



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

On Label Granularity and Object Localization

Citation for published version:

Cole, E, Wilber, K, Van Horn, G, Yang, X, Fornoni, M, Perona, P, Belongie, S, Howard, A & Mac Aodha, O 2022, On Label Granularity and Object Localization. in S Avidan, G Brostow, M Cissé, GM Farinella & T Hassner (eds), Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X. Lecture Notes in Computer Science, vol. 13670, Springer, Cham, pp. 604-620, European Conference on Computer Vision 2022, Tel Aviv, Israel, 23/10/22. https://doi.org/10.1007/978-3-031-20080-9_35

Digital Object Identifier (DOI):

[10.1007/978-3-031-20080-9_35](https://doi.org/10.1007/978-3-031-20080-9_35)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



On Label Granularity and Object Localization

Elijah Cole^{1,†} Kimberly Wilber² Grant Van Horn³ Xuan Yang²
 Marco Fornoni² Pietro Perona¹ Serge Belongie⁴
 Andrew Howard² Oisin Mac Aodha⁵

¹Caltech ²Google ³Cornell University
⁴University of Copenhagen ⁵University of Edinburgh

[†]ecole@caltech.edu

Abstract. Weakly supervised object localization (WSOL) aims to learn representations that encode object location using only image-level category labels. However, many objects can be labeled at different levels of granularity. Is it an animal, a bird, or a great horned owl? Which image-level labels should we use? In this paper we study the role of label granularity in WSOL. To facilitate this investigation we introduce iNat-Loc500, a new large-scale fine-grained benchmark dataset for WSOL. Surprisingly, we find that choosing the right training label granularity provides a much larger performance boost than choosing the best WSOL algorithm. We also show that changing the label granularity can significantly improve data efficiency.

1 Introduction

For many problems in computer vision, it is not enough to know *what* is in an image, we also need to know *where* it is. Examples can be found in many domains, including ecological conservation [21], autonomous driving [57], and medical image analysis [31]. The most popular paradigm for locating objects in images is object *detection*, which aims to predict a bounding box for every instance of every category of interest. Object *localization* is special case of detection where each image is assumed to contain exactly one object instance of interest, and the category of that object is known.

Standard approaches to object detection and localization require bounding boxes for training, which are expensive to collect at scale [39]. Weakly supervised object localization (WSOL) methods aim to sidestep this obstacle by learning to localize objects using only image-level labels at training time. The potential reduction in annotation cost which could result from effective weakly supervised methods has stimulated significant interest in WSOL over the last few years [62].

In this paper we explore the role of label granularity in WSOL. The *granularity* of a category is the degree to which it is specific, which can vary from coarse-grained (e.g. “animal”) to fine-grained (e.g. “great horned owl”) [56]. When we work with benchmark datasets in computer vision, we often take the given level of label granularity for granted. However, it is usually possible to make those

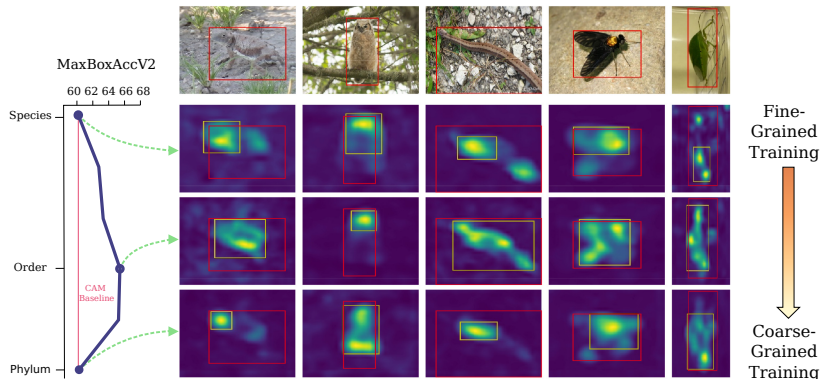


Fig. 1. Label granularity is a critical but understudied factor in weakly supervised object localization (WSOL). We show five hand-picked examples from our iNatLoc500 dataset. Below each image we show class activation maps (CAMs) [65] derived from training a classifier at different granularity levels, with ground truth bounding boxes (red) and WSOL-based bounding boxes (yellow) superimposed. Conventional training does not consider label granularity and can lead to inferior localization performance (red line). Better WSOL results can be achieved by training with coarse (i.e. “order”) labels, as opposed to fine-grained (i.e. “species”) ones.

labels more general or more specific. It is worth asking whether the label granularity we are given is the best one to use for a certain task. Label granularity matters for WSOL because the first step in most WSOL algorithms is to train a classifier using image-level category labels. By choosing a label granularity we are choosing which training images are grouped into categories. This affects the discriminative features learned by the classifier and ultimately determines the bounding box predictions. Is it possible to improve WSOL performance by controlling label granularity?

Unfortunately, it is difficult to explore label granularity in WSOL due to the limitations of existing datasets. The field of WSOL largely relies on CUB [54] and ImageNet [43]. CUB has a consistent label hierarchy (i.e. one that can be used to measure label granularity), but it is small ($\sim 6k$ training images) and homogeneous (only bird categories). ImageNet is large and diverse, but lacks a consistent label hierarchy (see Sec. 4.2). Furthermore, [12] recently found that many purported algorithmic advances in WSOL over the last few years – which were based on these two datasets – perform no better than baselines when they are evaluated fairly. This calls for the development of more diverse and challenging benchmarks for WSOL.

Our primary contributions are as follows:

1. We explore the effect of label granularity on WSOL, and show that training at coarser levels of granularity leads to surprisingly large performance gains across many different WSOL methods compared to conventional training e.g. +5.1 MaxBoxAccV2 for CAM and +6.6 MaxBoxAccV2 for CutMix (see Fig. 3)

2. We demonstrate that training on coarse labels is more data efficient than conventional training. For instance, training at a coarser level achieves the same performance as conventional CAM with $\sim 15\times$ fewer labels (see Fig. 4).
3. We introduce the iNaturalist Localization 500 (iNatLoc500) dataset, which consists of 138k images for weakly supervised training and 25k images with manually verified bounding boxes for validation and testing. iNatLoc500 covers 500 diverse categories with a consistent hierarchical label space.

2 Related Work

Here we primarily focus on literature related to WSOL. See [62] for a broader overview of related techniques such as weakly supervised object detection [6,7,49].

Weakly Supervised Object Localization. The goal of WSOL is to determine the location of single objects in images using only image-level labels at training time. Early attempts at WSOL explored a variety of different approaches, such as adapting boosting-based methods [36], framing the problem as multiple instance learning [20,22], and applying latent deformable part-based formulations [38].

Some foundational work in deep learning investigated the degree to which object localization comes “for free” when training supervised CNNs for image classification tasks [61,37,65]. In particular, the Class Activation Mapping (CAM) method of [65] showed that CNNs can capture some object location information even when they are trained using only image-level class labels. This inspired a large body of work (e.g. [63,47,64,13,27,28]) that attempted to address some of the shortcomings of CAM, e.g. by preventing the underlying model from only focusing on the most discriminative parts of an object [60] or increasing the spatial resolution of its outputs [45,10].

Recently, [12] showed that when state-of-the-art WSOL methods are fairly compared (e.g. by controlling for the backbone architecture and operating thresholds), they are no better than the standard CAM [65] baseline. Thus, despite its simplicity, CAM is still a surprisingly effective baseline for WSOL. Subsequent work has explored further techniques for improving CAM-based methods [2,29] and alternative approaches for estimating model coefficients [25].

Task Granularity and Localization. Despite the considerable interest in WSOL in recent years, many open questions remain. Examples include the effect of label granularity (e.g. coarse-grained labels like “bird” vs. fine-grained labels indicating the specific species of bird) and the effect of training set size. In the context of supervised object detection, [53] showed that *coarsening* category labels at training time can improve the localization performance of *object detectors*. It is unclear if the same phenomenon holds for WSOL. [55] explored the impact of label granularity for object detection on the OpenImages [30] dataset and observed a small performance improvement when training on finer labels. In the semi-supervised detection setting, [59] trained object detectors on OpenImages and ImageNet using both coarse-grained bounding box annotations and fine-grained image-level labels. [51] also explored semi-supervised detection with an approach that generates object proposals across multiple hierarchical

levels. Unlike our work, these detection-based methods require bounding box information at training time. In addition, the label hierarchies for datasets like ImageNet and OpenImages are not necessarily good proxies for visual similarity or concept granularity (see Sec. 4.2).

For WSOL, [28] showed that aggregating class attribution maps at coarser hierarchical levels (e.g. “dog”) results in more spatial coverage of the objects of interest, whereas maps for finer-scale concepts (e.g. “Afghan hound”) only focus on subparts of the object. However, their analysis does not explore the impact of training at different granularity levels. It is also worth noting that their aggregation method only improves performance on CUB. Regarding data quantity, [12] studied the number of supervised examples used to tune the hyperparameters of CAM, but did not consider the impact of the number of examples used to train the image classifier.

Though not directly related to our work, we note that label granularity has been studied in many contexts other than object localization, including action recognition [46], knowledge tracing [14], animal face alignment [26], and fashion attribute recognition [23]. In the context of image classification, prior work has tackled topics like analyzing the emergence of hierarchical structure in trained classifiers [5], identifying patterns in visual concept generalization [44], and training finer-grained image classifiers using only coarse-grained labels [48,42,58,50].

Datasets for Object Localization. Early work in WSOL (e.g. [36,20,34]) focused on relatively simple and small-scale datasets such as Caltech4 [19], the Weizmann Horse Database [8], or subsets of PASCAL-VOC [18]. With the rise of deep learning-based methods, CUB [54] and ImageNet [17,43] became the standard benchmarks for this task. CUB [54] consists of images of 200 different categories of birds, where each image contains a single bird instance. ImageNet [17,43] contains 1000 diverse categories and has significantly more images than CUB (>1M compared to ~6k). [12] proposed OpenImages30k, a 100-category localization-focused subset of the OpenImages V5 dataset [30]. An overview of these datasets is presented in Table 1.

These existing datasets are valuable, but they have shortcomings. CUB is small and homogeneous (only birds). OpenImages30k, as presented in [12], is not actually evaluated as a bounding box localization task. It is instead a per-pixel foreground object segmentation task where the ground truth also features some “ignore” regions that are excluded from the evaluation. Finally, while both OpenImages30k and ImageNet have label hierarchies, they do not reflect concept granularity in a consistent way. As a result, it is difficult to use them to better understand the relationship between concept granularity and localization. We discuss these issues in greater detail in Sec. 4.2. To address these shortcomings we introduce iNatLoc500, a new WSOL dataset composed of images from 500 fine-grained visual categories and equipped with a consistent label hierarchy.

Table 1. Comparison of datasets for WSOL. The vast majority of WSOL papers use only CUB and ImageNet. The OpenImages30k dataset was introduced by [12], which also defines the splits we use for CUB and ImageNet. For each split we provide the minimum, maximum, and mean number of images per category, along with the total number of images in the split. Means are rounded to the nearest integer. The properties of these four datasets are discussed in detail in Sec. 4.2.

Dataset	# Cat.	train-weaksup (D_w)				train-fullsup (D_f)				test (D_{test})			
		Min	Max	Mean	Total	Min	Max	Mean	Total	Min	Max	Mean	Total
CUB [54]	200	29	30	30	6k	3	6	5	1k	11	30	29	5.8k
ImageNet [17]	1000	732	1300	1281	1.28M	10	10	10	10k	10	10	10	10k
OpenImages30k [3,12]	100	230	300	298	30k	25	25	25	2.5k	50	50	50	5k
iNatLoc500	500	149	307	276	138k	25	25	25	12.5k	25	25	25	12.5k

3 Background

3.1 Weakly Supervised Object Localization (WSOL)

We begin by formalizing the WSOL setting. Let D_w be a set of *weakly labeled* images, i.e. $D_w = \{(x_i, y_i)\}_{i=1}^{N_w}$ where $x_i \in \mathbb{R}^{H \times W \times 3}$ is an image and $y_i \in \{1, \dots, C\}$ is an image-level label corresponding to one of C categories. Let D_f be a set of *fully labeled* images, i.e. $D_f = \{(x_i, y_i, \mathbf{b}_i)\}_{i=1}^{N_f}$ where x_i and y_i are defined as before and $\mathbf{b}_i \in \mathbb{R}^4$ is a bounding box for an instance of category y_i . In practice $N_w \gg N_f$. WSOL approaches typically comprise three steps:

(1) **Train.** Use D_w to train an image classifier $h_\theta : \mathbb{R}^{H \times W \times 3} \rightarrow [0, 1]^C$ by solving

$$\hat{\theta}(D_w) = \operatorname{argmin}_{\theta} \frac{1}{|D_w|} \sum_{(x_i, y_i) \in D_w} \mathcal{L}(h_\theta(x_i), y_i)$$

where \mathcal{L} is some training loss and θ represents the parameters of h . Different WSOL methods are primarily distinguished by the loss functions and training protocols they use to train h .

(2) **Localize.** For each $(x_i, y_i, \mathbf{b}_i) \in D_f$, predict a bounding box

$$\hat{\mathbf{b}}_i = g(x_i, y_i | h_{\hat{\theta}(D_w)})$$

according to some procedure $g : \mathbb{R}^{H \times W \times 3} \times \{1, \dots, C\} \rightarrow \mathbb{R}^4$. Typically g is a simple sequence of image processing operations applied to the feature maps of the trained classifier $h_{\hat{\theta}(D_w)}$.

(3) **Evaluate.** Let E denote a suitable WSOL error metric which compares the predicted boxes $\{\hat{\mathbf{b}}_i\}_{i=1}^{N_f}$ against the ground-truth boxes $\{\mathbf{b}_i\}_{i=1}^{N_f}$. Use the validation error $E(D_f | D_w)$ for model selection and hyperparameter tuning and then use a held-out test set D_{test} (which is fully labeled like D_f) to measure test error $E(D_{\text{test}} | D_w)$. See [12] for a discussion of WSOL performance metrics.

The role of low-shot supervised localization. Without the fully labeled images D_f , the WSOL problem becomes ill-posed [12]. Since WSOL therefore

requires at least a small number of bounding box annotations for validation, it is natural to ask how WSOL compares to few-shot object localization? For our purposes, we define few-shot object localization methods as those which use only D_f for training and validation. Under this definition, the few-shot methods (which use only D_f) actually require strictly *less* data than WSOL (which requires both D_w and D_f). Since WSOL and few-shot object localization are practical alternatives, it is important to consider them together as in [12].

3.2 Label Hierarchies and Label Granularity

We define a *label hierarchy* (on a label set L) to be a directed rooted tree H whose leaf nodes (i.e. nodes $v \in H$ with no children) correspond to the labels in L . Edges in H represent “is-a” relationships, so a directed edge from $u \in H$ to $v \in H$ means that v (e.g. “bird”) is a kind of u (e.g. “animal”). We overload L to refer to the label set and to the corresponding set of nodes in H . Let r denote the root node of H and let $d(u, v)$ denote the number of edges on the path from $u \in H$ to $v \in H$.

Coarsening a label. Because there is a unique path from the root node r to any leaf node $\ell \in L$, we can “coarsen” the label ℓ in a well-defined way by merging it with its parent node. We define the *coarsening operator* $c_k : H \rightarrow H$, which takes any node in the label hierarchy and returns the node which is k edges closer to the root. Thus, $c_0(\ell) = \ell$, $c_1(\ell)$ is the parent of ℓ , $c_2(\ell)$ is the grandparent of ℓ , and so on, with $c_k(\ell) = r$ for all $k \geq d(r, \ell)$.

Coarsening a dataset. We can describe a general “coarsened” version of $D_w = \{(x_i, y_i)\}_{i=1}^{N_w}$ as $D_w^{\mathbf{k}} = \{(x_i, c_{k_i}(y_i))\}_{i=1}^{N_w}$ where $\mathbf{k} = (k_1, \dots, k_{|D_w|})$. If we allow the entries of \mathbf{k} to be chosen completely independently, then we can encounter problems e.g. images with multiple valid labels. To prevent these cases, we require \mathbf{k} to be chosen such that $c_{k_i}(y_i) \in H$ is not a descendant of $c_{k_j}(y_j) \in H$ for any $i, j \in \{1, \dots, N_w\}$.

Problem statement. We can now formalize our key questions: How does \mathbf{k} affect $E(D_{\text{test}}|D_w^{\mathbf{k}})$? Are there choices of \mathbf{k} such that $E(D_{\text{test}}|D_w^{\mathbf{k}}) < E(D_{\text{test}}|D_w)$?

4 The iNatLoc500 Dataset

In this section we introduce the iNaturalist Localization 500 (iNatLoc500) dataset, a large-scale fine-grained dataset for weakly supervised object localization. We first detail the process of building the dataset and cleaning the localization annotations. We then discuss the key properties of the dataset and highlight the advantages of iNatLoc500 compared to three WSOL datasets that are currently commonly used (CUB, ImageNet, and OpenImages30k).

iNatLoc500 has three parts: **train-weaksup** (D_w), **train-fullsup** (D_f), and **test** (D_{test}). Each image in the weakly supervised training set (D_w) has one image-level category label. Each image in the fully supervised validation set (D_f) and test set (D_{test}) has one image-level category label *and* one bounding box annotation. All bounding boxes have been manually validated. Split statistics

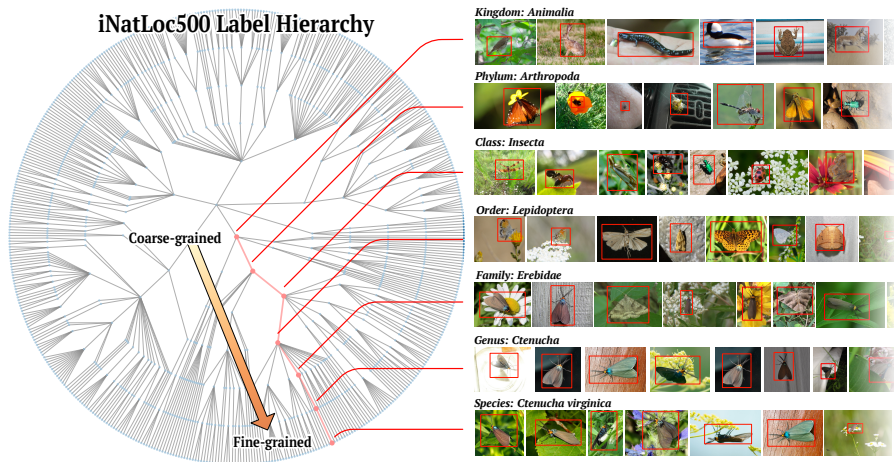


Fig. 2. Sample images from iNatLoc500 at different levels of the label hierarchy, from coarse (“kingdom”) to fine (“species”). Random images from coarse levels of the hierarchy tend to be much more varied than random images ones from finer levels.

are presented in Table 1 and sample images from the dataset can be found in Fig. 2. The dataset is publicly available.¹

4.1 Dataset Construction

The iNatLoc500 dataset is derived from two existing datasets: iNat17 [53] and iNat21 [52]. Both datasets contain images of plants and animals collected by the citizen science platform iNaturalist [1]. iNat21 is much larger than iNat17 (2.7M images, 10k species vs. 675k images, 5k species), but iNat17 has crowdsourced bounding box annotations. We draw from iNat21 for D_w and we draw from iNat17 for D_f and D_{test} .

Full details on the process of constructing iNatLoc500 can be found in the supplementary material, but we note two important design choices here. First, iNat17 did not collect bounding boxes for plant categories because it is often unclear how to draw bounding boxes for plants. Consequently, iNatLoc500 does not contain any plant categories. Second, we set very high quality standards for the bounding boxes. Five computer vision researchers manually reviewed $\sim 65k$ images to ensure the quality of the bounding boxes for D_f and D_{test} , of which only 51% met our quality standards. Explicit quality criteria and examples of removed images can be found in the supplementary material.

4.2 Dataset Properties

iNatLoc500 is fine-grained, large-scale, and visually diverse. Moreover, iNatLoc500 has a consistent label hierarchy which serves as a reliable proxy for

¹ https://github.com/visipedia/inat_loc/

label granularity. We now discuss the importance of each of these properties and contrast iNatLoc500 with existing WSOL datasets.

Fine-grained categories. Each category in iNatLoc500 corresponds to a different species, and the differences between species can be so subtle as to require expert-level knowledge [53]. While there are challenging images in ImageNet and OpenImages30k, most of the categories are coarse-grained i.e. relatively few pairs of categories are highly visually similar. For instance, the reptile categories in OpenImages30k (`lizard`, `snake`, `frog`, `crocodile`) are typically easy to distinguish. In iNatLoc-500 there are 107 reptile species, some of which are highly similar (e.g. `Chihuahuan spotted whiptail` vs. `Common spotted whiptail`).

Consistent label hierarchy. The label hierarchy for iNatLoc500 consists of the following seven tiers, ordered from coarsest to finest: kingdom, phylum, class, order, family, genus, and species. All of the species in iNatLoc500 are animals, so the “kingdom” tier only has one node (`Animalia`), which is the root node of the label hierarchy. Every species lies at the same distance from the root. The iNatLoc500 label hierarchy is *consistent* in the sense that all nodes at a given level of the hierarchy correspond to concepts with similar levels of specificity. This means that depth in the label hierarchy measures label granularity. The label hierarchy for CUB is also consistent. However, the taxonomies that underlie ImageNet and OpenImages30k are considerably more arbitrary. For instance, in OpenImages30k some categories are far from the root of the label hierarchy (e.g. `entity/vehicle/land_vehicle/car/limousine` or `entity/animal/mammal/carnivore/fox`) while others are close to the root (e.g. `entity/bicycle_wheel` or `entity/human_ear`) despite the fact that there is no obvious difference in concept specificity.

Unambiguous label semantics. The categories in iNatLoc500 are well-defined in the sense that (for most species) there is little room for debate about what “counts” as an instance of that species. While the distinctions between species can be quite subtle, each species is a well-defined category. CUB shares this advantage for the most part, but ImageNet and OpenImages30k do not. For instance, OpenImages30k contains the categories `wine` and `bottle`. To which category does a bottle of wine belong? (In fact, we find bottles of wine in both categories.) ImageNet is known to have similar issues with ambiguous and overlapping category definitions [4].

Visual diversity. Like ImageNet and OpenImages30k, iNatLoc500 has a category set which exhibits a high degree of visual diversity. CUB is much more homogeneous, consisting of only birds. Combined with its consistent label hierarchy, the visual diversity of iNatLoc500 enables future work on e.g. how localization ability generalizes across categories as a function of taxonomic distance.

Large scale. iNatLoc500 is a large-scale dataset, both in terms of the number of categories and the number of training images. CUB and OpenImages30k are considerably smaller on both counts. Large training sets are valuable because they simplify supervised learning. Large training sets also enable research on self-supervised representation learning, which has received little attention thus far in WSOL. We provide a summary of the key dataset statistics in Table 1.

5 Experiments

In this section we present WSOL results on iNatLoc500 as well as existing benchmark datasets. We also consider few-shot learning baselines based on segmentation and detection architectures. Finally, we use the unique properties of iNatLoc500 to study how label granularity affects localization performance and data efficiency. A summary of the different WSOL datasets can be found in Table 1.

5.1 Implementation Details

Performance metrics. All WSOL performance numbers in this paper are **MaxBoxAccV2**, which is defined in [12]. The only exceptions are the results for OpenImages30k in Table 2, which are given in **PxAP** as defined in [12].

Fixed-granularity training. In Sec. 5.3 we probe the effect of granularity on WSOL by training on “coarsened” versions of D_w . In the notation of Sec. 3.2, these can be written $D_w^{k \cdot \mathbf{1}}$ for $k = 1, 2, \dots$, where $\mathbf{1}$ denotes the “all ones” vector. This corresponds to merging all leaves with their parent k times. We then run the entire WSOL pipeline from scratch to compute $E(D_{\text{test}} | D_w^{k \cdot \mathbf{1}})$ for each k . To the best of our knowledge this is compatible with all existing WSOL methods.

Fixed-granularity CAM aggregation. We also consider a second method for using label hierarchy information to improve WSOL, inspired by [28]. Just like traditional CAM, the first step is to train an image classifier using the standard (most fine-grained) label set. However, instead of returning only the CAM for the species labeled in the input image, we return a CAM for each species in the same genus / family / ... / phylum and average them. This “aggregated” CAM is then evaluated as normal. We abbreviate this method as CAM-Agg.

Hyperparameter search for WSOL methods. Each time we train a WSOL method we re-tune the learning rate over the set $\{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$ and choose the one that leads to the best **MaxBoxAccV2** performance on the fully supervised validation set D_f . We then report the **MaxBoxAccV2** performance for the selected model on D_{test} . We leave all other hyperparameters fixed. Full training details can be found in the supplementary material.

Non-WSOL methods. We provide results for the baselines proposed in [12] (Center, FSL-Seg), as well as a new few-shot detection baseline (FSL-Det). “Center” is a naive baseline that simply assumes a centered Gaussian activation map for all images. “FSL-Seg” is a supervised baseline that is trained on the D_f split of each dataset. The architecture is based on models for saliency mask prediction [35]. Finally, we introduce “FSL-Det”, a few-shot detection baseline for WSOL that is also trained on D_f . It uses Faster-RCNN [41] with the same backbone as other methods (i.e. ImageNet-pretrained ResNet-50 [24]). Full implementation details can be found in the supplementary material.

5.2 Baseline Results

We follow [12] and evaluate six recent WSOL methods and two non-WSOL methods (Center and FSL-Seg) on iNatLoc500. The results can be found in Table 2.

Table 2. Comparison of WSOL methods. Numbers are **MaxBoxAccV2** for ImageNet, CUB, and iNatLoc500 and **PxAP** for OpenImages30k. All results use an ImageNet-pretrained ResNet-50 [24] backbone with an input resolution of 224x224. WSOL numbers for ImageNet, CUB, and OpenImages30k are the updated results from [11]. WSOL numbers for iNatLoc500 are our own, as are the numbers for the baselines (Center, FSL-Seg, FSL-Det). FSL baselines use 10 images / class for ImageNet, 5 images / class for CUB, 25 images / class for OpenImages30k, and 25 images / class for iNatLoc500. We do not report FSL-Det for OpenImages30k because the evaluation protocol for that dataset requires segmentation masks.

Method	ImageNet	CUB	OpenImages30k	iNatLoc500
CAM [65]	63.7	63.0	58.5	60.2
HaS [47]	63.4	64.7	55.9	60.0
ACoL [63]	62.3	66.5	57.3	55.3
SPG [64]	63.3	60.4	56.7	60.7
ADL [13]	63.7	58.4	55.2	58.9
CutMix [60]	63.3	62.8	57.7	60.1
Center	53.4	56.8	46.0	42.8
FSL-Seg	68.7	89.4	75.2	78.6
FSL-Det	70.4	95.4	-	83.6

We focus our observations on ImageNet, CUB, and iNatLoc500 since OpenImages30k is evaluated using a different task and evaluation metric. We first note that our findings on iNatLoc500 reinforce the main results from [12], namely that (a) none of the WSOL methods performs substantially better than CAM and (b) FSL-Seg significantly outperforms all WSOL methods. Second, if we consider the performance gap between CAM and the Center baseline, we see that simple centered boxes are not as successful on iNatLoc500 (-17.2 **MaxBoxAccV2**) as they are on CUB (-6.2 **MaxBoxAccV2**) and ImageNet (-10.3 **MaxBoxAccV2**). This indicates that iNatLoc500 is a more challenging dataset for benchmarking WSOL. Finally, we provide results for our few-shot detection baseline (FSL-Det). For ImageNet, CUB, and iNatLoc500 we find that FSL-Det is a stronger baseline than FSL-Seg. Like FSL-Seg, FSL-Det directly trains on the boxes in D_f , whereas the WSOL methods only use those boxes to tune their hyperparameters. However, FSL-Det sets a new ceiling for localization performance on these datasets, indicating that current WSOL methods have considerable room for improvement.

5.3 Label Granularity and Localization Performance

iNatLoc500 is equipped with a consistent label hierarchy which allows us to directly study the relationship between label granularity and localization performance. The traditional approach to WSOL on iNatLoc500 would begin by training a classifier on the *species-level* labels, i.e. the finest level in the label hierarchy. However, our hypothesis is that training at the most fine-grained level may not lead to the best localization performance. To study this, we use the fixed-granularity training method discussed in Sec. 5.1. In particular, we “re-label” D_w at each level of the label hierarchy using successively coarser categories. We

then use each of these re-labeled datasets to train and evaluate different WSOL methods. The results in Fig. 3(left) show that coarsening the labels of D_w can significantly boost WSOL performance (e.g. up to +5.1 **MaxBoxAccV2** for CAM). The numerical values plotted in Fig. 3(left) can be found in the supplementary materials. Note that it would be difficult to draw similar conclusions by studying ImageNet or OpenImages30k because their label hierarchies do not measure how fine-grained different categories are – see Sec. 4.2 for a discussion. Our conceptually simple coarsening approach results in large performance improvements across five different WSOL methods, without any modifications to the model architectures or training losses.

Coarse training beyond iNatLoc500. Fig. 3(left) shows that coarse training significantly improves WSOL performance on iNatLoc500. We study the effect of coarse training on FGVC-Aircraft [32], CUB [54], and ImageNet [17] in the supplementary material. As expected, FGVC-Aircraft and CUB (which have consistent label hierarchies) both benefit from coarse training while ImageNet (which lacks a consistent label hierarchy) does not.

Localization performance vs. classification performance. In Fig. 3(right) we show the image classification performance for each WSOL method in Fig. 3(left) at each granularity level. We see that classification performance and WSOL performance are not necessarily correlated. WSOL performance increases before decreasing at the coarsest level of granularity. Classification performance increases with label coarsening, even at the coarsest level of granularity.

An alternative method for incorporating label granularity. We also present the performance of CAM-Agg, an alternative method for incorporating granularity information in WSOL (see Sec. 5.1). In our experiments, CAM-Agg underperforms vanilla CAM at every granularity level. As a point of comparison, [28] finds that CAM-Agg is better than CAM for CUB but worse than CAM for ImageNet. Our findings suggest that training the model with coarse categories leads to much better localization performance when compared to aggregating the localization outputs for multiple similar fine-grained categories.

5.4 Label Granularity and Data Efficiency

Most WSOL work makes D_w as large as possible by default, so there has been little attention paid to how the size of D_w trades off against localization performance. In this section we analyze the performance of CAM-based WSOL as a function of the size of D_w . We are particularly interested in how label granularity interacts with data efficiency. To study this question, we first pick a granularity level and generate subsampled versions of D_w by choosing, uniformly at random, 50, 100, or 200 images from each category. Note that the size of each subsampled version of D_w depends on the granularity level. For instance, if the categories are the 317 genera, then 50 images per category is $50 \times 317 = 15,850$, compared to $50 \times 61 = 3,050$ images if the categories are the 61 orders. We present WSOL results for four granularity levels in Fig. 4. We find that by training at a coarser level, we can obtain better performance with fewer labels. All of the square markers above the dashed line in Fig. 4 correspond to cases where we can

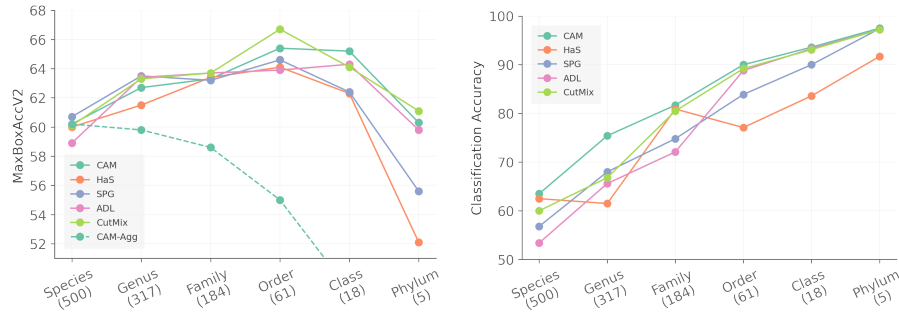


Fig. 3. Effect of label granularity of D_w on WSOL performance (left) and classification accuracy (right) for iNatLoc500. The number of categories at each tier is given in parentheses. **(Left)** Localization performance suffers when the category labels are either too fine (e.g. Species) or too coarse (e.g. Phylum). The results on the very left of the plot are the same as those in Table 2. Note that ACoL is excluded due to poor performance – we suspect it requires more epochs of training than the standard protocol allows for iNatLoc500. We also show results for CAM-Agg (Sec. 5.1), an alternative method for aggregating hierarchy information in WSOL. **(Right)** Each WSOL method trains the image classifier in a different way, but classification accuracy generally increases as the labels become more coarse. Naturally it is easier to distinguish between coarser categories, but it is interesting to note that classification performance is excellent at the phylum level, despite poor localization performance.

achieve better performance than the standard species-level CAM approach using fewer labels. To take one example, by training at the family level we can match the performance of the standard CAM approach by training with 50 images per family (9200 images), a training set reduction of $\sim 15\times$.

6 Discussion

Why does performance increase as we coarsen the labels? In Fig. 3(left) we see that five different WSOL algorithms perform better as we coarsen the labels in D_w , up until the coarsest level when performance drops. What accounts for this behavior? Our analysis of CAM in Fig. 5 provides some clues. Fig. 5(left) shows that the area of the predicted box tends to be larger than the area of the ground truth box, and that their ratio *decreases* towards unity as we coarsen the labels (black curve). That is, the predicted box size gets closer to the true box size as we coarsen the labels. This casts doubt on a common intuition (which as far as we know has not been empirically investigated before now) that WSOL methods predict smaller boxes for more fine-grained categories [28].

Why does performance drop at the coarsest level of granularity? In Fig. 5(left) we see that as we coarsen the labels the concentration of *activation* in the ground truth box *increases* before collapsing at the coarsest level (red curve). Fig. 5(right) shows that the activation maps become highly fragmented

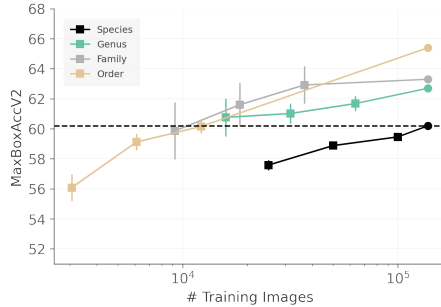


Fig. 4. Effect of the number of training images (N_w) on CAM performance for iNatLoc500. The dashed line corresponds to the performance of species-level CAM with the entirety of D_w . Each color corresponds to a different label granularity for D_w . Circles at the right of the graph indicate performance using all of D_w . Squares represent subsampled datasets which use a fixed number of images per category: 50, 100, or 200. All squares have error bars indicating the standard deviation over 5 runs with different randomly sampled subsets of D_w .

at coarser levels. Taken together, these two findings suggest that at the coarsest level the activation maps tend to focus more on global image characteristics (e.g. land vs. water) than the properties of the foreground object. Note that these features are still useful for image classification, as is shown in Fig. 3(right).

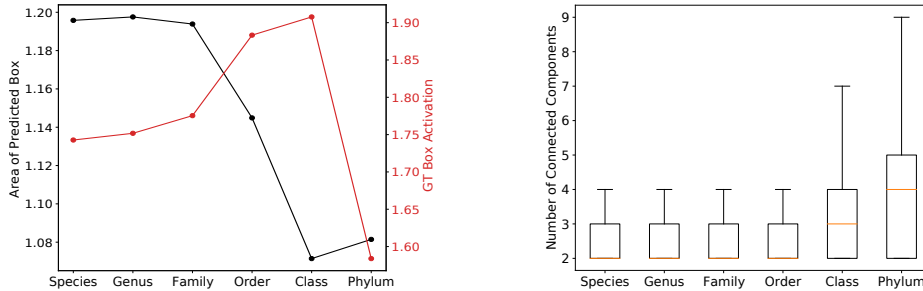


Fig. 5. Analysis of CAM-based WSOL on the D_f split of iNatLoc500. **(Left)** *Black*: Ratio of the area of the predicted box to the area of the ground truth box. *Red*: Ratio of the activation inside the ground truth box to the activation of background pixels. Both curves show medians over the 12.5k images in D_f at each granularity level. **(Right)** Number of connected components in the binarized activation maps at each granularity level. Each box plot shows the distribution over the 12.5k images in D_f . See the supplementary material for full details on the construction of these plots.

Limitations. The iNatLoc500 dataset has several limitations. First, it contains only animal categories. These categories are highly diverse, but they are not representative of all visual domains. Second, it is possible that there are errors in the image-level labels provided by the iNaturalist community, though this is expected to be rare as each image has been labeled by multiple people [52]. Third, many real fine-grained problems have a long-tailed class distribution but, like other localization datasets, iNatLoc500 is approximately balanced (at the

species level). Finally, there is a conceptual limitation in our experiments: the use of a single granularity level across the entire dataset. In fact, it is likely that different images are best treated at different granularity levels. Our work does not address this important topic which we leave for future work.

iNatLoc500 can be used to investigate numerous research agendas beyond traditional WSOL. For example, D_w was designed to be large enough for self-supervised learning, which has received surprisingly little attention in the WSOL community [9]. We are also interested in using iNatLoc500 to study whether self-supervised learning methods can be improved by using WSOL methods to select crops [40], especially in the context of fine-grained data [15]. For the object detection community, the clean boxes in iNatLoc500 can (i) serve as a test set for object detectors trained on the noisy iNat17 boxes, (ii) be used to study the problem of learning multi-instance detectors from one box per image, and (iii) be used to analyze the role of label granularity in object detection. Finally, we have seen that hierarchical reasoning can significantly improve localization performance. In the future, we aim to explore methods for automatically determining the most appropriate level of coarseness required for generating representations that best encode object location.

7 Conclusion

We have shown that substantial improvements in WSOL performance can be achieved by modulating the granularity of the training labels, and that coarser-grained training leads to more data-efficient WSOL. We also presented iNatLoc500, a new large-scale fine-grained dataset for WSOL. Despite the gains in performance from coarse-level training, iNatLoc500 remains a challenging localization task which we hope will motivate additional progress in WSOL.

Acknowledgements. We thank the iNaturalist community for sharing images and species annotations. This work was supported by the Caltech Resnick Sustainability Institute, an NSF Graduate Research Fellowship (grant number DGE1745301), and the Pioneer Centre for AI (DNRF grant number P1).

References

1. iNaturalist, www.inaturalist.org, accessed Mar 7 2022
2. Bae, W., Noh, J., Kim, G.: Rethinking class activation mapping for weakly supervised object localization. In: ECCV (2020)
3. Benenson, R., Popov, S., Ferrari, V.: Large-scale interactive object segmentation with human annotators. In: CVPR (2019)
4. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., Oord, A.v.d.: Are we done with imagenet? arXiv:2006.07159 (2020)
5. Bilal, A., Jourabloo, A., Ye, M., Liu, X., Ren, L.: Do convolutional neural networks learn class hierarchy? IEEE transactions on visualization and computer graphics **24**(1), 152–162 (2017)
6. Bilen, H., Pedersoli, M., Tuytelaars, T.: Weakly supervised object detection with convex clustering. In: CVPR (2015)

7. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: CVPR (2016)
8. Borenstein, E., Ullman, S.: Class-specific, top-down segmentation. In: ECCV (2002)
9. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
10. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Gradcam++: Generalized gradient-based visual explanations for deep convolutional networks. In: WACV (2018)
11. Choe, J., Oh, S.J., Chun, S., Akata, Z., Shim, H.: Evaluation for weakly supervised object localization: Protocol, metrics, and datasets. arXiv:2007.04178 (2020)
12. Choe, J., Oh, S.J., Lee, S., Chun, S., Akata, Z., Shim, H.: Evaluating weakly supervised object localization methods right. In: CVPR (2020)
13. Choe, J., Shim, H.: Attention-based dropout layer for weakly supervised object localization. In: CVPR (2019)
14. Choi, Y., Lee, Y., Shin, D., Cho, J., Park, S., Lee, S., Baek, J., Bae, C., Kim, B., Heo, J.: Ednet: A large-scale hierarchical dataset in education. In: International Conference on Artificial Intelligence in Education. pp. 69–73. Springer (2020)
15. Cole, E., Yang, X., Wilber, K., Mac Aodha, O., Belongie, S.: When does contrastive visual representation learning work? In: CVPR (2022)
16. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9268–9277 (2019)
17. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
18. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. IJCV (2010)
19. Fergus, R., Perona, P., Zisserman, A.: Object class recognition by unsupervised scale-invariant learning. In: CVPR (2003)
20. Galleguillos, C., Babenko, B., Rabinovich, A., Belongie, S.: Weakly supervised object localization with stable segmentations. In: ECCV (2008)
21. van Gemert, J.C., Verschoor, C.R., Mettes, P., Epema, K., Koh, L.P., Wich, S.: Nature conservation drones for automatic localization and counting of animals. In: ECCV (2014)
22. Gokberk Cinbis, R., Verbeek, J., Schmid, C.: Multi-fold mil training for weakly supervised object localization. In: CVPR (2014)
23. Guo, S., Huang, W., Zhang, X., Srikhanta, P., Cui, Y., Li, Y., Adam, H., Scott, M.R., Belongie, S.: The imaterialist fashion attribute dataset. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
24. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
25. Jung, H., Oh, Y.: Towards better explanations of class activation mapping. In: ICCV (2021)
26. Khan, M.H., McDonagh, J., Khan, S., Shahabuddin, M., Arora, A., Khan, F.S., Shao, L., Tzimiropoulos, G.: Animalweb: A large-scale hierarchical dataset of annotated animal faces. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6939–6948 (2020)
27. Ki, M., Uh, Y., Lee, W., Byun, H.: In-sample contrastive learning and consistent attention for weakly supervised object localization. In: ACCV (2020)

28. Kim, J.M., Choe, J., Akata, Z., Oh, S.J.: Keep calm and improve visual feature attribution. In: ICCV (2021)
29. Kim, J., Choe, J., Yun, S., Kwak, N.: Normalization matters in weakly supervised object localization. In: ICCV (2021)
30. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4. IJCV (2020)
31. Li, Z., Dong, M., Wen, S., Hu, X., Zhou, P., Zeng, Z.: Clu-cnns: Object detection for medical images. Neurocomputing (2019)
32. Maji, S., Kannala, J., Rahtu, E., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. Tech. rep. (2013)
33. Miller, G.A.: Wordnet: a lexical database for english. Communications of the ACM (1995)
34. Nguyen, M.H., Torresani, L., De La Torre, F., Rother, C.: Weakly supervised discriminative localization and classification: a joint learning process. In: ICCV (2009)
35. Oh, S.J., Benenson, R., Khoreva, A., Akata, Z., Fritz, M., Schiele, B.: Exploiting saliency for object segmentation from image level labels. In: CVPR (2017)
36. Opelt, A., Pinz, A.: Object localization with boosting and weak supervision for generic object recognition. In: Scandinavian Conference on Image Analysis (2005)
37. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Is object localization for free?-weakly-supervised learning with convolutional neural networks. In: CVPR (2015)
38. Pandey, M., Lazebnik, S.: Scene recognition and weakly supervised object localization with deformable part-based models. In: ICCV (2011)
39. Papadopoulos, D.P., Uijlings, J.R., Keller, F., Ferrari, V.: Extreme clicking for efficient object annotation. In: ICCV (2017)
40. Peng, X., Wang, K., Zhu, Z., Wang, M., You, Y.: Crafting better contrastive views for siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16031–16040 (2022)
41. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. NeurIPS (2015)
42. Robinson, J., Jegelka, S., Sra, S.: Strength from weakness: Fast learning using weak supervision. In: ICML (2020)
43. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV (2015)
44. Sariyildiz, M.B., Kalantidis, Y., Larlus, D., Alahari, K.: Concept generalization in visual representation learning. In: ICCV (2021)
45. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: ICCV (2017)
46. Shao, D., Zhao, Y., Dai, B., Lin, D.: Finegym: A hierarchical video dataset for fine-grained action understanding. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2616–2625 (2020)
47. Singh, K.K., Lee, Y.J.: Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: ICCV (2017)
48. Taherkhani, F., Kazemi, H., Dabouei, A., Dawson, J., Nasrabadi, N.M.: A weakly supervised fine label classifier enhanced by coarse supervision. In: ICCV (2019)
49. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: CVPR (2017)
50. Touvron, H., Sablayrolles, A., Douze, M., Cord, M., Jégou, H.: Graft: Learning fine-grained image representations with coarse labels. In: ICCV (2021)

51. Uijlings, J., Popov, S., Ferrari, V.: Revisiting knowledge transfer for training object class detectors. In: CVPR (2018)
52. Van Horn, G., Cole, E., Beery, S., Wilber, K., Belongie, S., Mac Aodha, O.: Benchmarking representation learning for natural world image collections. In: CVPR (2021)
53. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: CVPR (2018)
54. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011)
55. Wang, R., Mahajan, D., Ramanathan, V.: What leads to generalization of object proposals? In: ECCV Workshops (2020)
56. Wei, X.S., Song, Y.Z., Mac Aodha, O., Wu, J., Peng, Y., Tang, J., Yang, J., Belongie, S.: Fine-grained image analysis with deep learning: A survey. PAMI (2021)
57. Wu, B., Iandola, F., Jin, P.H., Keutzer, K.: Squeezedet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In: CVPR Workshops (2017)
58. Xu, Y., Qian, Q., Li, H., Jin, R., Hu, J.: Weakly supervised representation learning with coarse labels. In: ICCV (2021)
59. Yang, H., Wu, H., Chen, H.: Detecting 11k classes: Large scale object detection without fine-grained bounding boxes. In: ICCV (2019)
60. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: ICCV (2019)
61. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV (2014)
62. Zhang, D., Han, J., Cheng, G., Yang, M.H.: Weakly supervised object localization and detection: A survey. PAMI (2021)
63. Zhang, X., Wei, Y., Feng, J., Yang, Y., Huang, T.S.: Adversarial complementary learning for weakly supervised object localization. In: CVPR (2018)
64. Zhang, X., Wei, Y., Kang, G., Yang, Y., Huang, T.: Self-produced guidance for weakly-supervised object localization. In: ECCV (2018)
65. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)

A Additional Experiments

A.1 Does coarse training help on other datasets?

In the main paper we show that training on coarser labels significantly improves WSOL performance on iNatLoc500. It is reasonable to ask whether we can replicate this pattern on other datasets. In Fig. A1 we present results on FGVC-Aircraft [32], CUB [54], and ImageNet [17]. We give dataset-specific details below.

FGVC-Aircraft. The FGVC-Aircraft dataset consists of images of different kinds of aircraft, which are organized into a label hierarchy with the following tiers, ordered from coarsest to finest: manufacturer, family, and variant. Fig. A1(top) shows that training with coarser labels improves WSOL performance (+4.3 **MaxBoxAccV2**). This shows that the benefits of coarse training are not limited to natural world datasets.

Training details: We use the best hyperparameters for CUB from [12], except that we train for 10 epochs and decay the learning rate every 3 epochs.

CUB. Fig. A1(middle) shows that training with coarser labels improves WSOL performance (+4.4 **MaxBoxAccV2**). This indicates that our observations on iNatLoc500 in the main paper generalize to images collected under different protocols i.e. iNaturalist user photos vs. iconic images crawled from Flickr. Unlike iNatLoc500, we do not see a drop in performance at the coarsest level. This is consistent with our previous findings because CUB contains only birds, so its hierarchy terminates before reaching the level of granularity where iNatLoc500 performance drops.

Training details: We use the filtered version of CUB as described in Sec. C and train with the best hyperparameters for CUB from [12].

ImageNet. Fig. A1(bottom) shows no benefit to coarsening the labels for ImageNet. This is consistent with our claim in the main paper that the ImageNet hierarchy does not measure granularity, and motivates the development of better label hierarchies for datasets like ImageNet in future work.

Training details: Unlike iNatLoc500, FGVC-Aircraft, and CUB, the ImageNet label hierarchy has leaf nodes at many different depths. To accommodate this, we must modify the coarsening procedure from the main paper. Instead of coarsening every leaf node at every step, we only coarsen leaf nodes which are at the deepest level of the hierarchy. Each granularity level is named **cX** where $X \in \{0, 1, 2, \dots\}$ is the number of times this coarsening has been applied. To speed up training we sample 200 images per category for D_w . We use the filtered version of ImageNet as described in Sec. C and train with the best hyperparameters for ImageNet from [12].

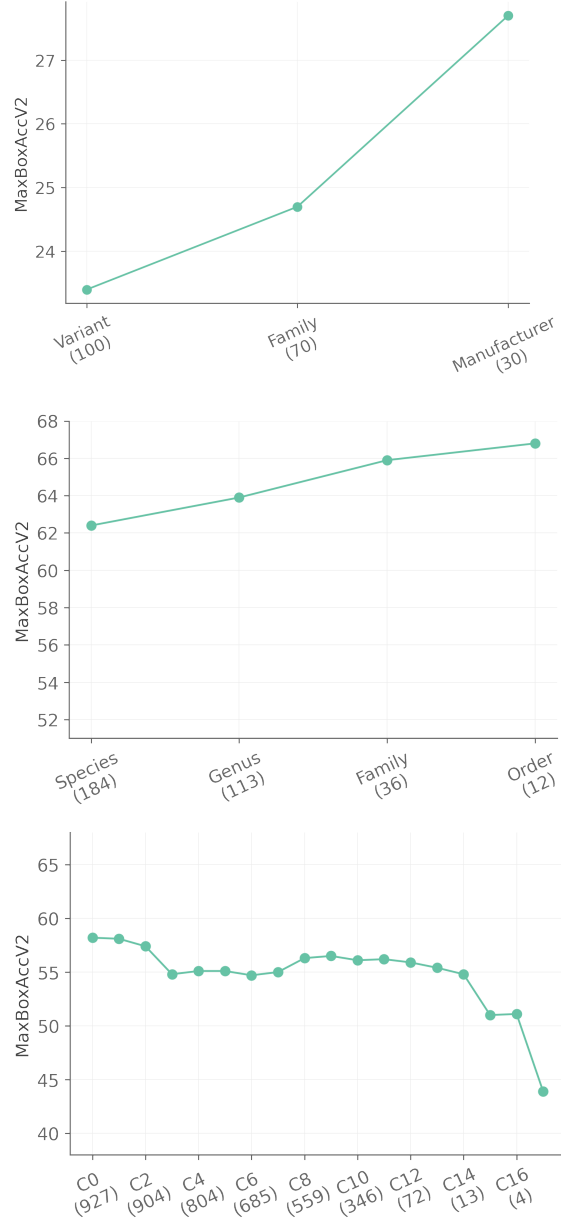


Fig. A1. WSOL performance as a function of training set label granularity for FGVC-Aircraft (top), CUB (middle) and ImageNet (bottom). Like iNatLoc500, FGVC-Aircraft and CUB show significant gains at coarser granularities. There is no apparent benefit for ImageNet, which lacks a consistent label hierarchy.

A.2 What is the impact of longer training schedules?

In [12], the authors design their WSOL training schedules so that CUB and OpenImages30k use the same computational budget. They use a budget of 300k images processed, which equates to 50 epochs for CUB and 10 epochs for OpenImages30k. To respect this budget, in the main paper we train on iNatLoc500 for 2 epochs (276k images processed). In Fig. A2 we see that a longer training schedule can improve performance slightly (Family, Phylum) or significantly (Species, Genus, Order, Class). However, the pattern is the same whether we train for 2 epochs or 10 epochs, i.e. performance drops for labels that are too coarse or too fine.

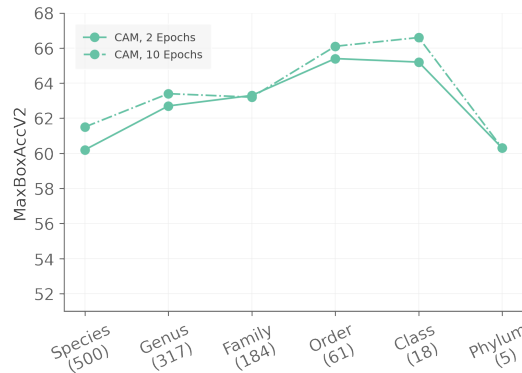


Fig. A2. Comparison of our standard training schedule (2 epochs, reducing learning rate after 1 epoch) and a longer training schedule (10 epochs, reducing learning rate every 3 epochs) for CAM on iNatLoc500. Training for longer does not change the observation that coarser labels result in better localization.

A.3 How does WSOL performance depend on the IoU threshold?

Throughout the main paper we use the `MaxBoxAccV2` metric proposed by [12]. This metric averages performance over three IoU thresholds: 30%, 50%, and 70%. In Fig. A3 we show the performance of CAM on iNatLoc500 separately for each IoU threshold. Not surprisingly, we see that performance decreases significantly as the IoU threshold becomes more demanding (i.e. larger). We also observe that, regardless of the IoU threshold, the best performance is obtained at a label granularity that is neither too fine nor too coarse. In the right panel of Fig. A3 we see that the relative performance improvement is larger for more demanding IoU thresholds. This may be because there is less room to improve for “easier” IoU thresholds.

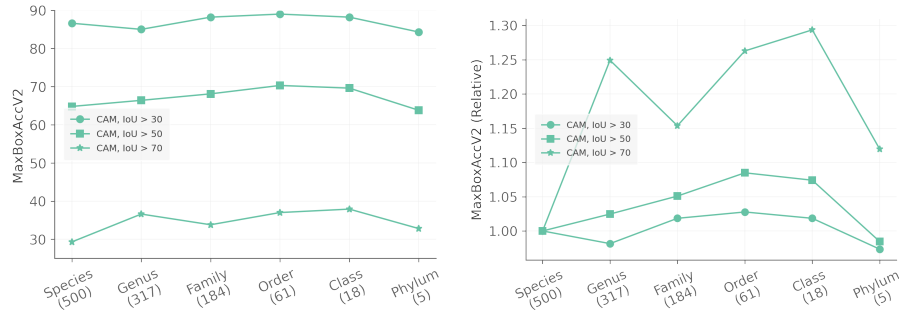


Fig. A3. CAM performance on iNatLoc500 as a function of label granularity and IoU threshold. The left panel shows absolute performance. The right panel shows performance relative to the species-level performance, which is the traditional baseline approach. More specifically, the right panel is generated by normalizing each curve in the left panel by its left-most endpoint.

A.4 How stable are the CAM results?

Each result in the main paper is the result of a single run, so it is important to quantify how much test performance varies when we re-train. In Fig. A4 we show the results of re-training CAM on iNatLoc500 five times at each granularity level with identical hyperparameters. The standard deviations at different granularity levels range from ~ 0.2 to ~ 0.8 , which is much smaller than the effect sizes we discuss in the main paper. Interestingly, training seems to be most stable for the best-performing coarse-grained levels (order and class), and least stable for the genus level.

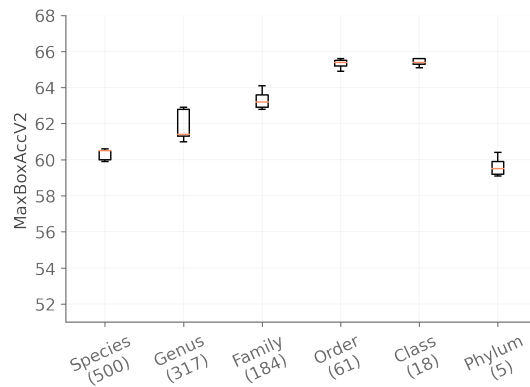


Fig. A4. Distribution of CAM performance at each granularity level of iNatLoc500 for five runs with identical hyperparameters. The orange line denotes the mean.

A.5 What is the effect of additional hyperparameter tuning?

In their paper, [12] searches over 30 random hyperparameter sets for each WSOL method. We use a less computationally intensive protocol. For iNatLoc500, we start from their best hyperparameters for ImageNet and re-optimize the learning rate by searching over $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$. To quantify the performance difference between our reduced hyperparameter search and the full search in [12], we conduct both procedures on iNatLoc500 using CAM at the species level. The full hyperparameter search (30 hyperparameter sets) achieves 60.8 **MaxBoxAccV2** compared to 60.2 **MaxBoxAccV2** for our abbreviated hyperparameter search (5 hyperparameter sets). As expected, the additional hyperparameter optimization provides an improvement for CAM but the difference is surprisingly modest. We would expect a similar boost to occur for any granularity level. We also note that the gap may be greater for methods with more hyperparameters to tune. We provide the learning rates used in our paper in Table A3.

B Dataset Construction Details for iNatLoc500

In this section we detail the process of merging and cleaning data from iNat17 and iNat21 to produce iNatLoc500.

Species matching. In total there are 4486 species names that occur in both iNat17 and iNat21. We discard any images which do not correspond to a species shared by both datasets. We also omit any species that does not have bounding box annotations. In particular, this means that we discard all plant species, since iNat17 does not have any bounding boxes for plants.

Removing duplicate observations. Each image on the iNaturalist platform has an associated `observation_id` which corresponds to a unique encounter with an individual plant or animal. We find all observation IDs which occur in both iNat17 and iNat21 and we remove all of the corresponding images from iNat21. It is important to remove duplicates at the `observation_id` level instead of the image level, since an iNaturalist observation may be associated with multiple similar but distinct images of the same individual organism.

Instance count filtering. Since our focus is object localization (as opposed to detection), any images with multiple bounding box annotations are discarded.

Box size filtering. Any image whose box is smaller than 32 pixels in either dimension is removed. In addition, any image whose box width (height) is more than 96% of the image width (height) is removed. This step speeds up the annotation process by filtering out a significant number of “bad” images. Very small boxes are problematic because annotators are more likely to make mistakes, while very large boxes tend to be extreme close-ups.

Split considerations. While the majority of observations on iNaturalist are associated with only one image, some do have multiple images. When splitting the fully supervised images into D_f and D_{test} we ensure that all of the images for one observation go into exactly one split. This is important because images from the same observation can be highly similar.

Manual Annotation. Well-annotated validation and test sets are essential for reliable model selection and benchmarking. The image-level fine-grained class labels reflect the consensus of the iNaturalist community, and like prior iNaturalist datasets [53,52] we assume they are correct. However, the bounding box annotations were crowd-sourced with non-expert workers. We therefore manually validate the bounding box annotations for the images in the D_f and D_{test} splits. Images with any of the issues listed below were excluded from the dataset. Note that the distribution of images in D_f and D_{test} is likely to be somewhat different than the distribution of images in D_w due to this cleaning process. Examples of problematic images are given in the supplementary material.

- *Missing instances.* Images with multiple bounding box annotations are filtered out before annotation cleaning. Unfortunately, the bounding box annotations for an image are sometimes incomplete, which means that an image with one bounding box annotation for a species can contain multiple instances of that species. Images with multiple instances of the labeled species are removed.

- *Inaccurate bounding boxes.* Some bounding boxes are too large or too small, e.g. boxes which miss appendages such as legs or tails or boxes which only contain the face of the animal. Images with inaccurate bounding boxes are removed. We also remove any images for which it is unclear whether or not the bounding box is correct, which may occur when an image is blurry or poorly illuminated.

- *Indirect evidence.* iNaturalist accepts images showing *indirect evidence* of an animal (e.g. footprints, feathers, droppings), not just images of the animal itself. We omit images which show only indirect evidence of an animal. We also omit images of animal carcasses, which are not uncommon for e.g. deer.

- *Body part close-ups.* Some images in iNaturalist are clearly intended to show the structure of some specific body part in scientific detail, such as an image of a paw next to a ruler. We omit these images.

C Label Hierarchies

We visualize the label hierarchies for CUB, ImageNet, and iNatLoc500 in Fig. A5. Producing our final hierarchies for CUB and ImageNet required some care. We give details below.

CUB. CUB was not released with a label hierarchy, so we constructed one. We start by attempting to map each category to a node on the tree of life, like iNatLoc500. CUB consists of 200 bird categories, where some of these categories correspond to species (e.g. **Black-footed Albatross**) and some do not (e.g. the genus **Sayornis** or umbrella terms like **frigatebird**). We discard any CUB category whose name could not be unambiguously mapped to a single species. By checking these species names against the iNaturalist taxonomy, we obtained the genus, family, order, and class for each species. All bird species belong to the class **Aves**, so this is the root node of the label hierarchy. Since we retained only species-level categories, every leaf node lies at the same distance from the root node. Our CUB label hierarchy has 184 leaf nodes.

ImageNet. ImageNet is equipped with a label hierarchy based on WordNet [33]. One problem with this hierarchy is that some nodes have multiple parents, which violates the assumptions of the label coarsening procedure outlined in the main paper. We remedy this using a simple greedy approach in which we iterate over the nodes with multiple parents in some fixed (but arbitrary) order and delete all but one parent node. In particular, for each node with multiple parents we perform the following operations:

1. Choose a parent and compute the number of leaf nodes that are still reachable if that parent is retained and the others are deleted. Repeat for each parent.
2. Keep the parent node for which the greatest number of leaf nodes remain reachable from the root.
3. Delete the other parent nodes.
4. Delete any descendants of deleted nodes which are no longer reachable from the root.

After executing this process, we obtain a label hierarchy in which each (non-root) node has a unique parent. Our ImageNet label hierarchy has 927 leaf nodes.

D Descriptive Statistics

D.1 Class Imbalance and Label Granularity

We give basic statistics on the distribution of images over categories for iNatLoc500 at different granularity levels in Table A1. We also visualize the distribution of images over categories at different granularity levels in Fig. A7. At the species level, the categories are approximately balanced, but the spread between the largest and the smallest category is much larger for coarser label sets.

D.2 Box Size

In Fig. A6 we compare the box size distributions for iNatLoc500, ImageNet, and CUB. For each curve in Fig. A6 we compute the area of each box, divide the box areas by the corresponding image sizes, and compute the CDF. The box distribution for iNatLoc500 seems to interpolate between the the box distributions for CUB and ImageNet, e.g. iNatLoc500 has more “small” boxes than CUB but not as many as ImageNet.

E Performance Scores

For ease of comparison we provide the raw **MaxBoxAccV2** scores for each WSOL method (and CAM-Agg) at each granularity level in Table A2.

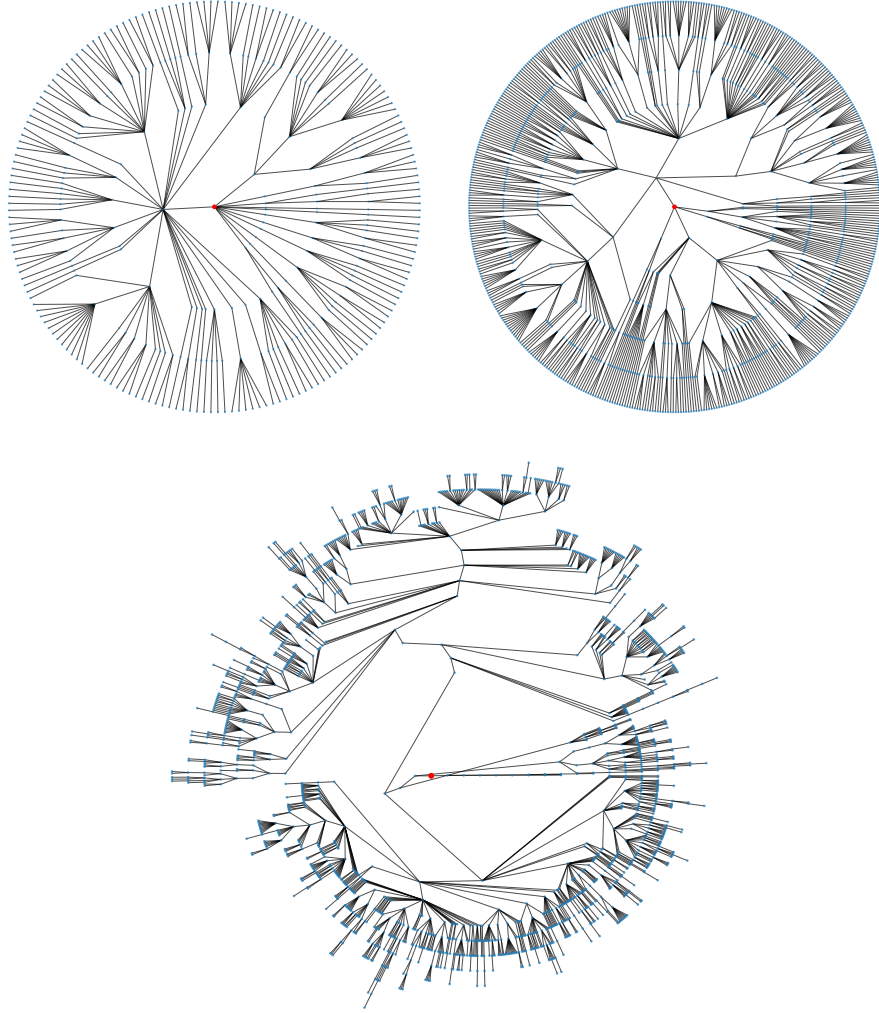


Fig. A5. Label hierarchies for CUB (top left), iNatLoc500 (top right), and ImageNet (bottom). Root nodes are shown in red. The hierarchies for CUB and iNatLoc500 have uniform depths (4 and 6, respectively). The hierarchy for ImageNet is considerably more irregular.

F Implementation Details

F.1 Qualitative Analysis of WSOL

In this section we define the terms used in the qualitative analysis figures in the main paper. In what follows we choose the threshold t to be the optimal threshold for an IoU of 0.50, as defined by [12].

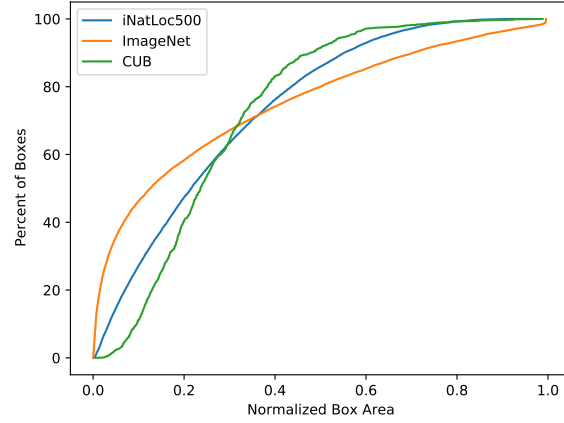


Fig. A6. Comparison of CDFs of box sizes for iNatLoc500, ImageNet, and CUB. All box sizes are normalized by the size of the image.

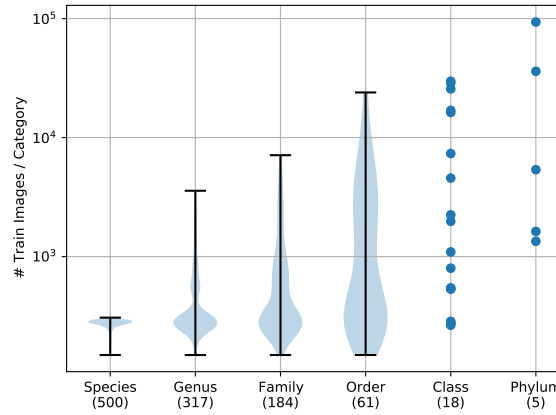


Fig. A7. Distribution of images over categories at different granularity levels for iNatLoc500. We show violin plots for Species, Genus, Family, and Order. Class and Phylum contain only a small number of categories so we can show each point individually. See Table A1 for summary statistics at each granularity level.

- **Area of Predicted Box:** The area of the predicted box divided by the area of the ground truth box. The predicted box is computed using a threshold t .
- **GT Box Activation:** The sum of the heatmap pixels inside the ground truth box divided by the sum of the heatmap pixels outside the ground truth box.

Table A1. Summary statistics for iNatLoc500 at each granularity level. For each granularity level, we provide the number of categories as well as the minimum, maximum, and mean number of images per category. We also calculate the imbalance factor, which is the size of the largest class divided by the size of the smallest class [16]. Refer to Fig. A7 for a visualization of the distribution of images over categories at each granularity level.

Granularity	# Categories	Min	Max	Mean	Imbalance
Species	500	149	307	276	2.1
Genus	317	149	3575	435	24.0
Family	184	149	7113	750	47.7
Order	61	149	23947	2262	160.7
Class	18	265	29741	7666	112.2
Phylum	5	1345	93576	27599	69.6

Table A2. **MaxBoxAccV2** scores for different WSOL methods trained at different levels of granularity. As in the main paper, ACoL is excluded due to poor performance. We also include provide scores for CAM-Agg, an alternative method of using using granularity information for WSOL described in the main paper.

Method	Species	Genus	Family	Order	Class	Phylum
CAM	60.2	62.7	63.3	65.4	65.2	60.3
HaS	60.0	61.5	63.4	64.1	62.3	52.1
ACoL	-	-	-	-	-	-
SPG	60.7	63.5	63.2	64.6	62.4	55.6
ADL	58.9	63.4	63.7	63.9	64.3	59.8
CutMix	60.1	63.3	63.7	66.7	64.1	61.1
CAM-Agg	60.2	59.8	58.6	55.0	48.8	40.1

- **Number of Connected Components:** The number of connected components in the predicted heatmap after it has been binarized with threshold t .

F.2 WSOL Methods

We consider six standard WSOL methods in this work: CAM [65], HaS [47], ACoL [63], SPG [64], ADL [13], and CutMix [60]. We leave the details of those methods to their respective papers. For each WSOL method, we use the training procedures and optimal hyperparameters used by [12] for ImageNet. The only exceptions are as follows. First, we always use enlarged 28×28 feature maps, instead of letting the choice between 14×14 and 28×28 be an additional hyperparameter. Second, we use a weight decay of 10^{-5} instead of 10^{-4} . Third, we always re-optimize the learning rate by searching over the set $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ and using the value of **MaxBoxAccV2** on D_f to select the best one. The best learning rate values for each WSOL method and granularity level are provided in Table A3 and the method-specific hyperparameters can be found in Table A4. We summarize the rest of the training details, which match [12], below.

Table A3. Best learning rates for WSOL methods at different granularity levels. Learning rates are selected from $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ based on the value of **MaxBoxAccV2** on **train-fullsup**. Note that learning rates for ACoL at coarser granularity levels are omitted due to poor performance.

Method	Species	Genus	Family	Order	Class	Phylum
CAM	10^{-2}	10^{-1}	10^{-2}	10^{-2}	10^{-2}	10^{-2}
HaS	10^{-2}	10^{-1}	10^{-2}	10^{-2}	10^{-2}	10^{-2}
ACoL	10^{-3}	-	-	-	-	-
SPG	10^{-1}	10^{-1}	10^{-1}	10^{-1}	10^{-1}	10^{-2}
ADL	10^{-1}	10^{-1}	10^{-1}	10^{-2}	10^{-2}	10^{-2}
CutMix	10^{-2}	10^{-1}	10^{-2}	10^{-2}	10^{-2}	10^{-2}

Table A4. Method-specific hyperparameters used for iNatLoc500. These are the same hyperparameters used by [12] for ImageNet.

Method	Hyperparameters
CAM	N/A
HaS	<code>drop_rate = 0.09</code> , <code>drop_area = 31</code>
ACoL	<code>erasing_threshold = 0.79</code>
SPG	$\delta_l^{B1} = 0.02$, $\delta_h^{B1} = 0.03$, $\delta_l^{B2} = 0.05$, $\delta_h^{B2} = 0.47$, $\delta_l^C = 0.29$, $\delta_h^C = 0.36$
ADL	<code>drop_rate = 0.68</code> , <code>erasing_threshold = 0.93</code>
CutMix	<code>size_prior = 0.10</code> , <code>mix_rate = 0.93</code>

Architecture. All methods use an ImageNet-pretrained ResNet-50 backbone with an input resolution of 224×224 .

Image preprocessing. Training images are resized to 256×256 , randomly cropped to 224×224 , then horizontally flipped with probability 0.5. At test time, images are simply resized to 224×224 . All images are normalized according to ImageNet statistics.

Optimization. We train using SGD with Nesterov momentum, a momentum parameter of 0.9, and a batch size of 32. The learning rate for the final linear classifier layer is set to be $10\times$ larger than the learning rate for the rest of the network. For fairness, [12] trains on each dataset for a number of epochs which equates to processing 300k images. To respect this criterion, we train iNatLoc500 for 2 epochs (276k images processed) and decay the learning rate by a factor of 10 after the first epoch.

Evaluation. The search space for the optimal heatmap threshold consists of 1000 linearly spaced values between 0 and 1. Note that all heatmaps are min-maxed normalized before evaluation, so their values fall in $[0, 1]$. Final **MaxBoxAccV2** numbers are an average over three IoU thresholds: 30, 50, and 70.

F.3 Center Baseline

We perform baseline experiments using the “center” baseline for WSOL introduced in [12], which simply generates a centered Gaussian heatmap for each image. Since [12] did not fully specify the implementation of their center baseline,

our re-implementation may differ slightly. We opt for a simple implementation which does not depend on the image shape. Specifically, we generate an image $C \in \mathbb{R}^{M \times M}$ where

$$C_{i,j} = \exp \left(-\frac{((i - \frac{M-1}{2})^2 + (j - \frac{M-1}{2})^2)}{2\sigma^2} \right)$$

for the pixel in row i and column j . We then apply min-max normalization to C . We set $M = 224$ and $\sigma = M/4$.

Note that in the continuous domain, the value of σ would not matter, since, for any σ , a square centered box of any size could be obtained by choosing the right heatmap threshold. In practice, the heatmap threshold is optimized over a fixed grid of values. In this case, each value of σ yields a different collection of centered boxes, which results in different performance numbers.

F.4 FSL-Seg: Few-Shot Localization via Segmentation

The FSL-Seg baseline for WSOL was introduced in [12], but they did not fully specify the implementation details so our approach may differ. Our training protocols are identical to those we use for WSOL methods, except for the modifications described below.

Architecture. Like the WSOL methods, we begin with an ImageNet-pretrained ResNet-50 with an input resolution of 224×224 . We modify the network by replacing the final fully connected layer with a 1×1 convolution layer with a sigmoid activation. Since the feature maps have shape $2048 \times 28 \times 28$, the output of this modified ResNet-50 is a single “score map” $S \in [0, 1]^{28 \times 28}$.

Loss. We train using a weighted per-pixel binary cross-entropy loss given by

$$\sum_{ij} \left[\frac{Y_{ij}}{\|Y\|_0} \log S_{ij} + \frac{(1 - Y_{ij})}{\|1 - Y\|_0} \log(1 - S_{ij}) \right]$$

where $Y \in \{0, 1\}^{28 \times 28}$ is a binary label mask and $\|Y\|_0$ denotes the number of nonzero values in Y . This weighting has the effect of equally balancing positive and negative labels. For OpenImages30k, binary label masks are directly available. However, CUB, ImageNet, and iNatLoc500 only have bounding box annotations. For these three datasets we compute Y by converting the bounding box annotations into binary masks. Note that these masks are noisy because most objects do not completely fill their bounding boxes.

Optimization. For each dataset we train for 10 epochs and decay the learning rate by a factor of 10 every 3 epochs.

F.5 FSL-Det: Few-Shot Localization via Detection

For FSL-Det we use a Faster-RCNN object detection architecture. We use an off-the-shelf TensorFlow Object Detection API training configuration file originally meant for training a Faster-RCNN model on COCO. Other than changing the input image size and the dataset, we do not modify the architecture or any training procedures. The configuration can be found here:

https://github.com/tensorflow/models/blob/65407126c5adc216d606d360429fe12ed3c3f187/research/object_detection/configs/tf2/faster_rcnn_resnet50_v1_640x640_coco17_tpu-8.config

G Manual Annotation

We performed extensive filtering and quality control to produce **train-fullsup** and **test** splits data for iNatLoc500. We show randomly selected examples from iNatLoc500 in Fig. A8. We also show examples of images which were rejected and give the reason in each case in Fig. A9.

H Qualitative Examples

We show some hand-picked predictions for CAM-based WSOL in Fig. A10.

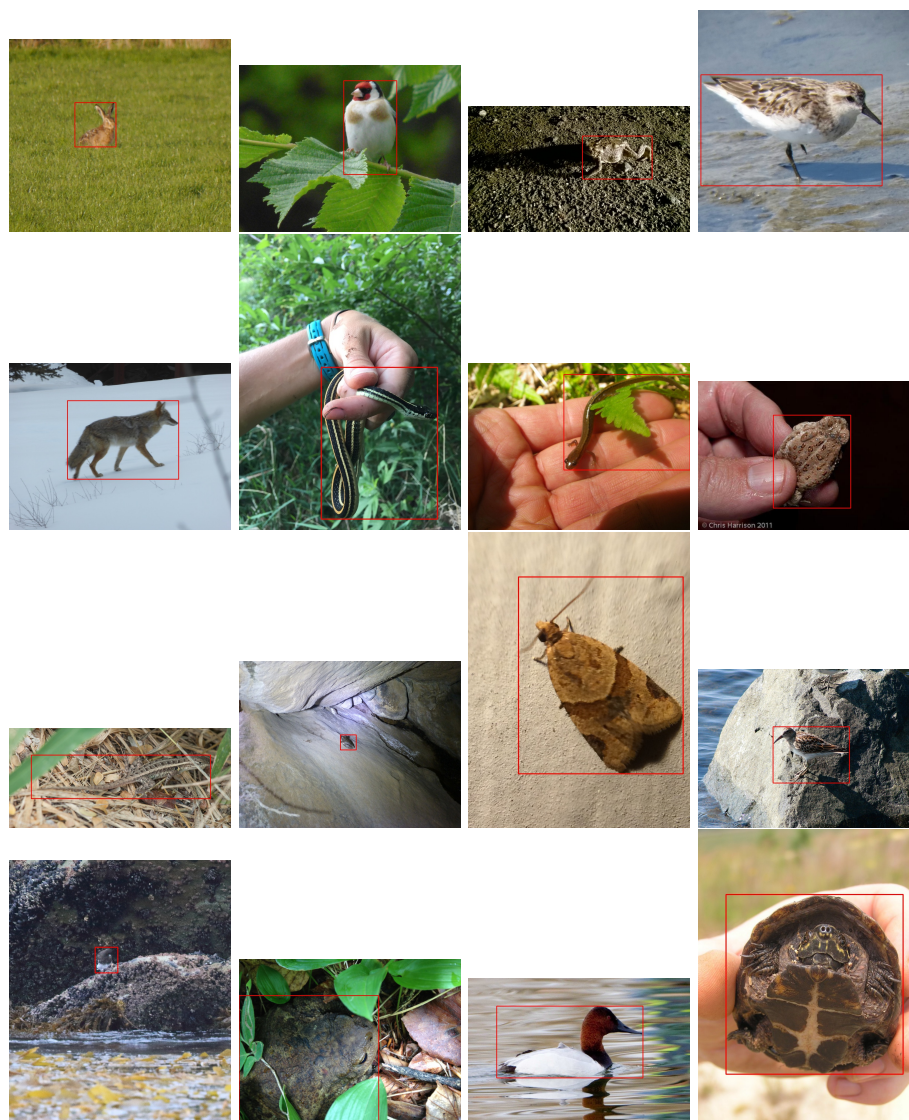


Fig. A8. Randomly selected sample images from the iNatLoc500 dataset.



Fig. A9. Examples of problematic images from iNat17 which were filtered out of the iNatLoc500 dataset. Each image is identified with a tuple (i, j) where i is the row and j is the column. We now describe the problem in each image. (1, 1): The box is too small. (1, 2): The box is too large. (1, 3): The target class is the crab, not the otter, so the box is too large. (1, 4): The box is too large and there are multiple instances of the target species. (2, 1): The box is too large and there are multiple instances of the target species. (2, 2): The image is an extreme close-up. (2, 3): The correct box is ambiguous due to blurring. (2, 4): The box is too large.

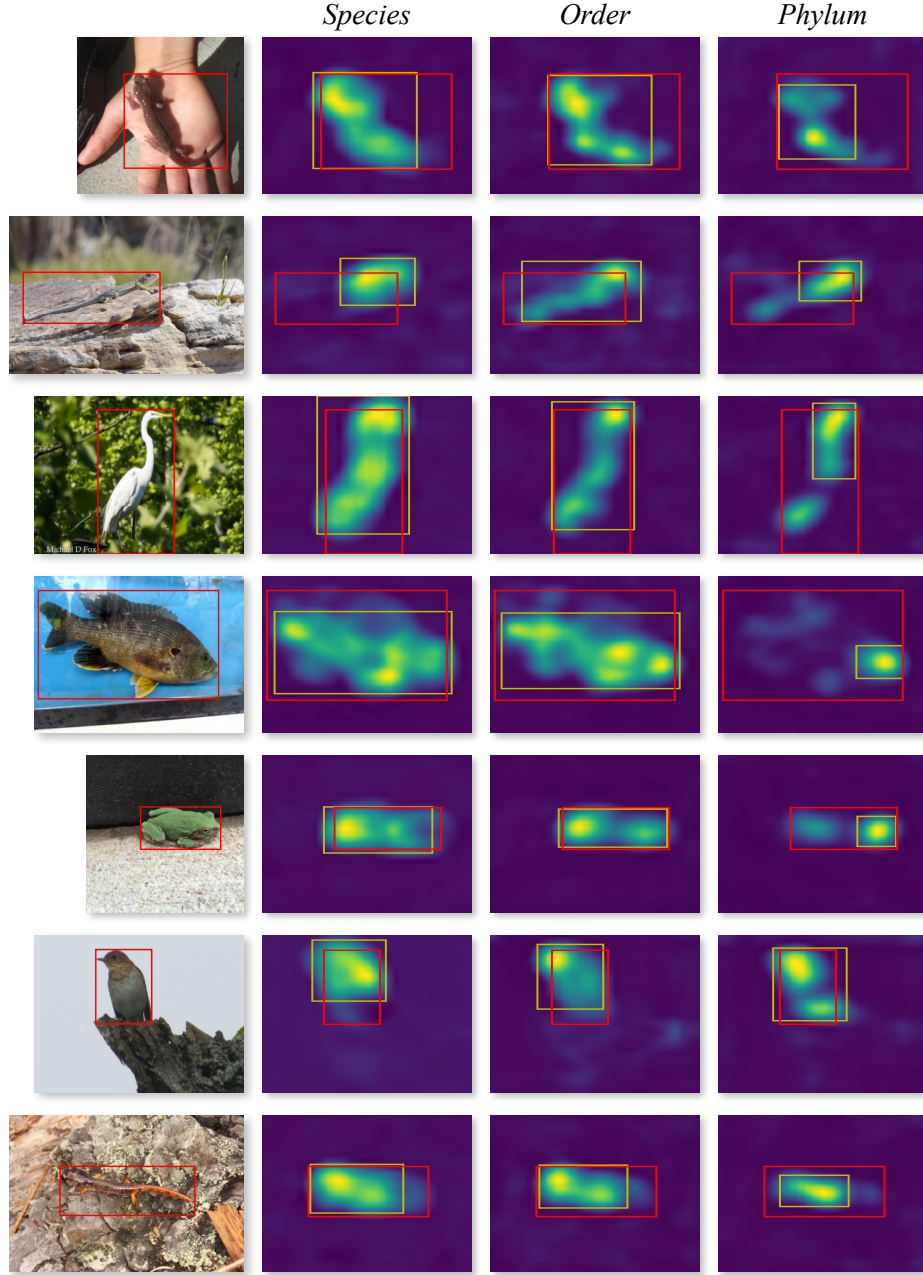


Fig. A10. Examples of CAM-based WSOL predictions at different levels of granularity. In each row we provide activation map for classifiers trained at the phylum, order, and species level. Each activation map shows the ground truth bounding box (red) and WSOL-based bounding box (yellow).