* illustrates tight bounding boxes tracking the tail (orange arrows) and foot (blue arrows). The *airplane example* shows how partial objects are annotated (first frame) and how objects are tracked across changes in perspective, occlusions and camera cuts. Note in the *zebra example* how the same object is tracked across multiple frames and other objects of the same class are ignored. BAR CHART: Number of detections/segmentations in various image (red; dots) and video (blue; lines) data sets. The present data set, YT-BB, is at the bottom. TABLE: Human annotation statistics for some classes in YT-BB. The first three columns are counts for: classification annotations, bounding boxes, and unique videos with bounding boxes. Negative classifications are not counted here. The “Motion” column shows the RMS of the distance the box travels from one frame to the next, in frame-relative coordinates, to show that the objects exhibit significant motion (see Section 4.2 for details and Supplementary Table 3 for other measures of movement). The “NONE” tag annotates frames that did not have any of the 23 classes.

itly.

The persistence and temporal consistency of objects present in natural video scenes call for a different kind of labeling, whereby objects of interest are tracked across frames and precisely localized. To this day, there is no human-curated large-scale data set that provides classification and detection annotations for objects of several classes in a wide variety of videos.

This work attempts to address this issue by providing a large body of video annotations with manually curated bounding boxes of objects tracked for relatively long durations on the order of 100 frames. The size of the data set makes it suitable for training large deep neural networks and explore visio-temporal modeling approaches in a realistic setting.

## 2. Related work

Several video data sets are already available to the community. Below are some of the most relevant, highlighting how they differ from YouTube-BoundingBoxes:

- Many data sets such as the HMDB-51 data set [29] and the UCF-101 data set [47] provide segment-level annotations for a variety of human action categories; richer annotations including fine-grained temporal and localization information were provided as part of the THUMOS [17] challenge.
- TRECVID [3] is a yearly set of competitions centered on video retrieval and indexing, hosting a variety of video data sets. For 2016, they provide a localization test set with 1000 videos annotated with bounding boxes for 10 classes; each video may or may not contain a box.
- VOT [27] and MOT [30] are yearly visual object tracking challenges with associated data sets. These are small and extensively curated in order to provide controlled frequencies of various common difficulties for object tracking (such as a occlusions, illumination changes or size changes).
- The Sports-1M data set [25, 26] consists of segment-level annotations for a variety of sports, with temporal localization.
- The YouTube-Objects data set [38, 39] consists of a number of frames queried from YouTube with a few hundred curated bounding box annotations.
- The Caltech Pedestrian Detection data set [10] consists of 350,000 bounding boxes of pedestrians annotated from a vehicle driving through an urban environment.
- The YouTube-8M data set [2] consists of a very large set of frame-level, automatically generated annotations of YouTube videos. The labels were generated using state-of-the-art deep networks to classify thousands of possible entities.
- ImageNet 2015 [43] has a video object detection data set with 5,400 videos.

Still-image detection data sets are larger and more abun-

dant. They vary in detail from bounding boxes (Caltech-101 [14], MIT-CSAIL [52], ImageNet [9, 42], PASCAL VOC [12, 13], SUN2012 [56, 57]) to pixel-level segmentations (Berkeley Segmentation Data Set [33], Caltech-101 [14], PASCAL VOC [12, 13], Microsoft COCO [32]).

The bar chart in Figure 1 puts our data set in context: YT-BB is the largest human-annotated detection data set in existence so far. Specifically for the case of video, it exceeds other data sets in size by more than an order of magnitude.

## 3. Methods

### 3.1. Data mining

In order to provide a low entry-bar to video for researchers that have models pre-trained on static images, we chose as our labels 23 classes that form a subset of the detection classes in the COCO data set [32]. Due to its particular importance, we included the “person” class and gave it preferential treatment in terms of total volume and in terms of how the videos were mined (details below). The other classes are all common objects or animals (first column in Supplementary Table 1).

Many academic data sets are made artificially easy compared to real-world problem settings because they have a closed set of labels to choose from, whereas most data collected “in the wild” can’t be expected to correspond to a well-defined category. This is particularly important for detection and localization tasks. To directly address this problem, we added a “NONE” class that marks frames that do not have any of the 23 object classes.

We sampled public YouTube videos and used object-agnostic signals to reduce the set obtained to a size suitable for human annotation. We calculated an estimate of the entropy across frames and removed those below a particular threshold, reducing the frequency of slide shows and other videos with minimal motion. Requiring that videos have fewer than 100 views notably reduced the number of professionally edited clips. More generally, this limit on the view count helps protect against the bias of a plain internet search result, which would yield preferentially videos that are likable (good lighting, centered characters, stable cameras, *etc.*) Finally, a camera-cut detector helped remove videos that had unusually short scenes, which are indicative of a high degree of post-processing. We then split the remaining videos into short non-overlapping clips (mean length = 18.7s, sd = 1.00s). All these restrictions together resulted in a collection of video segments typical of what a hand-held camera would record in a natural setting.

This data mining procedure proved satisfactory for the “person” class, but was too inefficient for classes that occur infrequently in the YouTube corpus. To compensate, we ran image classifiers at 1 frame per second across our video sample. We retained the top 1 million videos, discarding

those deemed by the classifiers as too unlikely to contain any of our 23 classes<sup>1</sup>. When possible, we exploited the WordNet hierarchy [15] to associate multiple fine-grain image labels with a given class label in YT-BB.

The “person” class is especially important. In particular, the research community has a vested interest in detecting people in videos. While our initial approach of sampling random YouTube videos for this class may provide a fairly unbiased data set, it may also produce one that is too homogeneous. The most popular videos (like a recent music album) are not a problem because they were removed by the initial object-agnostic filters. However, there is still a class of videos that may be *filmed* very frequently even if they are not viewed many times by other users. This would include the sort of things most of us care about, like birthday parties, graduation ceremonies, and the like. As an attempt to compensate for that, we enriched our random sample of “person” videos with a comparable yet smaller number of videos mined from entities that correlated with “person” well (“person”, “bicycle”, “crowd” and—surprisingly—“elephant” are examples). Finally, we intentionally mined a disproportionately large number of videos for this class with the goal that the “person” subset of our data set may stand on its own.

Balancing the time spent on human annotation and the yield required focusing on segments that usually contain only one class. We felt this was preferable to the huge sacrifice in volume that would have been necessary to label segments containing multiple classes. Namely, mining for videos with several classes results in a much lower yield and the alternative of mining them with higher recall produces too many false negatives which in turn increases the human annotation time too much.

### 3.2. Human annotations

Like other large data sets before us [9, 32], we used human annotation pipelines to label our data. As in [32], we set up a cascade of stages that successively refine the quality of the results. This strategy is standard [7] and has been found to improve results [4]. We used four stages:

1. Five frames from each ( $\sim 19$  s) video segment, evenly sampled in time, were simultaneously presented to one human rater (i.e. annotator), who had to determine whether a specific class was present in any of them. Negative segments were discarded.
2. Each full segment was presented to three different annotators as a “movie-roll”, sampled at 1 frame per second. The

<sup>1</sup>We intentionally ran the image classifiers at low thresholds for each class in order to avoid the pitfall of selecting easy-to-classify examples. Specifically, we ranked candidate segments according to the confidence of the classifiers and set the threshold for a given image classifier to operate at low rates of precision, as judged by one-off experiments. Our selection of threshold had the goal of leaving plenty of work for human annotators to do as far as discriminating the presence or absence of each class.

annotators had to indicate whether the class was present in each frame. The majority vote produced our (intermediate) classification data set. To find frames for the “NONE” class, we asked the raters explicitly about each of the 23 classes to ensure they were absent. Such annotations are precise but very time-consuming, and so the frequency of the “NONE” class is limited (see table in Figure 1 or Supplementary Table 1). Segments with at least one positive frame for a given class were used in stages 3 and 4.

3. For each segment, a single human annotator overlaid a bounding box tightly around an object of the given class in each of the frames in the segment, at 1 frame per second. Every appearance of a single object was annotated throughout the segment. (Other objects of the same class were to be ignored). The annotator also had the option of assigning an *absent tag* to a frame if the object could not be seen there. To resolve corner cases, they followed the guidelines in Supplementary Section 3, which address issues of box tightness, partial objects, occlusion, *etc.* Incidentally, these rules may help readers clarify peculiarities of our data set.

4. Each annotation from stage 3 was verified by one (training and validation data sets) or three (testing data set) different human annotators. Boxes or absent-tags with negative majority votes were discarded.

We employed *Amazon Mechanical Turk* for the first two stages of human annotation as in [9, 32]. This allowed for quick progress [5], yet suffered from the widely known drawbacks of crowd-sourcing, including difficulty motivating raters and poor quality of individual annotations [23]. This can be partly curbed through replication [37, 44], as we did in stage 2. While there exist more sophisticated methods for analyzing replicated labels [37, 44, 55], we opted for the majority vote because it was simple, we only had 3 labels per example, and the data was going to be further refined by stages 3 and 4 anyway. In order to harness the benefits of annotator training [11], for stages 3 and 4 we switched to our internal human annotation system. We employed human raters that read a written manual describing the task in detail and went over it during class sessions. During the annotation process, they were also able to escalate questions when they felt unsure about corner cases.

Another important aspect of human computation is the annotator’s interaction with the data. We designed user interfaces (Supplementary Section 1) in keeping with the principle that “the simpler, the better” [16]. Especially for *Mechanical Turk* tasks, it was important to phrase the questions well, striking a balance between reducing ambiguity and keeping the operator’s attention (details in Supplementary Section 2).

To fine tune the pipeline, we frequently inspected the data by eye. Stages 1 and 2 were finalized only after the resulting classification accuracy was estimated to be above 0.95 for each class. Stages 3 and 4 were optimized by giving

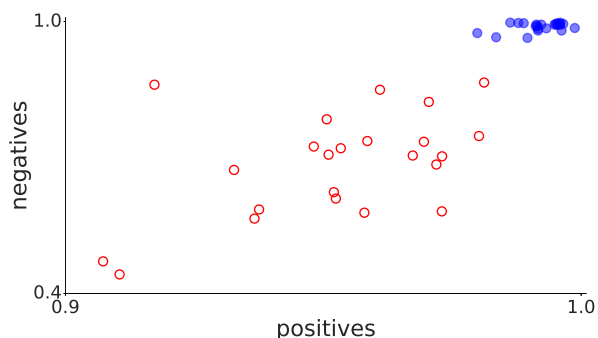


Figure 2. Unanimous fraction per class. The fraction of answers for which the operators voted unanimously, for classification (empty red circles) and detection (filled blue circles). The axes distinguish positives from negatives. For the classifications, this is based on the frame-level annotations, which were done by three human annotators. For the detections, the bounding box drawing stage was done by only one annotator, so this is based on the verification stage of the testing data set, which was done by three annotators too.

feedback to annotators based on (i) examples where stage 4 showed the most disagreement and (ii) randomly sampled examples. As we increased the size of the annotation batches, rarer examples of type (i) appeared. The quality seemed to improve with annotator experience too, so the validation and testing subsets were done last. The resulting quality after stage 4 is discussed below.

## 4. Results

### 4.1. Data set size

This process yielded a data set of 5.6 million frames annotated with bounding boxes from 240,000 unique YouTube videos. We also provide additional absent detection tags in 1 million frames from 55,000 unique videos. A superset of those videos contain classification annotations too: 9.6 million positives and 1 million negatives, with a similar distribution over unique videos. This is presented in detail in Supplementary Tables 1 and 2 and succinctly in the table in Figure 1, together with some examples.

### 4.2. Quality assessment

In order to assess the quality of the classifications, we measured the fraction of answers that were unanimous. Stage 1 strongly biases our sample toward positives (Section 3.2), which results in a higher fraction of false negative classifications. The use of untrained, unvetted raters also seriously reduces the accuracy of the answers (Figure 2). While this could be improved, our main goal of classifying the videos is to filter them in order to draw the bounding boxes, and so we did not optimize stages 1 and 2 further. Nevertheless, we make them available in our data set.

For the detections, we asked raters to verify the bounding boxes (or the absent tag). The frequency of correct verifications is an indication of the quality of the boxes. By this measure, each class had at least 98% correct bounding boxes and at least 98% correct absent-tags. In the case of the testing data set (for which we employed three raters), we can consider the harsher criterion of requiring a unanimously correct verification vote (instead of just a majority-correct verification vote): this gave that both, boxes and absent-tags, are still at least 98% correct for most classes and all classes are above 95% correct (Figure 2).

Annotation quality aside, a concern is that the objects in the videos exhibit movement. Otherwise, the data set would be equivalent to static images. We measured the RMS of the distance the center of the bounding boxes travels from one frame to the next and found that there is indeed significant motion. A few values are quoted in the table in Figure 1. Results for all classes can be found in Supplementary Table 3. Other statistics are also listed there, such as the fractional size change of the box per second (min: 7.2% for train, max: 19% for skateboard), how often it enters and exists the field of view, how much area it covers and how frequently it is present.

### 4.3. Data set splits

The final annotations were split into training, validation, and testing subsets, as is standard for machine learning applications. The validation and testing subsets comprise 10% of the total, and this fraction is constant across classes. The splits were done such that no YouTube video can straddle two subsets. Part of the testing subset will be withheld in order to provide a quality measure for future public challenges based on YT-BB.

## 5. Baseline models

We measured the performance on YT-BB of image classification and object detection models trained on the COCO data set and vice-versa. This is possible because YT-BB’s labels are a subset of COCO’s, and both data sets classify and localize objects. The goal of this analysis two-fold: (1) to establish the relative difficulty of either task on the two datasets and (2) to provide a point of comparison for future network architectures.

### 5.1. Image classification

We started by comparing the relative difficulty of two instances of the same image classification model, one trained on YT-BB (“the YT-BB model”) and one trained on COCO (“the COCO model”). Our data set has explicit classification annotations. For COCO, we treated the presence or absence of any object localization of a class as either a positive or negative label, respectively. Both models employed

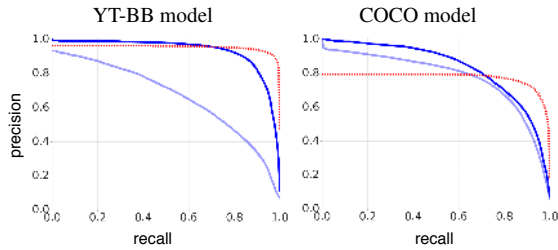


Figure 3. Baseline comparison on image classification. Mean precision-recall curve across 23 classes trained on the YT-BB (left) and the COCO [32] data sets (right). The dark line is the test set for the same data set; the light line is the test set for the opposite data set. The red dashed line is the evaluation where the classification scores are averaged across a video segment.

Train. data	Eval. data	smooth?	mAP	AUC
COCO	COCO	-	0.83	0.84
COCO	YT-BB	no	0.77	0.78
COCO	YT-BB	yes	0.77	0.78
YT-BB	YT-BB	no	0.93	0.94
YT-BB	YT-BB	yes	0.95	0.96
YT-BB	COCO	-	0.66	0.67

Table 1. Summary of image classification baselines. mAP and AUC are calculated across the 23 object classes (excluding “NONE”). The “smooth?” column indicates whether the predictions of YT-BB were averaged in time.

an Inception-v3<sup>2</sup> architecture [49] with logistic regression, implemented in TensorFlow[1]. The choice of logistic regression reflects the fact that multiple labels may be associated with a single image. Both models were initialized with the weights of an Inception-v3 image classification system pre-trained on the ImageNet 2012 Challenge data set [9] and subsequently fine-tuned on YT-BB/COCO individually.

We measured the mean precision-recall curve across all 23 classes (excluding the “NONE” class since it is not available in COCO). These results are shown in the dark curves in Figure 3. We find that training on YT-BB (mAP = 0.93) is easier than on COCO (mAP = 0.83), which could reflect the larger amount of training data per-class available in YT-BB.

One open question is the difficulty of domain transfer—i.e. training on one data set and evaluating on the other. We assessed this by measuring the mean precision-recall curve across 23 classes for the COCO model on YT-BB data (Figure 3, right panel, light curve) and vice-versa (Figure 3, left panel, light curve). We find that a COCO model evaluated on YT-BB (mAP = 0.77) was worse than one evaluated on COCO data (mAP = 0.83). The analogous claim is true for a model trained on YT-BB (Table 1). These results indicate that images in YT-BB are diverse and not just a subset of those in COCO.

<sup>2</sup> See Supplementary Section ?? for GitHub locations.

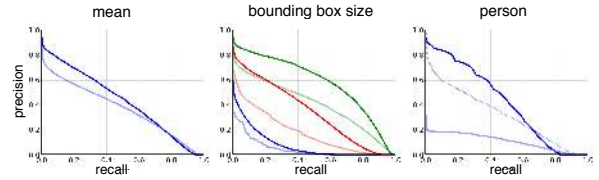


Figure 4. Baseline comparison of the COCO detection model. Precision-recall curve when evaluated on COCO [32] (dark line) and on YT-BB (light line). Left: Mean across the 23 classes. Center: Delineated by the bounding box size into large (green; top 2 curves), medium (red; middle 2 curves) and small (blue; bottom 2 curves). Right: Detection of the person class on COCO (dark top line), on YT-BB (light bottom line), and on person-specific subset of YT-BB (dashed middle line)

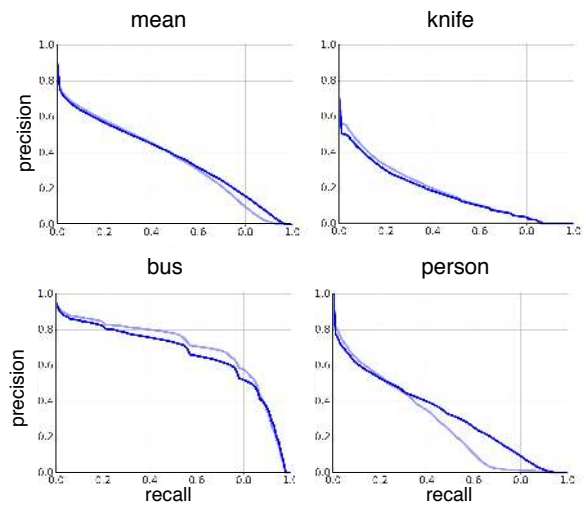


Figure 5. Effect of simple temporal smoothing on localization baselines. The dark curve shows precision-recall curve for single-frame object localization (mAP = 0.37, 0.18, 0.66, 0.33) and the light curve shows precision-recall employing simple temporal smoothing (mAP = 0.36, 0.19, 0.69, 0.28, respectively). See text for details.

## 5.2. Object detection

We then compared the relative difficulty of YT-BB and COCO for object detection. We used two instances of a Faster-RCNN<sup>2</sup> detection proposal architecture [22] paired with an Inception-ResNet-v2<sup>2</sup> feature network [41, 48]. Increasing the number of detection proposal results in improved object localizations at the expense of more computationally expensive inference and training. We selected a set of hyperparameters that resulted in 1400b FLOPs per frame for inference. Both instances were partially initialized with the weights of an Inception-ResNet-v2 image classification system trained on the ImageNet 2012 Challenge data set [9]. One instance was subsequently trained further on YT-BB (“the YT-BB model”) and another on COCO (“the COCO

Training data	Evaluation data	mAP	mAP @50%	mAP @75%	mAP small	mAP medium	mAP large
COCO*	COCO*	0.33	0.54	0.34	0.06	0.29	0.49
COCO	COCO	0.43	0.67	0.47	0.08	0.35	0.58
COCO	YT-BB	0.37	0.56	0.41	0.05	0.18	0.41
YT-BB	YT-BB	0.59	0.81	0.66	0.07	0.31	0.63
YT-BB	COCO	0.31	0.47	0.34	0.01	0.16	0.46

Table 2. Summary of image detection baselines between COCO and YT-BB across the 23 label classes. mAP is the mean average precision, averaged over multiple categories, multiple scales and multiple IOU fractions ranging between 0.5 and 0.95. mAP is the COCO competition metric. mAP @50 is the same but restricted to IOU  $\geq 50\%$ . mAP @75 is the same but restricted to IOU  $\geq 75\%$ . mAP (size) is the same but restricted to small, medium and large objects. For reference, the top row (COCO\*) highlights a single-crop, no-ensemble version of the winning entry to the COCO competition as measured across all 80 COCO classes.

model”).

We evaluated the performance of each model by measuring the mean precision-recall curve across all 23 classes. These results are summarized across several standard calculations of mAP for object detection in Table 2 and delineated by class in Supplementary Table 4. We find that training on YT-BB (mAP = 0.59) is easier than on COCO (mAP = 0.43). This result is consistent when measured across a range of detection box sizes.

In parallel with the image classification baseline, we also measured the relative difficulty of the two data sets by considering the problem of domain transfer. Again, we assessed this by examining the mean precision-recall curve across the 23 classes for the COCO model evaluated on YT-BB data. The degree to which the COCO model performance on YT-BB was worse than COCO reflects the relative difficulty and domain shift of the YT-BB data set. We found that a COCO model evaluated on YT-BB (mAP = 0.37) was indeed worse than when evaluated on COCO data (mAP = 0.43). This result was consistent across bounding box sizes and ranges of overlap assessment (Table 2, Figure 4). Notably, the COCO model was particularly poor at localizing medium and large YT-BB objects (Figure 4, middle panel). The analogous claim is true for a model trained on YT-BB (Table 2).

We next focused on the “person” class. The COCO model performed significantly worse in this case (mAP = 0.41 vs mAP = 0.12). At the lowest possible threshold, the COCO model fails to identify more than  $\sim 62\%$  of the “person” detections in YT-BB frames (Figure 4, right panel, light curve). At high thresholds, the COCO model exhibits low precision for “person”. This may be due to the fact that YT-BB is not exhaustively labeled. Unlabeled people may appear in videos which have been annotated for other classes. This may result in high false-positive scores and systematically lower precision. To mitigate this artifact we restrict the evaluation of the COCO model to a subset of YT-BB frames that have been labeled with a bounding box for “person” (Figure 4, right panel, dashed curve). The precision-recall curve was lifted as a result of the re-

moval of many unlabeled people, but remained below the precision-recall curve evaluated on COCO data. This analysis is however imperfect since images annotated with a “person” localization might contain additional people. Future work will be needed to determine how much of the additional difficulty ascribed to the YT-BB “person” label is due to the diversity of the “person” poses available in the YT-BB data set.

### 5.3. Exploiting temporal information in videos

All our baselines up to this point treated the frames as individual images. A salient aspect of the YT-BB data set is, however, that these frames exist within contiguous segments of video. Such video sequence can help regularize and improve video frame predictions. Devising better learning architectures for this purpose is an area of intense research interest [35, 40, 45]. As a demonstration of this data set’s potential, we performed several simple manipulations that indicate that temporal information exists and may be used by a learning system.

For the image classification task, we replaced the prediction for each label with the mean prediction for each label across each YT-BB video segment. The result of this temporal smoothing is shown in the dashed red line in Figure 3 and summarized in Table 1. Although the mAP and AUC do not change significantly (Table 1), the precision-recall curves do highlight that the temporally-averaged prediction systematically surpasses the single-frame prediction in the high recall regime (*e.g.* recall  $> \sim 0.7$ ). In principle, one could therefore build an improved system which achieves the *envelope* of the single-frame and temporally-averaged prediction scores.

For the object detection task, we down-weighted spurious weak object detections that appeared in single video frames but not in neighboring frames. Specifically, we artificially multiplied by  $\frac{1}{100}$  the confidence scores of detected objects that did not overlap significantly with previous and subsequent frame detections (IOU  $< 0.1$ ; confidence  $< 0.5$ ). Figure 5 shows the effects of this manipulation on the precision-recall curves. When aggregat-

ing across all classes, temporal smoothing slightly reduces model performance (mAP = 0.36 vs mAP = 0.37). We broke down this result to expose the diversity of behavior across labels. This revealed elevated precision-recall curves for some classes (e.g. “knife”, “bus”) but lowered curves for other classes (e.g. “person”). In principle, a unified model could at least learn which classes benefit and apply the correction only to those. Despite the mixed results, this analysis suggests that taking into account the temporal structure of the video could result in better detection models.

## 6. Discussion

In this paper, we introduced YT-BB, a new data set with 380,000 video segments, annotated with 5.6 million human-drawn bounding boxes tracking everyday objects in 23 categories. This represents an unprecedentedly large video detection data set (Section 2). First, we described the data mining process that led to minimally-edited videos and the human annotation stages that produced tight and precise bounding boxes, as well as precise tags indicating object absence (Sections 3.1 and 3.2). Then we presented relevant statistics and measures of annotation quality for each class. In particular, basic metrics of bounding box motion indicate that the objects or the camera exhibit significant changes throughout the segment (Section 4). Finally, we showed baselines for classification and detection trained and evaluated both on this data set and on COCO. These baselines demonstrate the potential for information in the video sequence to improve upon the basic inferences that can be done from single frames alone (Section 5).

Future work could refine YT-BB in various ways, most notably by adding more classes. With only 23 classes, we were able to pay special attention to each (Supplementary Section 2). Scaling up, classes would have to be treated more generically, as was done in [9]. This would magnify the challenges of crowd-sourcing schemes (“paradigm A”), in which low annotator accountability produces initial answers that often have poor quality [23], requiring significant additional effort to get to the final labels [37, 44, 55]. In this work, we observed such challenges in stages 1 and 2 (Section 3.2 and Supplementary Section 2). Alternatively, one could use a group of dedicated annotators who are committed to the project (“paradigm B”), as we did in stages 3 and 4 (Section 3.2). While we never carried out a proper A/B test, anecdotally we found paradigm B much more satisfying for a large-scale project. This can be traced back to the ability to train the annotators [11] and to provide them with feedback over time, resolving each problem encountered “once-and-for-all”.

Another direction for improvement could be to gather more bounding boxes. Increasing the sheer number does not seem critical as our baselines show no signs of overfitting. On the other hand, exhaustively labeling the exist-

ing videos may prove helpful, especially within the testing subset. While this would render the annotation task more complex, simplicity could be regained by introducing additional stages. Cascading stages have been found useful before [32]. In our case, it allowed the tuning of the annotation tool’s user interface to each task (Supplementary Section 1), rendering the first stage as much as 50 times faster than the last one. This, in turn, allowed for more negative examples to be present at the input since they could be easily discarded, and therefore the initial data mining stage could be more permissive. User interface optimization sometimes yielded unexpected results. For example, it turned out that providing default guesses for the bounding box locations was often *not* faster. Moreover, the annotators may find it easier to leave the default unchanged, which could bias the results toward such automatically generated defaults. Removing these defaults also made the tool simpler, which is generally known to be advantageous [16].

The baseline results suggest that there exists headroom for improving the quality of models on this data set. In particular, the data affords two distinct research directions. One is that the human annotation results identified individual video frames that are *hard negatives*, i.e. individual frames in the video that did not contain the object of interest even though surrounding frames did. These hard negatives might provide useful training and evaluation examples for future visual models.

The second research direction is to build models that harness the information in the temporal sequence of frames in a computationally efficient manner. Our baseline results indicate that even performing naive manipulations that incorporate such temporal aspects may contribute to better object classification and detection in video. The ability to build tractable, scalable models that exploit sequential information by keeping an internal memory state (e.g. [8, 21]) would likely lead toward better object detection and tracking (e.g. [31, 54]).<sup>3</sup>

## 7. Acknowledgements

We wish to thank Matthias Grundmann, John Gregg, Christian Falk and especially Thomas Silva for early efforts at bounding box annotations; Susanna Ricco, Sanketh Shetty for general advice about mining video data; George Toderici, Rahul Sukthankar for advice in many aspects of this work, Sami Abu-El-Haija, Manfred Georg for enormous efforts and generous advice about harvesting and annotating YouTube videos; Mir Shabber Ali Khan, Ashwin Kakarla and many others for the human annotations; and the larger Google Brain team for support with TensorFlow and training vision models.

<sup>3</sup>The data is available at <https://research.google.com/youtube-bb>



## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark, 2016.
- [3] G. Awad, J. Fiscus, M. Michel, D. Joy, W. Kraaij, A. F. Smeaton, G. Qunot, M. Eskevich, R. Aly, G. J. F. Jones, R. Ordeman, B. Huet, and M. Larson. Trecvid 2016: Evaluating video search, video event detection, localization, and hyperlinking. In *Proceedings of TRECVID 2016*. NIST, USA, 2016.
- [4] M. S. Bernstein, G. Little, R. C. Miller, B. Hartmann, M. S. Ackerman, D. R. Karger, D. Crowell, and K. Panovich. Soy-lent: a word processor with a crowd inside. *Communications of the ACM*, 58(8):85–94, 2015.
- [5] M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon’s mechanical turk a new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1):3–5, 2011.
- [6] Caffe Model Zoo. <http://github.com/BVLC/caffe/wiki/Model-Zoo>. [Accessed 19-Oct-2016].
- [7] J. J. Chen, N. J. Menezes, A. D. Bradley, and T. North. Opportunities for crowdsourcing research on amazon mechanical turk. *Interfaces*, 5(3), 2011.
- [8] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR*, abs/1406.1078, 2014.
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.
- [10] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: A benchmark. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 304–311. IEEE, 2009.
- [11] S. Dow, A. Kulkarni, B. Bunge, T. Nguyen, S. Klemmer, and B. Hartmann. Shepherding the crowd: managing and providing feedback to crowd workers. In *CHI’11 Extended Abstracts on Human Factors in Computing Systems*, pages 1669–1674. ACM, 2011.
- [12] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, 2015.
- [13] 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*, 88(2):303–338, 2010.
- [14] L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 106(1):59–70, 2007.
- [15] C. Fellbaum. *WordNet: An Electronic Lexical Database*. Bradford Books, 1998.
- [16] A. Finnerty, P. Kucherbaev, S. Tranquillini, and G. Convertino. Keep it simple: Reward and task design in crowdsourcing. In *Proceedings of the Biannual Conference of the Italian Chapter of SIGCHI*, page 14. ACM, 2013.
- [17] A. Gorban, H. Idrees, Y.-G. Jiang, A. Roshan Zamir, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://www.thumos.info>, 2015.
- [18] D. Grangier and S. Bengio. A discriminative kernel-based model to rank images from text queries. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2008.
- [19] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. 2007.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- [21] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, Nov. 1997.
- [22] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, et al. Speed/accuracy trade-offs for modern convolutional object detectors. *arXiv preprint arXiv:1611.10012*, 2016.
- [23] P. G. Ipeirotis, F. Provost, and J. Wang. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD workshop on human computation*, pages 64–67. ACM, 2010.
- [24] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22Nd ACM International Conference on Multimedia, MM’14*, pages 675–678, New York, NY, USA, 2014. ACM.
- [25] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [26] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. The YouTube Sports-1M Dataset. <http://github.com/gtoderici/sports-1m-dataset>, 2014. [Accessed 19-Oct-2016].
- [27] M. Kristan, J. Matas, A. Leonardis, M. Felsberg, L. Cehovin, G. Fernandez, T. Vojir, G. Hager, G. Nebehay, and R. Pflugfelder. The visual object tracking vot2015 challenge results. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1–23, 2015.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [29] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recog-

- dition. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2011.
- [30] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler. MOTChallenge 2015: Towards a benchmark for multi-target tracking. *arXiv:1504.01942 [cs]*, Apr. 2015. arXiv: 1504.01942.
- [31] H. Li, Y. Li, and F. Porikli. Deeptack: Learning discriminative feature representations online for robust visual tracking. *CoRR*, abs/1503.00072, 2015.
- [32] 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*, pages 740–755. Springer, 2014.
- [33] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 416–423. IEEE, 2001.
- [34] S. A. Nene, S. K. Nayar, H. Murase, et al. Columbia object image library (coil-20). Technical report, Technical report CUCS-005-96, 1996.
- [35] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Computer Vision and Pattern Recognition*, 2015.
- [36] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. Corrado, and J. Dean. Zero-shot learning by convex combination of semantic embeddings. *CoRR*, abs/1312.5650, 2013.
- [37] P. Paritosh. Human computation must be reproducible. 2012.
- [38] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3282–3289. IEEE, 2012.
- [39] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Youtube-Objects dataset. <https://data.vision.ee.ethz.ch/cvl/youtube-objects/>, 2012. [Accessed 19-Oct-2016].
- [40] M. Ranzato, A. Szlam, J. Bruna, M. Mathieu, R. Collobert, and S. Chopra. Video (language) modeling: a baseline for generative models of natural videos. *CoRR*, abs/1412.6604, 2014.
- [41] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [42] O. Russakovsky, J. Deng, Z. Huang, A. C. Berg, and L. Fei-Fei. Detecting avocados to zucchinis: what have we done, and where are we going? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2064–2071, 2013.
- [43] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [44] V. Sheng, F. Provost, and P. Ipeirotis. Get another label? improving data quality and data mining. 2008.
- [45] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. *CoRR*, abs/1406.2199, 2014.
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [47] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [48] C. Szegedy, S. Ioffe, and V. Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016.
- [49] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. *arXiv preprint arXiv:1512.00567*, 2015.
- [50] K. Tang, R. Sukthankar, J. Yagnik, and L. Fei-Fei. Discriminative segment annotation in weakly labeled video. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR 2013)*, 2013.
- [51] TensorFlow-Slim image classification library. <http://github.com/tensorflow/models/tree/master/slim>. [Accessed 19-Oct-2016].
- [52] A. Torralba, K. P. Murphy, and W. T. Freeman. Sharing features: efficient boosting procedures for multiclass object detection. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–762. IEEE, 2004.
- [53] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. *CoRR*, abs/1411.4555, 2014.
- [54] L. Wang, W. Ouyang, X. Wang, and H. Lu. Visual tracking with fully convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [55] T. P. Waterhouse. Pay by the bit: an information-theoretic metric for collective human judgment. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 623–638. ACM, 2013.
- [56] J. Xiao, K. A. Ehinger, J. Hays, A. Torralba, and A. Oliva. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, pages 1–20, 2014.
- [57] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, pages 3485–3492. IEEE, 2010.