

The Unreasonable Effectiveness of Noisy Data for Fine-Grained Recognition

Jonathan Krause¹✉, Benjamin Sapp², Andrew Howard², Howard Zhou², Alexander Toshev², Tom Duerig², James Philbin², and Li Fei-Fei¹

¹ Stanford University, Stanford, USA
{[jkrause](mailto:jkrause@cs.stanford.edu), [feifeili](mailto:feifeili@cs.stanford.edu)}@cs.stanford.edu

² Google, Mountain View, USA

benjamin.sapp@gmail.com, howarda@google.com, howardzhou@google.com,
toshev@google.com, tduerig@google.com, philbinj@gmail.com

Abstract. Current approaches for fine-grained recognition do the following: First, recruit experts to annotate a dataset of images, optionally also collecting more structured data in the form of part annotations and bounding boxes. Second, train a model utilizing this data. Toward the goal of solving fine-grained recognition, we introduce an alternative approach, leveraging free, noisy data from the web and simple, generic methods of recognition. This approach has benefits in both performance and scalability. We demonstrate its efficacy on four fine-grained datasets, greatly exceeding existing state of the art without the manual collection of even a single label, and furthermore show first results at scaling to more than 10,000 fine-grained categories. Quantitatively, we achieve top-1 accuracies of 92.3 % on CUB-200-2011, 85.4 % on Birdsnap, 93.4 % on FGVC-Aircraft, and 80.8 % on Stanford Dogs without using their annotated training sets. We compare our approach to an active learning approach for expanding fine-grained datasets.

1 Introduction

Fine-grained recognition refers to the task of distinguishing very similar categories, such as breeds of dogs [27, 36], species of birds [4, 5, 57, 59], or models of cars [30, 69]. Since its inception, great progress has been made, with accuracies on the popular CUB-200-2011 bird dataset [59] steadily increasing from 10.3 % [59] to 84.6 % [68].

The predominant approach in fine-grained recognition today consists of two steps. First, a dataset is collected. Since fine-grained recognition is a task inherently difficult for humans, this typically requires either recruiting a team of experts [37, 57] or extensive crowd-sourcing pipelines [4, 30]. Second, a method

Work done while J. Krause was interning at Google.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46487-9_19](https://doi.org/10.1007/978-3-319-46487-9_19)) contains supplementary material, which is available to authorized users.




Fig. 1. There are more than 14,000 species of birds in the world. In this work we show that using noisy data from publicly-available online sources can not only improve recognition of categories in today’s datasets, but also scale to very large numbers of fine-grained categories, which is extremely expensive with the traditional approach of manually collecting labels for fine-grained datasets. Here we show 4,225 of the 10,982 categories recognized in this work.

for recognition is trained using these expert-annotated labels, possibly also requiring additional annotations in the form of parts, attributes, or relationships [5, 26, 35, 74]. While methods following this approach have shown some success [5, 28, 35, 74], their performance and scalability is constrained by the paucity of data available due to these limitations. With this traditional approach it is prohibitive to scale up to all 14,000 species of birds in the world (Fig. 1), 278,000 species of butterflies and moths, or 941,000 species of insects [24].

In this paper, we show that it is possible to train effective models of fine-grained recognition using noisy data from the web and simple, generic methods of recognition [53, 54]. We demonstrate recognition abilities greatly exceeding current state of the art methods, achieving top-1 accuracies of 92.3% on CUB-200-2011 [59], 85.4% on Birdsnap [4], 93.4% on FGVC-Aircraft [37], and 80.8% on Stanford Dogs [27] *without using a single manually-annotated training label from the respective datasets*. On CUB, this is nearly at the level of human experts [6, 57]. Building upon this, we scale up the number of fine-grained classes recognized, reporting first results on over 10,000 species of birds and 14,000 species of butterflies and moths.

The rest of this paper proceeds as follows: After an overview of related work in Sect. 2, we provide an analysis of publicly-available noisy data for fine-grained recognition in Sect. 3, analyzing its quantity and quality. We describe a more traditional active learning approach for obtaining larger quantities of fine-grained data in Sect. 4, which serves as a comparison to purely using noisy data. We present extensive experiments in Sect. 5, and conclude with discussion in Sect. 6.

2 Related Work

Fine-Grained Recognition. The majority of research in fine-grained recognition has focused on developing improved models for classification [1, 3, 5, 7–9, 14, 16, 18, 20–22, 28, 29, 35, 36, 40, 41, 48–50, 65, 67, 68, 70–72, 74–77]. While these works have made great progress in modeling fine-grained categories given the limited data available, very few works have considered the impact of that data [57, 67, 68]. Xu *et al.* [68] augment datasets annotated with category labels and parts with web images in a multiple instance learning framework, and Xie *et al.* [67] do multitask training, where one task uses a ground truth fine-grained dataset and the other does not require fine-grained labels. While both of these methods have shown that augmenting fine-grained datasets with additional data can help, in our work we present results which completely forgo the use of any curated ground truth dataset. In one experiment hinting at the use of noisy data, Van Horn *et al.* [57] show the possibility of learning 40 bird classes from Flickr images. Our work validates and extends this idea, using similar intuition to significantly improve performance on existing fine-grained datasets and scale fine-grained recognition to over ten thousand categories, which we believe is necessary in order to fully explore the research direction.

Considerable work has also gone into the challenging task of curating fine-grained datasets [4, 27, 30, 31, 57–59, 64, 69] and developing interactive methods for recognition with a human in the loop [6, 60–62]. While these works have demonstrated effective strategies for collecting images of fine-grained categories, their scalability is ultimately limited by the requirement of manual annotation. Our work provides an alternative to these approaches.

Learning from Noisy Data. Our work is also inspired by methods that propose to learn from web data [10, 11, 15, 19, 34, 44] or reason about label noise [38, 42, 51, 57, 66]. Works that use web data typically focus on detection and classification of a set of coarse-grained categories, but have not yet examined the fine-grained setting. Methods that reason about label noise have been divided in their results: some have shown that reasoning about label noise can have a substantial effect on recognition performance [65], while others demonstrate little change from reducing the noise level or having a noise-aware model [42, 51, 57]. In our work, we demonstrate that noisy data can be surprisingly effective for fine-grained recognition, providing evidence in support of the latter hypothesis.

3 Noisy Fine-Grained Data

In this section we provide an analysis of the imagery publicly available for fine-grained recognition, which we collect via web search.¹ We describe its quantity, distribution, and levels of noise, reporting each on multiple fine-grained domains.

¹ Google image search: <http://images.google.com>.




Fig. 2. Distributions of the number of images per category available via image search for the categories in CUB, Birdsnap, and L-Bird (far left), FGVC and L-Aircraft (middle left), and L-Butterfly (middle right). At far right we aggregate and plot the average number of images per category in each dataset in addition to the training sets of each curated dataset we consider, denoted CUB-GT, Birdsnap-GT, and FGVC-GT.

3.1 Categories

We consider four domains of fine-grained categories: birds, aircraft, Lepidoptera (a taxonomic order including butterflies and moths), and dogs. For birds and Lepidoptera, we obtained lists of fine-grained categories from Wikipedia, resulting in 10,982 species of birds and 14,553 species of Lepidoptera, denoted L-Bird (“Large Bird”) and L-Butterfly. For aircraft, we assembled a list of 409 types of aircraft by hand (including aircraft in the FGVC-Aircraft [37] dataset, abbreviated FGVC). For dogs, we combine the 120 dog breeds in Stanford Dogs [27] with 395 other categories to obtain the 515-category L-Dog. We evaluate on two other fine-grained datasets in addition to FGVC and Stanford Dogs: CUB-200-2011 [59] and Birdsnap [4], for a total of four evaluation datasets. CUB and Birdsnap include 200 and 500 species of common birds, respectively, FGVC has 100 aircraft variants, and Stanford Dogs contains 120 breeds of dogs. In this section we focus our analysis on the categories in L-Bird, L-Butterfly, and L-Aircraft in addition to the categories in their evaluation datasets.

3.2 Images from the Web

We obtain imagery via Google image search results, using all returned images as images for a given category. For L-Bird and L-Butterfly, queries are for the scientific name of the category, and for L-Aircraft and L-Dog queries are simply for the category name (*e.g.* “Boeing 737-200” or “Pembroke Welsh Corgi”).

Quantifying the Data. How much fine-grained data is available? In Fig. 2 we plot distributions of the number of images retrieved for each category and report aggregates across each set of categories. We note several trends: Categories in existing datasets, which are typically common within their fine-grained domain, have more images per category than the long-tail of categories present in the larger L-Bird, L-Aircraft, or L-Butterfly, with the effect most pronounced in L-Bird and L-Butterfly. Further, domains of fine-grained categories have substantially different distributions, *i.e.* L-Bird and L-Aircraft have more images



Fig. 3. Examples of cross-domain noise for birds, butterflies, airplanes, and dogs. Images are generally of related categories that are outside the domain of interest, *e.g.* a map of a bird’s typical habitat or a t-shirt containing the silhouette of a dog.

per category than L-Butterfly. This makes sense – fine-grained categories and domains of categories that are more common and have a larger enthusiast base will have more imagery since more photos are taken of them. We also note that results tend to be limited to roughly 800 images per category, even for the most common categories, which is likely a restriction placed on public search results.

Most striking is the large difference between the number of images available via web search and in existing fine-grained datasets: even Birdsnap, which has an average of 94.8 images per category, contains only 13% as many images as can be obtained with a simple image search. Though their labels are noisy, web searches unveil an order of magnitude more data which can be used to learn fine-grained categories.

In total, for all four datasets, we obtained 9.8 million images for 26,458 categories, requiring 151.8 GB of disk space. All urls will be released.

Noise. Though large amounts of imagery are freely available for fine-grained categories, focusing only on scale ignores a key issue: *noise*. We consider two types of label noise, which we call *cross-domain* noise and *cross-category* noise. We define cross-domain noise to be the portion of images that are not of any category in the same fine-grained domain, *i.e.* for birds, it is the fraction of images that do not contain a bird (examples in Fig. 3). In contrast, *cross-category* noise is the portion of images that have the wrong label within a fine-grained domain, *i.e.* an image of a bird with the wrong species label.

To quantify levels of cross-domain noise, we manually label a 1,000 image sample from each set of search results, with results in Fig. 4. Although levels of noise are not too high for any set of categories (max. 34.2% for L-Butterfly), we notice an interesting correlation: cross-domain noise decreases moderately as the number of images per category (Fig. 2) increases. We hypothesize that categories with many search results have a corresponding large pool of images to draw results from, and thus actual search results will tend to be higher-precision.

In contrast to cross-domain noise, cross-category noise is much harder to quantify, since doing so effectively requires ground truth fine-grained labels of query results. To examine cross-category noise from at least one vantage point,




Fig. 4. The cross-domain noise in search results for each domain.




Fig. 5. The percentage of images retained after filtering.

we show the confusion matrix of given versus predicted labels on 30 categories in the CUB [59] test set and their web images in Fig. 6, left and right, which we generate via a classifier trained on the CUB training set, acting as a noisy proxy for ground truth labels. In these confusion matrices, cross-category noise is reflected as a strong off-diagonal pattern, while cross-domain noise would manifest as a diffuse pattern of noise, since images not of the same domain are an equally bad fit to all categories. Based on this interpretation, the web images show a moderate amount more cross-category noise than the clean CUB test set, though the general confusion pattern is similar.

We propose a simple, yet effective strategy to reduce the effects of cross-category noise: exclude images that appear in search results for more than one category. This approach, which we refer to as *filtering*, specifically targets images for which there is explicit ambiguity in the category label (examples in Fig. 7). As we demonstrate experimentally, filtering can improve results while reducing training time via the use of a more compact training set – we show the portion of images kept after filtering in Fig. 5. Agreeing with intuition, filtering removes more images when there are more categories. Anecdotally, we have also tried a few techniques to combat cross-domain noise, but initial experiments did not see any improvement in recognition so we do not expand upon them here. While reducing cross-domain noise should be beneficial, we believe that it is not as important as cross-category noise in fine-grained recognition due to the absence of out-of-domain classes during testing.

4 Data via Active Learning

In this section we briefly describe an active learning-based approach for collecting large quantities of fine-grained data. Active learning and other human-in-the-loop systems have previously been used to create datasets in a more cost-efficient way than manual annotation [12, 46, 73], and our goal is to compare this more traditional approach with simply using noisy data, particularly when considering the application of fine-grained recognition. In this paper, we apply active learning to the 120 dog breeds in the Stanford Dogs [27] dataset.




Fig. 6. Confusion matrices of the predicted label (column) given the provided label (row) for 30 CUB categories on the CUB test set (left) and search results for CUB categories (right). For visualization purposes we remove the diagonal.



Fig. 7. Examples of images removed via ltering and the categories whose results they appeared in. Some share similar names (left examples), while others share similar locations (right examples).

Our system for active learning begins by training a classifier on a seed set of input images and labels (*i.e.* the Stanford Dogs training set), then proceeds by iteratively picking a set of images to annotate, obtaining labels with human annotators, and re-training the classifier. We use a convolutional neural network [25, 32, 53] for the classifier, and now describe the key steps of sample selection and human annotation in more detail.

Sample Selection. There are many possible criterion for sample selection [46]. We employ confidence-based sampling: For each category c , we select the $b\hat{P}(c)$ images with the top class scores $f_c(x)$ as determined by our current model, where $\hat{P}(c)$ is a desired prior distribution over classes, b is a budget on the number of images to annotate, and $f_c(x)$ is the output of the classifier. The intuition is as follows: even when $f_c(x)$ is large, false positives still occur quite frequently – in Fig. 8 left, observe that the false positive rate is about 20% at the highest confidence range, which might have a large impact on the model. This contrasts with approaches that focus sampling in uncertain regions [2, 17, 33, 39]. We find that images sampled with uncertainty criteria are typically ambiguous and difficult or even impossible for both models *and* humans to annotate correctly, as demonstrated in Fig. 8 bottom row: unconfident samples are often heavily occluded, at unusual viewpoints, or of mixed, ambiguous breeds, making it unlikely that they can be annotated effectively. This strategy is similar to the “expected model change” sampling criteria [47], but done for each class independently.

Human Annotation. Our interface for human annotation of the selected images is shown in Fig. 9. Careful construction of the interface, including the addition of both positive and negative examples, as well as hidden “gold standard” images for immediate feedback, improves annotation accuracy considerably (see Supplementary Material for quantitative results). Final category decisions are made via majority vote of three annotators.

5 Experiments

5.1 Implementation Details

The base classifier we use in all noisy data experiments is the Inception-v3 convolutional neural network architecture [54], which is among the state of the art methods for generic object recognition [23, 43, 52]. Learning rate schedules are determined by performance on a holdout subset of the training data, which is 10% of the training data for control experiments training on ground truth datasets, or 1% when training on the larger noisy web data. Unless otherwise noted, all recognition results use as input a single crop in the center of the image.

Our active learning comparison uses the Yahoo Flickr Creative Commons 100M dataset [55] as its pool of unlabeled images, which we first pre-filter with a binary dog classifier and localizer [53], resulting in 1.71 million candidate dogs. We perform up to two rounds of active learning, with a sampling budget B of $10\times$ the original dataset size per round². For experiments on Stanford Dogs, we use the CNN of [25], which is pre-trained on a version of ILSVRC [13, 43] with dog data removed, since Stanford Dogs is a subset of ILSVRC training data.




Fig. 8. **Left:** Classifier confidence versus false positive rate on 100,000 randomly sampled from Flickr images (YFCC100M [55]) with dog detections. Even the most confident images have a 20% false positive rate. **Right:** Samples from Flickr. Rectangles below images denote correct (green), incorrect (red), or ambiguous (yellow). **Top row:** Samples with high confidence for class “Pug” from YFCC100M. **Bottom row:** Samples with low confidence score for class “Pug”. (Color figure online)



Fig. 9. Our tool for binary annotation of fine-grained categories. Instructional positive images are provided in the upper left and negatives are provided in the lower left.

² To be released.

Table 1. Comparison of data source used during training with recognition performance, given in terms of Top-1 accuracy. “CUB-GT” indicates training only on the ground truth CUB training set, “Web (raw)” trains on all search results for CUB categories, and “Web (filtered)” applies filtering between categories within a domain (birds). L-Bird denotes training first on L-Bird, then fine-tuning on the subset of categories under evaluation (*i.e.* the filtered web images), and L-Bird + CUB-GT indicates training on L-Bird, then fine-tuning on Web (filtered), and finally fine-tuning again on CUB-GT. Similar notation is used for the other datasets. “(MC)” indicates using multiple crops at test time (see text for details). We note that only the rows with “-GT” make use of the ground truth training set; all other rows rely solely on noisy web imagery.

Training data	Acc.	Dataset	Training data	Acc.	Dataset
CUB-GT	84.4	CUB [59]	FGVC-GT	88.1	FGVC [37]
Web (raw)	87.7		Web (raw)	90.7	
Web (filtered)	89.0		Web (filtered)	91.1	
L-Bird	91.9		L-Aircraft	90.9	
L-Bird(MC)	92.3		L-Aircraft(MC)	93.4	
L-Bird + CUB-GT	92.2		L-Aircraft + FGVC-GT	94.5	
L-Bird + CUB-GT(MC)	92.8		L-Aircraft + FGVC-GT(MC)	95.9	
Birdsnap-GT	78.2		Birdsnap [4]	Stanford-GT	
Web (raw)	76.1	Web (raw)		78.5	
Web (filtered)	78.2	Web (filtered)		78.4	
L-Bird	82.8	L-Dog		78.4	
L-Bird(MC)	85.4	L-Dog(MC)		80.8	
L-Bird + Birdsnap-GT	83.9	L-Dog + Stanford-GT		84.0	
L-Bird + Birdsnap-GT(MC)	85.4	L-Dog + Stanford-GT(MC)		85.9	

5.2 Removing Ground Truth from Web Images

One subtle point to be cautious about when using web images is the risk of inadvertently including images from ground truth test sets in the web training data. To deal with this concern, we performed an aggressive deduplication procedure with all ground truth test sets and their corresponding web images. This process follows Wang *et al.* [63], which is a state of the art method for learning a similarity metric between images. We tuned this procedure for high near-duplicate recall, manually verifying its quality. More details are included in the Supplementary Material.

5.3 Main Results

We present our main recognition results in Table 1, where we compare performance when the training set consists of either the ground truth training set, raw web images of the categories in the corresponding evaluation dataset, web images after applying our filtering strategy, all web images of a particular domain, or all images including even the ground truth training set.

On CUB-200-2011 [59], the smallest dataset we consider, even using raw search results as training data results in a better model than the annotated training set, with filtering further improving results by 1.3%. For Birdsnap [4],

the largest of the ground truth datasets we evaluate on, raw data mildly underperforms using the ground truth training set, though filtering improves results to be on par. On both CUB and Birdsnap, training first on the very large set of categories in L-Bird results in dramatic improvements, improving performance on CUB further by 2.9% and on Birdsnap by 4.6%. This is an important point: even if the end task consists of classifying only a small number of categories, training with more fine-grained categories yields significantly more effective networks. This can also be thought of as a form of transfer learning within the same fine-grained domain, allowing features learned on a related task to be useful for the final classification problem. When permitted access to the annotated ground truth training sets for additional fine-tuning and domain transfer, results increase by another 0.3% on CUB and 1.1% on Birdsnap.

For the aircraft categories in FGVC, results are largely similar but weaker in magnitude. Training on raw web data results in a significant gain of 2.6% compared to using the curated training set, and filtering, which did not affect the size of the training set much (Fig. 5), changes results only slightly in a positive direction. Counterintuitively, pre-training on a larger set of aircraft does not improve results on FGVC. Our hypothesis for the difference between birds and aircraft in this regard is this: since there are many more species of birds in L-Bird than there are aircraft in L-Aircraft (10,982 vs. 409), not only is the training size of L-Bird larger, but each training example provides stronger information because it distinguishes between a larger set of mutually-exclusive categories. Nonetheless, when access to the curated training set is available for fine-tuning, performance dramatically increases to 94.5%. On Stanford Dogs we see results similar to FGVC, though for dogs we happen to see a mild loss when comparing to the ground truth training set, not much difference with filtering or using L-Dog, and a large boost from adding in the ground truth training set.

An additional factor that can influence performance of web models is domain shift – if images in the ground truth test set have very different visual properties compared to web images, performance will naturally differ. Similarly, if category names or definitions within a dataset are even mildly off, web-based methods will be at a disadvantage without access to the ground truth training set. Adding the ground truth training data fixes this domain shift, making web-trained models quickly recover, with a particularly large gain if the network has already learned a good representation, matching the pattern of results for Stanford Dogs.

Limits of Web-Trained Models. To push our models to their limits, we additionally evaluate using 144 image crops at test time, averaging predictions across each crop, denoted “(MC)” in Table 1. This brings results up to 92.3%/92.8% on CUB (without/with CUB training data), 85.4%/85.4% on Birdsnap, 93.4%/95.9% on FGVC, and 80.8%/85.9% on Stanford Dogs. We note that this is close to human expert performance on CUB, which is estimated to be between 93% [6] and 95.6% [57].

Table 2. Comparison with prior work on CUB-200-2011 [59]. We only include methods which use no annotations at test time. Here “GT” refers to using Ground Truth category labels in the training set of CUB, “BBox” indicates using bounding boxes, and “Parts” additionally uses part annotations.

Method	Training annotations	Acc.
Alignments [21]	GT	53.6
PDD [50]	GT + BB + Parts	60.6
PB R-CNN [74]	GT + BB + Parts	73.9
Weak Sup. [77]	GT	75.0
PN-DCN [5]	GT + BB + Parts	75.7
Two-Level [65]	GT	77.9
Consensus [48]	GT + BB + Parts	78.3
NAC [49]	GT	81.0
FG-Without [29]	GT + BB	82.0
STN [26]	GT	84.1
Bilinear [35]	GT	84.1
Augmenting [68]	GT + BB + Parts + Web	84.6
Noisy Data + CNN [54]	Web	92.3

Comparison with Prior Work. We compare our results to prior work on CUB, the most competitive fine-grained dataset, in Table 2. While even our baseline model using only ground truth data from Table 1 was at state of the art levels, by forgoing the CUB training set and only training using noisy data from the web, our models greatly outperform all prior work. On FGVC, which is more recent and fewer works have evaluated on, the best prior performing method we are aware of is the Bilinear CNN model of Lin *et al.* [35], which has accuracy 84.1% (ours is 93.4% without FGVC training data, 95.9% with), and on Birdsnap, which is even more recent, the best performing method we are aware of that uses no extra annotations during test time is the original 66.6% by Berg *et al.* [4] (ours is 85.4%). On Stanford Dogs, the most competitive related work is [45], which uses an attention-based recurrent neural network to achieve 76.8% (ours is 80.8% without ground truth training data, 85.9% with).

We identify two key reasons for these large improvements: The first is the use of a strong generic classifier [54]. A number of prior works have identified the importance of having well-trained CNNs as components in their systems for fine-grained recognition [5, 26, 29, 35, 74], which our work provides strong evidence for. On all four evaluation datasets, our CNN of choice [54], trained on the ground truth training set alone and without any architectural modifications, performs at levels at or above the previous state-of-the-art. The second reason for improvement is the large utility of noisy web data for fine-grained recognition, which is the focus of this work.

We finally remind the reader that our work focuses on the application-level problem of recognizing a given set of fine-grained categories, which might not come with their own expert-annotated training images. The use of existing test sets serves to provide an accurate measure of performance and put our work in a larger context, but results may not be strictly comparable with prior work that operates within a single given dataset.

Comparison with Active Learning. We compare using noisy web data with a more traditional active learning-based approach (Sect. 4) under several different settings in Table 3. We first verify the efficacy of active learning itself: when training the network from scratch (*i.e.* no fine-tuning), active learning improves performance by up to 15.6%, and when fine-tuning, results still improve by 1.5%.

How does active learning compare to using web data? Purely using filtered web data compares favorably to non-fine-tuned active learning methods (4.4% better), though lags behind the fine-tuned models somewhat. To better compare the active learning and noisy web data, we factor out the difference in scale by performing an experiment with subsampled active learning data, setting it to be the same size as the filtered web data. Surprisingly, performance is very similar, with only a 0.4% advantage for the cleaner, annotated active learning data, highlighting the effectiveness of noisy web data despite the lack of manual annotation. If we furthermore augment the filtered web images with the Stanford Dogs training set, which the active learning method notably used both as training data and its seed set of images, performance improves to even be slightly better than the manually-annotated active learning data (0.5% improvement).

Table 3. Active learning-based results on Stanford Dogs [27], presented in terms of top-1 accuracy. Methods with “(scratch)” indicate training from scratch and “(ft)” indicates fine-tuning from a network pre-trained on ILSVRC, with web models also fine-tuned. “subsample” refers to downsampling the active learning data to be the same size as the filtered web images. Note that Stanford-GT is a subset of active learning data, which is denoted “A.L.”.

Training Procedure	Acc.
Stanford-GT (scratch)	58.4
A.L., one round (scratch)	65.8
A.L., two rounds (scratch)	74.0
Stanford-GT (ft)	80.6
A.L., one round (ft)	81.6
A.L., one round (ft, subsample)	78.8
A.L., two rounds (ft)	82.1
Web (filtered)	78.4
Web (filtered) + Stanford-GT	82.6

These experiments indicate that, while more traditional active learning-based approaches towards expanding datasets are effective ways to improve recognition performance given a suitable budget, simply using noisy images retrieved from the web can be nearly as good, if not better. As web images require no manual annotation and are openly available, we believe this is strong evidence for their use in solving fine-grained recognition.

Very Large-Scale Fine-Grained Recognition. A key advantage of using noisy data is the ability to scale to large numbers of fine-grained classes. However, this poses a challenge for evaluation – it is infeasible to manually annotate images with one of the 10,982 categories in L-Bird, 14,553 categories in L-Butterfly, and would even be very time-consuming to annotate images with the 409 categories in L-Aircraft. Therefore, we turn to an approximate evaluation, establishing a rough estimate on true performance. Specifically, we query Flickr for up to 25 images of each category, keeping only those images whose title strictly contains the name of each category, and aggressively deduplicate these images with our training set in order to ensure a fair evaluation. Although this is not a perfect evaluation set, and is thus an area where annotation of fine-grained datasets is particularly valuable [57], we find that it is remarkably clean on the surface: based on a 1,000-image estimate, we measure the cross-domain noise of L-Bird at only 1%, L-Butterfly at 2.3%, and L-Aircraft at 4.5%. An independent evaluation [57] further measures all sources of noise combined to be only 16% when searching for bird species. In total, this yields 42,115 testing images for L-Bird, 42,046 for L-Butterfly, and 3,131 for L-Aircraft.




Fig. 10. Classification results on very large-scale fine-grained recognition. From top to bottom, depicted are examples of categories in L-Bird, L-Butterfly, and L-Aircraft, along with their category name. The first examples in each row are correctly predicted by our models, while the last two examples in each row are errors, with our prediction in grey and correct category (according to Flickr metadata) printed below.

Given the difficulty and noise, performance is surprisingly high: On L-Bird top-1 accuracy is 73.1%/75.8% (1/144 crops), for L-Butterfly it is 65.9%/68.1%, and for L-Aircraft it is 72.7%/77.5%. Corresponding mAP numbers, which are better suited for handling class imbalance, are 61.9, 54.8, and 70.5, reported for the single crop setting. We show qualitative results in Fig. 10. These categories span multiple continents in space (birds, butterflies) and decades in time (aircraft), demonstrating the breadth of categories in the world that can be recognized using only public sources of noisy fine-grained data. To the best of our knowledge, these results represent the largest number of fine-grained categories distinguished by any single system to date.

How Much Data is Really Necessary? In order to better understand the utility of noisy web data for fine-grained recognition, we perform a control experiment on the web data for CUB. Using the filtered web images as a base, we train models using progressively larger subsets of the results as training data, taking the top ranked images across categories for each experiment. Performance versus the amount of training data is shown in Fig. 11. Surprisingly, relatively few web images are required to do as well as training on the CUB training set, and adding more noisy web images always helps, even when at the limit of search results. Based on this analysis, we estimate that one noisy web image for CUB categories is “worth” 0.507 ground truth training images [56].

Error Analysis. Given the high performance of these models, what room is left for improvement? In Fig. 12 we show the taxonomic distribution of the remaining errors on L-Bird. The vast majority of errors (74.3%) are made between very similar classes at the genus level, indicating that most of the remaining errors are indeed between extremely similar categories, and only very few errors (7.4%)




Fig. 11. Number of web images used for training vs. performance on CUB-200-2011 [59]. We vary the amount of web training data in multiples of the CUB training set size (5,994 images). Also shown is performance when training on the ground truth CUB training set (CUB-GT).




Fig. 12. The errors on L-Bird that fall in each taxonomic rank, represented as a portion of all errors made. For each error made, we calculate the taxonomic rank of the least common ancestor of the predicted and test category.

are made between dissimilar classes, whose least common ancestor is the “Aves” (*i.e.* Bird) taxonomic class. This suggests that most errors still made by the models are fairly reasonable, corroborating the qualitative results of Fig. 10.

6 Discussion

In this work we have demonstrated the utility of noisy data toward solving the problem of fine-grained recognition. We found that the combination of a generic classification model and web data, filtered with a simple strategy, was surprisingly effective at discriminating fine-grained categories. This approach performs favorably when compared to a more traditional active learning method for expanding datasets, but is even more scalable, which we demonstrated experimentally on up to 14,553 fine-grained categories. One potential limitation of the approach is the availability of imagery for categories either not found or not described in the public domain, for which an alternative method such as active learning may be better suited. Another limitation is the current focus on classification, which may be problematic if applications arise where multiple objects are present or localization is otherwise required. Nonetheless, with these insights on the unreasonable effectiveness of noisy data, we are optimistic for applications of fine-grained recognition in the near future.

Acknowledgments. We thank Gal Chechik, Chuck Rosenberg, Zhen Li, Timnit Gebru, Vignesh Ramanathan, Oliver Groth, and the anonymous reviewers for valuable feedback.

References

1. Angelova, A., Zhu, S., Lin, Y.: Image segmentation for large-scale subcategory flower recognition. In: Workshop on Applications of Computer Vision (WACV), pp. 39–45. IEEE (2013)

2. Balcan, M.-F., Broder, A., Zhang, T.: Margin based active learning. In: Bshouty, N.H., Gentile, C. (eds.) COLT 2007. LNCS (LNAI), vol. 4539, pp. 35–50. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-72927-3_5](https://doi.org/10.1007/978-3-540-72927-3_5)
3. Berg, T., Belhumeur, P.N.: Poof: Part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 955–962. IEEE (2013)
4. Berg, T., Liu, J., Lee, S.W., Alexander, M.L., Jacobs, D.W., Belhumeur, P.N.: Birdsnap: large-scale fine-grained visual categorization of birds. In: *Computer Vision and Pattern Recognition (CVPR)*, June 2014
5. Branson, S., Van Horn, G., Perona, P., Belongie, S.: Improved bird species recognition using pose normalized deep convolutional nets. In: *British Machine Vision Conference (BMVC)* (2014)
6. Branson, S., Van Horn, G., Wah, C., Perona, P., Belongie, S.: The ignorant led by the blind: a hybrid human-machine vision system for fine-grained categorization. *Int. J. Comput. Vision (IJCV)*, 1–27 (2014)
7. Chai, Y., Lempitsky, V., Zisserman, A.: Bicos: A bi-level co-segmentation method for image classification. In: *International Conference on Computer Vision (ICCV)*. IEEE (2011)
8. Chai, Y., Lempitsky, V., Zisserman, A.: Symbiotic segmentation and part localization for fine-grained categorization. In: *International Conference on Computer Vision (ICCV)*, pp. 321–328. IEEE (2013)
9. Chai, Y., Rahtu, E., Lempitsky, V., Gool, L., Zisserman, A.: TriCoS: a tri-level class-discriminative co-segmentation method for image classification. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012*. LNCS, vol. 7572, pp. 794–807. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33718-5_57](https://doi.org/10.1007/978-3-642-33718-5_57)
10. Chen, X., Gupta, A.: Webly supervised learning of convolutional networks. In: *International Conference on Computer Vision (ICCV)*. IEEE (2015)
11. Chen, X., Shrivastava, A., Gupta, A.: Neil: Extracting visual knowledge from web data. In: *International Conference on Computer Vision (ICCV)*, pp. 1409–1416. IEEE (2013)
12. Collins, B., Deng, J., Li, K., Fei-Fei, L.: Towards scalable dataset construction: an active learning approach. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*. LNCS, vol. 5302, pp. 86–98. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88682-2_8](https://doi.org/10.1007/978-3-540-88682-2_8)
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: *Computer Vision and Pattern Recognition (CVPR)* (2009)
14. Deng, J., Krause, J., Fei-Fei, L.: Fine-grained crowdsourcing for fine-grained recognition. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 580–587 (2013)
15. Divvala, S.K., Farhadi, A., Guestrin, C.: Learning everything about anything: webly-supervised visual concept learning. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 3270–3277. IEEE (2014)
16. Duan, K., Parikh, D., Crandall, D., Grauman, K.: Discovering localized attributes for fine-grained recognition. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 3474–3481. IEEE
17. Erkan, A.N.: Semi-supervised learning via generalized maximum entropy. Ph.D. thesis, New York University (2010)

18. Farrell, R., Oza, O., Zhang, N., Morariu, V.I., Darrell, T., Davis, L.S.: Birdlets: Subordinate categorization using volumetric primitives and pose-normalized appearance. In: International Conference on Computer Vision (ICCV), pp. 161–168. IEEE (2011)
19. Fergus, R., Fei-Fei, L., Perona, P., Zisserman, A.: Learning object categories from internet image searches. *Proc. IEEE* **98**(8), 1453–1466 (2010)
20. Gavves, E., Fernando, B., Snoek, C.G., Smeulders, A.W., Tuytelaars, T.: Fine-grained categorization by alignments. In: International Conference on Computer Vision (ICCV), pp. 1713–1720. IEEE
21. Gavves, E., Fernando, B., Snoek, C.G., Smeulders, A.W., Tuytelaars, T.: Local alignments for fine-grained categorization. *Int. J. Comput. Vision (IJCV)*, 1–22 (2014)
22. Goering, C., Rodner, E., Freytag, A., Denzler, J.: Nonparametric part transfer for fine-grained recognition. In: Computer Vision and Pattern Recognition (CVPR), pp. 2489–2496. IEEE (2014)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
24. Hinchliff, C.E., Smith, S.A., Allman, J.F., Burleigh, J.G., Chaudhary, R., Coghill, L.M., Crandall, K.A., Deng, J., Drew, B.T., Gazis, R., Gude, K., Hibbett, D.S., Katz, L.A., Laughinghouse, H.D., McTavish, E.J., Midford, P.E., Owen, C.L., Ree, R.H., Rees, J.A., Soltis, D.E., Williams, T., Cranston, K.A.: Synthesis of phylogeny and taxonomy into a comprehensive tree of life. *Proc. Nat. Acad. Sci.* (2015). <http://www.pnas.org/content/early/2015/09/16/1423041112.abstract>
25. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning (ICML) (2015)
26. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Neural Information Processing Systems (NIPS) (2015)
27. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: First Workshop on Fine-Grained Visual Categorization, Conference on Computer Vision and Pattern Recognition (CVPR), Colorado Springs, CO, June 2011
28. Krause, J., Gebu, T., Deng, J., Li, L.J., Fei-Fei, L.: Learning features and parts for fine-grained recognition. In: International Conference on Pattern Recognition (ICPR), Stockholm, Sweden, August 2014
29. Krause, J., Jin, H., Yang, J., Fei-Fei, L.: Fine-grained recognition without part annotations. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE
30. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13). IEEE (2013)
31. Kumar, N., Belhumeur, P.N., Biswas, A., Jacobs, D.W., Kress, W.J., Lopez, I.C., Soares, J.V.: Leafsnap: a computer vision system for automatic plant species identification. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) European Conference on Computer Vision (ECCV), vol. 7573, pp. 502–516. Springer, Heidelberg (2012)
32. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
33. Lewis, D.D., Catlett, J.: Heterogeneous uncertainty sampling for supervised learning. In: International Conference on Machine Learning (ICML), pp. 148–156 (1994)
34. Li, L.J., Fei-Fei, L.: Optimol: automatic online picture collection via incremental model learning. *Int. J. Comput. Vision (IJCV)* **88**(2), 147–168 (2010)

35. Lin, T.Y., RoyChowdhury, A., Maji, S.: Bilinear cnn models for fine-grained visual recognition. In: International Conference on Computer Vision (ICCV). IEEE
36. Liu, J., Kanazawa, A., Jacobs, D., Belhumeur, P.: Dog breed classification using part localization. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7572, pp. 172–185. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33718-5_13](https://doi.org/10.1007/978-3-642-33718-5_13)
37. Maji, S., Kannala, J., Rahtu, E., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. Technical report (2013)
38. Mnih, V., Hinton, G.E.: Learning to label aerial images from noisy data. In: International Conference on Machine Learning (ICML), pp. 567–574 (2012)
39. Mozafari, B., Sarkar, P., Franklin, M., Jordan, M., Madden, S.: Scaling up crowd-sourcing to very large datasets: a case for active learning. *Proc. VLDB Endowment* **8**(2), 125–136 (2014)
40. Nilsback, M.E., Zisserman, A.: A visual vocabulary for flower classification. In: Computer Vision and Pattern Recognition (CVPR), vol. 2, pp. 1447–1454. IEEE (2006)
41. Pu, J., Jiang, Y.-G., Wang, J., Xue, X.: Which looks like which: exploring inter-class relationships in fine-grained visual categorization. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8691, pp. 425–440. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10578-9_28](https://doi.org/10.1007/978-3-319-10578-9_28)
42. Reed, S., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., Rabinovich, A.: Training deep neural networks on noisy labels with bootstrapping (2014). arXiv preprint [arXiv:1412.6596](https://arxiv.org/abs/1412.6596)
43. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vision (IJCV)*, 1–42, April 2015
44. Schroff, F., Criminisi, A., Zisserman, A.: Harvesting image databases from the web. *Pattern Anal. Mach. Intell. (PAMI)* **33**(4), 754–766 (2011)
45. Sermanet, P., Frome, A., Real, E.: Attention for fine-grained categorization (2014). arXiv preprint [arXiv:1412.7054](https://arxiv.org/abs/1412.7054)
46. Settles, B.: Active learning literature survey. *Univ. Wis. Madison* **52**(55–66), 11 (2010)
47. Settles, B., Craven, M., Ray, S.: Multiple-instance active learning. In: Advances in Neural Information Processing Systems (NIPS), pp. 1289–1296 (2008)
48. Shih, K.J., Mallya, A., Singh, S., Hoiem, D.: Part localization using multi-proposal consensus for fine-grained categorization. In: British Machine Vision Conference (BMVC) (2015)
49. Simon, M., Rodner, E.: Neural activation constellations: unsupervised part model discovery with convolutional networks. In: ICCV (2015)
50. Simon, M., Rodner, E., Denzler, J.: Part detector discovery in deep convolutional neural networks. In: Asian Conference on Computer Vision (ACCV), vol. 2, pp.162–177 (2014)
51. Sukhbaatar, S., Fergus, R.: Learning from noisy labels with deep neural networks (2014). arXiv preprint [arXiv:1406.2080](https://arxiv.org/abs/1406.2080)
52. Szegedy, C., Ioffe, S., Vanhoucke, V.: Inception-v4, inception-resnet and the impact of residual connections on learning (2016). arXiv preprint [arXiv:1602.07261](https://arxiv.org/abs/1602.07261)
53. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Computer Vision and Pattern Recognition (CVPR) (2015)

54. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE (2016)
55. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: The new data and new challenges in multimedia research (2015). arXiv preprint [arXiv:1503.01817](https://arxiv.org/abs/1503.01817)
56. Torralba, A., Efros, A., et al.: Unbiased look at dataset bias. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 1521–1528. IEEE (2011)
57. Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P., Perona, P., Belongie, S.: Building a bird recognition app. and large scale dataset with citizen scientists: the fine print in fine-grained dataset collection. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE (2015)
58. Vedaldi, A., Mahendran, S., Tsogkas, S., Maji, S., Girshick, B., Kannala, J., Rahtu, E., Kokkinos, I., Blaschko, M.B., Weiss, D., Taskar, B., Simonyan, K., Saphra, N., Mohamed, S.: Understanding objects in detail with fine-grained attributes. In: *Computer Vision and Pattern Recognition (CVPR)* (2014)
59. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 dataset. Technical report CNS-TR-2011-001, California Institute of Technology (2011)
60. Wah, C., Belongie, S.: Attribute-based detection of unfamiliar classes with humans in the loop. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 779–786. IEEE (2013)
61. Wah, C., Branson, S., Perona, P., Belongie, S.: Multiclass recognition and part localization with humans in the loop. In: *International Conference on Computer Vision (ICCV)*, pp. 2524–2531. IEEE (2011)
62. Wah, C., Horn, G., Branson, S., Maji, S., Perona, P., Belongie, S.: Similarity comparisons for interactive fine-grained categorization. In: *Computer Vision and Pattern Recognition (CVPR)* (2014)
63. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1386–1393 (2014)
64. Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-UCSD Birds 200. Technical report CNS-TR-2010-001, California Institute of Technology (2010)
65. Xiao, T., Xu, Y., Yang, K., Zhang, J., Peng, Y., Zhang, Z.: The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE
66. Xiao, T., Xia, T., Yang, Y., Huang, C., Wang, X.: Learning from massive noisy labeled data for image classification. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE
67. Xie, S., Yang, T., Wang, X., Lin, Y.: Hyper-class augmented and regularized deep learning for fine-grained image classification. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE
68. Xu, Z., Huang, S., Zhang, Y., Tao, D.: Augmenting strong supervision using web data for fine-grained categorization. In: *International Conference on Computer Vision (ICCV)* (2015)
69. Yang, L., Luo, P., Loy, C.C., Tang, X.: A large-scale car dataset for fine-grained categorization and verification. In: *Computer Vision and Pattern Recognition (CVPR)*. IEEE

70. Yang, S., Bo, L., Wang, J., Shapiro, L.G.: Unsupervised template learning for fine-grained object recognition. In: *Advances in Neural Information Processing Systems (NIPS)*, pp. 3122–3130 (2012)
71. Yao, B., Bradski, G., Fei-Fei, L.: A codebook-free and annotation-free approach for fine-grained image categorization. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 3466–3473. IEEE (2012)
72. Yao, B., Khosla, A., Fei-Fei, L.: Combining randomization and discrimination for fine-grained image categorization. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 1577–1584. IEEE (2011)
73. Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J.: Construction of a large-scale image dataset using deep learning with humans in the loop (2015). arXiv preprint [arXiv:1506.03365](https://arxiv.org/abs/1506.03365)
74. Zhang, N., Donahue, J., Girshick, R., Darrell, T.: Part-based r-cnns for fine-grained category detection. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8689, pp. 834–849. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1_54](https://doi.org/10.1007/978-3-319-10590-1_54)
75. Zhang, N., Farrell, R., Darrell, T.: Pose pooling kernels for sub-category recognition. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 3665–3672. IEEE (2012)
76. Zhang, N., Farrell, R., Iandola, F., Darrell, T.: Deformable part descriptors for fine-grained recognition and attribute prediction. In: *International Conference on Computer Vision (ICCV)*, pp. 729–736. IEEE (2013)
77. Zhang, Y., Wei, X-S., Wu, J., Cai, J., Lu, J., Nguyen, V.A., Do, M.N.: Weakly supervised fine-grained image categorization (2015). arXiv preprint [arXiv:1504.04943](https://arxiv.org/abs/1504.04943)