

# Semi-Supervised Hierarchical Open-Set Classification

Erik Wallin<sup>1,2</sup>, Fredrik Kahl<sup>2</sup>, Lars Hammarstrand<sup>2</sup>  
<sup>1</sup>Saab AB, <sup>2</sup>Chalmers University of Technology

{walline,fredrik.kahl,lars.hammarstrand}@chalmers.se

## Abstract

*Hierarchical open-set classification handles previously unseen classes by assigning them to the most appropriate high-level category in a class taxonomy. We extend this paradigm to the semi-supervised setting, enabling the use of large-scale, uncurated datasets containing a mixture of known and unknown classes to improve the hierarchical open-set performance. To this end, we propose a teacher-student framework based on pseudo-labeling. Two key components are introduced: 1) subtree pseudo-labels, which provide reliable supervision in the presence of unknown data, and 2) age-gating, a mechanism that mitigates overconfidence in pseudo-labels. Experiments show that our framework outperforms self-supervised pretraining followed by supervised adaptation, and even matches the fully supervised counterpart when using only 20 labeled samples per class on the iNaturalist19 benchmark. Our code is available at <https://github.com/walline/semihoc>.*

## 1. Introduction

Out-of-distribution (OOD) detection is important for deploying deep learning in real-world settings. Most prior work studies the *binary* case, where a model is trained on a set of in-distribution (ID) classes, and the test-time task is to detect any sample not belonging to these classes [6, 14, 32].

More recently, researchers have explored the hierarchical setting of OOD detection [12, 24, 35], in which the ID classes are organized according to a known taxonomy or tree structure, and OOD samples that belong somewhere in this hierarchy (but not in the leaf ID classes) should be classified as the correct internal category. This setting, sometimes referred to as hierarchical novelty detection (HND), provides more informative OOD predictions than binary rejection. However, the problem is considerably more challenging, as it effectively becomes a multi-class classification problem over internal categories for which no labeled training data are available.

In parallel, works in domains such as open-set semi-

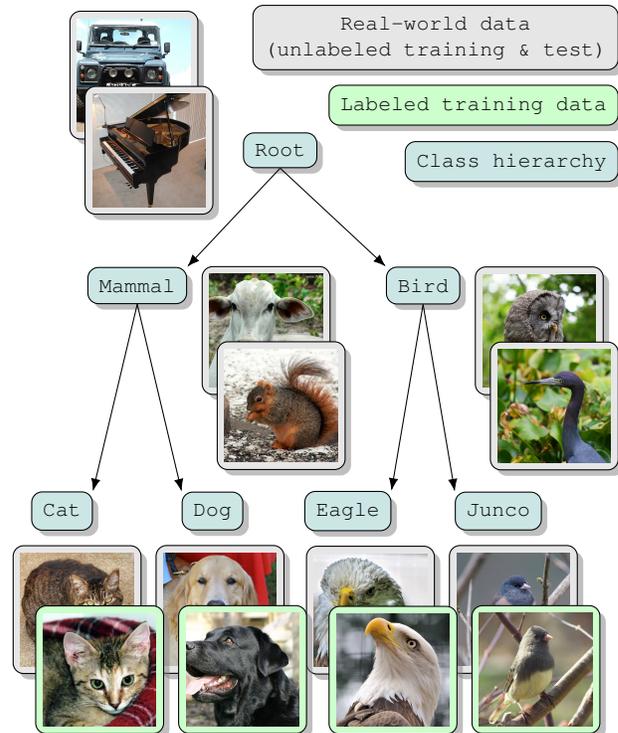


Figure 1. Semi-supervised hierarchical open-set classification. The goal is to predict data from the open-world distribution (gray) to the appropriate hierarchy nodes.

supervised learning [4, 34] and semantically coherent OOD detection [19, 36] have shown that OOD detection can benefit from exposing the model to OOD samples in the unlabeled training data. This raises the question: if real-world data inherently contain a mix of both ID and OOD samples, why not leverage these data to improve performance? Motivated by this, we introduce *semi-supervised hierarchical open-set classification*, a new problem setting that combines HND with the semi-supervised paradigm (see Fig. 1). We aim to explore both the opportunities and the challenges that arise in this setting.

From our investigations, we found that a key challenge is that standard confidence-based pseudo-labeling is unreliable when assigning OOD data to high-level categories, as OOD samples are difficult to distinguish from uncertain

ID samples. To address this, we introduce *subtree pseudo-labels* which, when assigned, indicate the likely membership within a specific subtree, without committing to a specific class, be it ID or OOD. This approach provides a more robust supervision signal for OOD samples.

Furthermore, we observe that OOD samples tend to receive increasingly deep pseudo-labels as training progresses, eventually overshooting the correct internal node and confusing it with some overly specific category. To counter these overpredictions, we propose *age-gating*: a mechanism that prevents pseudo-labels that are assigned unexpectedly late during training.

We integrate subtree pseudo-labels and age-gating into a teacher-student framework for semi-supervised hierarchical open-set classification, which we call *SemiHOC*. Experiments demonstrate that SemiHOC outperforms the baseline of self-supervised pretraining followed by supervised adaptation. Notably, on the iNaturalist19 benchmark, SemiHOC matches the performance of full ID supervision while using only 20 labels per class.

Our main contributions are as follows:

- We introduce semi-supervised hierarchical open-set classification, which extends HND to the semi-supervised setting, where the unlabeled data contain both ID and OOD samples.
- We identify two key challenges of this setting: 1) reliable pseudo-labeling of OOD data and 2) accumulation of overconfidence for OOD data, and address these with our *subtree pseudo-labels* and *age-gating*.
- We present SemiHOC, a teacher-student framework for semi-supervised open-set classification that integrates subtree pseudo-labels and age-gating, and show that it can match a fully supervised baseline with only a small number of labels per class.

## 2. Related work

### 2.1. OOD detection

OOD detection is a well-studied problem in deep learning. In its most common form, a model is trained on a set of in-distribution (ID) classes; during deployment, the goal is to detect inputs that do not belong to any of these classes [6, 14, 18, 37]. The standard paradigm addresses this in a binary fashion: a sample is either ID or OOD, disregarding semantic similarities between the OOD data and the ID classes. In contrast, the hierarchical settings of HND and this work allow for more informative predictions by assigning OOD samples to the most appropriate internal node of a given class hierarchy.

### 2.2. Class hierarchies for classification

Many classification problems involve classes that can be organized into a hierarchy, such as biological taxonomy or

product categorization. A common approach is to train models that optimize hierarchical metrics, such as tree distances or hierarchical precision [1, 5, 8–10, 15]. These methods account for the severity of classification errors within the hierarchy but assume that all test samples belong to known classes, and thus do not address OOD data.

### 2.3. OOD detection with class hierarchies

Several works have addressed the classification of OOD samples as internal nodes in a class hierarchy [12, 16, 23, 24, 27, 35]. This setting has been referred to as hierarchical novelty detection [12], hierarchical OOD detection [35], and fine-grained openset detection [24]. We use the term *hierarchical open-set classification*, which better fits the semi-supervised setting where the unknown classes appear in the unlabeled data (*i.e.*, they are not fully novel), and highlights that the goal is classification of unknown classes rather than binary rejection. There are two baseline approaches for this problem: top-down and flattening.

In the top-down approach [12, 16], local classifiers and OOD detectors are placed at each internal node in the hierarchy. At inference time, the model traverses from the root, predicting at each node whether the sample is locally OOD based on a thresholded score. The sample is assigned to the deepest node before exceeding the threshold. While intuitive, this approach has the drawback of requiring tuned thresholds for each OOD detector, which is impractical in real-world scenarios without access to labeled OOD data.

The flattening approach [12, 23, 27] instead trains a single flat classifier to directly predict among both ID leaf classes and internal OOD nodes. Since no OOD training data are available for the internal nodes, the usual workaround is to use data from descendant ID classes as a proxy. However, this may poorly represent the true distribution of OOD samples at internal nodes.

More recently, methods that more explicitly leverage the hierarchical structure have been proposed in ProHOC [35] and State-FGOD [24]. ProHOC utilizes the class hierarchy to factorize the distribution of predictions, such that each prediction becomes a product of conditional probabilities corresponding to the path from the root to the predicted node. The conditional distributions are modeled from ID classification networks trained at different levels of the hierarchy. State-FGOD proposes a state-based representation in which a network is trained to predict logits for each node in the hierarchy. The logits corresponding to valid states (paths in the class hierarchy) are multiplied and transformed to a predictive distribution using the softmax function.

A drawback of many existing methods for hierarchical novelty detection is their reliance on bias parameters to set the trade-off between ID and OOD performance. These biases are difficult to tune in practice without access to labeled OOD data. In this work, we therefore build on the ProHOC

methodology, which is designed to avoid such tuning.

## 2.4. Unlabeled ID/OOD exposure

OOD detection is motivated by the fact that real-world data often contains samples outside the target classes. Yet most traditional approaches assume access only to ID data during training, treating OOD detection purely as a test-time task. If, however, real-world data contain both ID and OOD, it is natural to exploit such unlabeled and uncurated data to improve both ID classification and OOD detection.

Recent research directions address this more realistic setting by leveraging an unlabeled mix of ID and OOD data during training. Examples include semantically coherent OOD (SC-OOD) detection [19, 36], open-set semi-supervised learning (OsSSL) [4, 13, 20, 34], and open-world semi-supervised learning (OwSSL) [2, 17, 25, 26].

SC-OOD detection highlights that many benchmarks use distinct datasets for ID and OOD, disregarding any semantic overlap between these datasets. For example, dogs from Tiny ImageNet are treated as OOD when the model is trained on CIFAR-10, which also contains dogs. SC-OOD trains on an unlabeled mix of both datasets, enabling the model to recognize overlapping categories as ID.

OsSSL assumes a small labeled ID set and a larger unlabeled mix of ID and OOD data, leveraging the unlabeled data to improve both ID classification and OOD detection. OwSSL has a similar data setup to OsSSL but differs in goal: OwSSL aims not only to detect OOD samples, but also to classify them into distinct classes. However, OwSSL typically relies on strong assumptions, such as balanced OOD class distributions or a known number of OOD classes, which limit its practicality.

Our work also utilizes an unlabeled ID/OOD mix but differs from existing approaches by:

- incorporating relationships between classes through a class hierarchy,
- performing OOD classification as opposed to binary detection, predicting OOD as internal categories of the hierarchy, and
- making no assumptions about the number or distribution of OOD classes, some categories may have abundant OOD data, others none.

## 3. Semi-supervised hierarchical open-set classification

### 3.1. Problem formulation

We consider a labeled training set with samples from domain  $\mathcal{X}$  with classes  $\mathcal{C}^{\text{id}}$ . These classes form the leaf nodes of a known directed rooted tree (hierarchy)  $\mathcal{H}$  with node set  $\mathcal{C}$  such that  $\mathcal{C}^{\text{id}} \subset \mathcal{C}$ . The labeled set is drawn from the ID distribution  $p^{\text{id}}(x, y)$  where  $x \in \mathcal{X}$  is the data and  $y \in \mathcal{C}^{\text{id}}$  is the label. The unlabeled training set and the test set

are drawn from an open-world distribution  $p^{\text{open}}(x, y)$  with  $y \in \mathcal{C}$  for which we only observe  $x \in \mathcal{X}$ . Thus,  $p^{\text{open}}$  covers both ID samples ( $y \in \mathcal{C}^{\text{id}}$ ) and OOD samples ( $y \in \mathcal{C} \setminus \mathcal{C}^{\text{id}}$ ), where OOD labels correspond to internal nodes of  $\mathcal{H}$ . Our objective is to learn a classifier,  $f : \mathcal{X} \rightarrow \mathcal{C}$ , that predicts the correct node in  $\mathcal{H}$ : a leaf if the sample belongs to an ID class, or the appropriate internal node if the sample is OOD.

### 3.2. ProHOC

This work builds on ProHOC [35], a method for HND trained with labeled data from  $p^{\text{id}}$  and evaluated on the open-set distribution  $p^{\text{open}}$ . ProHOC uses depth-specific classification networks, one per hierarchy depth, each trained with standard cross-entropy using ID data whose labels are remapped to their ancestors at the specified depth. The depth-specific predictions are fused into a hierarchical distribution by modeling conditional probabilities at each internal node, based on the confidences and entropies of the depth-specific predictions. This yields a probability distribution over the union of all possible OOD predictions (internal nodes) and ID classes (leaf nodes). In this work, we use ProHOC’s hierarchical predictions for pseudo-labeling, and we use these pseudo-labels to train the depth-specific classification models with standard cross-entropy losses.

### 3.3. Pseudo-labeling

A core technique for semi-supervised learning is pseudo-labeling [11]: using unlabeled samples with high-confidence predictions as training data. This requires a correlation between the prediction confidence and the accuracy, allowing us to assign pseudo-labels with high precision. This correlation is established for the ID case [30], but does not necessarily need to hold for OOD predictions in hierarchical open-set classification.

In the top panel of Fig. 2, we look at the accuracies as a function of confidence for predictions from a fully trained ProHOC model on iNaturalist19, where ID predictions (leaf nodes) and OOD predictions (internal nodes) are separated. For data predicted as an ID class, we see a clear correlation between prediction confidence and accuracy, indicating that these predictions can be used for pseudo-labeling. However, for data predicted as an internal node, there is no such correlation, making the naive application of confidence-based pseudo-labels unfit for these predictions.

However, as shown in Fig. 3, when analyzing mistakes for samples predicted as OOD, specifically by looking at the distances from the predictions to the ground truth, decomposed into overprediction and underprediction, we see that 59.2% of samples have an overprediction distance of zero, meaning these are predicted as the correct node or an ancestor of the correct node. The most common being predicting the parent of the ground truth node, which occurs in 39.9% of OOD predictions.

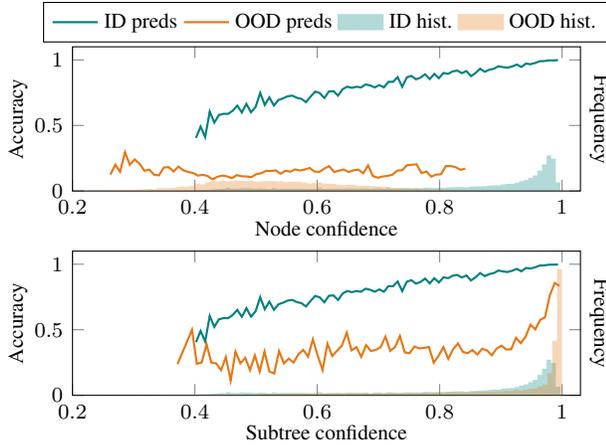


Figure 2. Average accuracy as a function of prediction confidence for samples predicted as ID and OOD. The shaded regions show the corresponding confidence histograms. The subtree confidence enables pseudo-labeling for OOD-predictions by achieving high accuracy for high-confidence predictions.

Motivated by these findings, we propose a new type of pseudo-labels for hierarchical open-set classification: *subtree pseudo-labels* (SPL). A subtree pseudo-label for a node in the hierarchy says that the sample likely belongs to the subtree starting from this node, *i.e.*, the sample is either OOD at this node or one of the descendant nodes, or belongs to one of the descendant leaves. Formally, we define the subtree as

$$\text{Subtree}(c) = \{c' \in \mathcal{C} \mid c' \text{ is a descendant of } c \text{ or is } c\}, \quad (1)$$

where  $\mathcal{C}$  is the set of nodes in the class hierarchy  $\mathcal{H}$ .

The confidence used for SPLs is calculated by summing the predicted probabilities,  $p(c|x)$ , belonging to the subtree. As shown in the bottom panel of Fig. 2, subtree confidence enables high accuracy at high confidence for OOD (non-leaf) predictions, allowing us to assign these less specific but more robust SPLs with high precision.

Samples get the subtree pseudo-label if the confidence exceeds a threshold,  $\tau$ . The subtree pseudo-labels can be formulated as

$$\hat{y}_c = \left[ \sum_{c' \in \text{Subtree}(c)} p(c'|x) \right] > \tau, \quad (2)$$

such that  $\hat{y}_u = 1$  if the condition holds and zero otherwise. We denote the set of subtree pseudo-labels for a sample  $\hat{\mathbf{y}}$ .

### 3.4. Overpredictions of OOD data

It is known from prior work on OsSSL [4] that applying standard confidence-based pseudo-labeling to unlabeled data containing unseen classes eventually causes the unknown samples to be assigned to their most similar ID classes. For hierarchical open-set classification, this risk

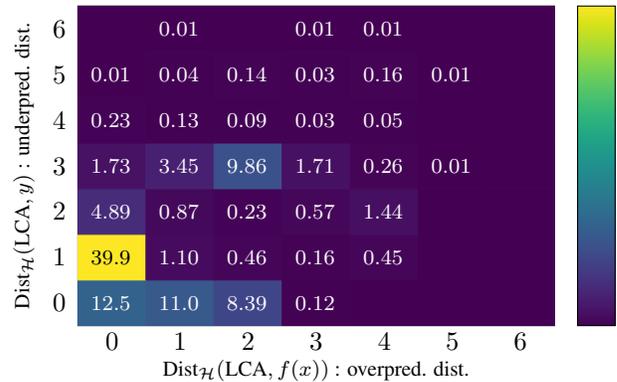


Figure 3. Distances (the number of edges in  $\mathcal{H}$ ) to the ground truth for data predicted as OOD (internal nodes), displayed as percentages. LCA denotes the Lowest Common Ancestor of the prediction and the ground-truth.

is even greater, as the unknown classes can appear deep in the hierarchy and thus be semantically and visually close to the ID classes. Moreover, as observed in Sec. 3.3, standard confidence-based pseudo-labeling struggles to reliably separate ID from OOD in the hierarchical setting.

Our experiments confirm this problem. Figure 7 of Sec. 3.4 shows the purity and average depth of subtree pseudo-labels assigned to OOD samples when training a standard SSL framework with subtree pseudo-labels. Here, purity is defined as the fraction of samples whose ground-truth label lies within the assigned subtree. We see that the model gradually becomes overconfident, assigning OOD samples to increasingly deep nodes in the hierarchy.

To further analyze this overconfidence, we study the dynamics of SPL assignments over training epochs in search of informative patterns. For this purpose, we log the first assignment per sample per node and visualize the resulting frequencies, decomposed into correct and incorrect assignments. The dynamics in Fig. 4 reveal that correct assignments occur early, while incorrect, overly deep assignments emerge later as distinct peaks. The overpredictions are largely driven by visually similar OOD samples that belong higher in the hierarchy. These results are from iNat19 with 10 labels per class.

This observation motivates our proposed *age-gating* strategy for SPL. The key idea is to prevent additional SPLs from being assigned to a node once the initial wave of correct assignments has passed. We implement this with a simple cutoff-detection method, detailed in Algorithm 1. The cutoff detection scans for the first peak in the SPL assignment frequency of each node and defines a cutoff when the frequency falls below a fixed fraction of the peak value. Cutoff detection is performed per node at the end of each epoch. When a cutoff is set, no new SPLs are assigned to that node. Additional results on the effectiveness of age-gating are presented in Sec. 4.5.

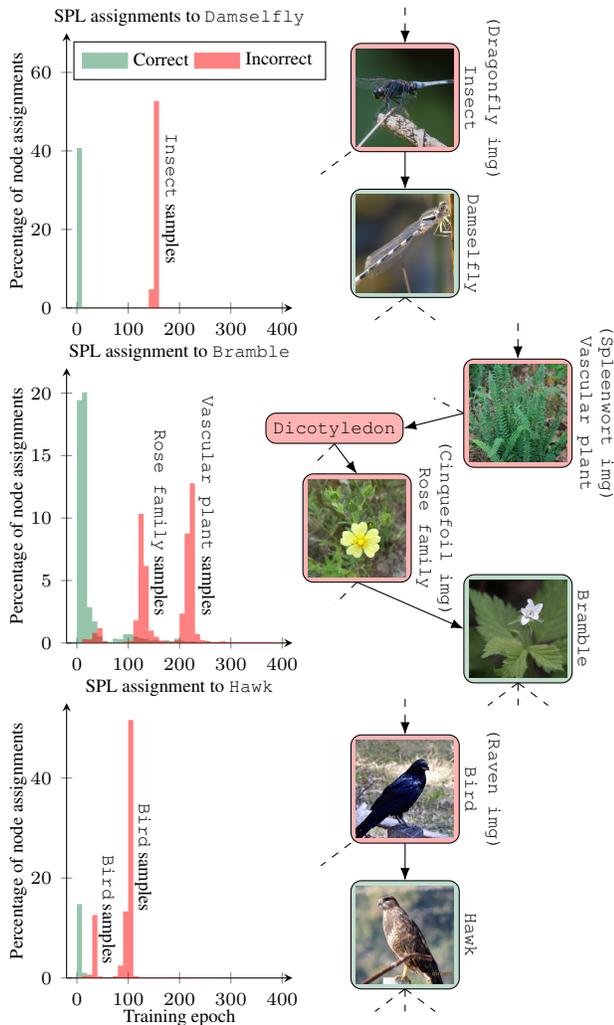


Figure 4. Frequencies of subtree pseudo-label assignments across training epochs for selected hierarchy nodes, split into correct and incorrect assignments. Each node’s local hierarchy structure and example images from correct and incorrect assignments are shown to the right. We see that the majority of correct SPLs are assigned early in training, while incorrect overpredictions of OOD data emerge later as distinct peaks.

### 3.5. Our framework: SemiHOC

Equipped with our proposed subtree pseudo-labels and age-gating strategy, we introduce *SemiHOC*, a framework for semi-supervised open-set classification. At its core, SemiHOC follows a teacher-student architecture, extended with our components specifically designed for the hierarchical open-set setting. We adopt the ProHOC method to produce hierarchical predictions, as it is the only existing HND method that does not require tuning of an ID/OOD trade-off parameter, a parameter that is difficult to set without access to labeled OOD data.

For the labeled data, we follow the same approach as

#### Algorithm 1: Detect cutoffs for age-gating

---

**Input:** List of pseudo-label assignment epochs for one node epochs; bin width  $w$ ; drop threshold  $\gamma$   
**Output:** Step value where a significant drop is detected, or  $\infty$

```

1 counts, edges  $\leftarrow$  BinCounts (epochs, bin width  $w$ )
2 max_count  $\leftarrow$  0
3 for  $i \leftarrow 0$  to length(counts) - 1 do
4   if counts[ $i$ ] > max_count then
5     max_count  $\leftarrow$  counts[ $i$ ]
6   else if counts[ $i$ ] <  $\gamma \cdot$  max_count then
7     return edges[ $i$ ]
8 return  $\infty$ 

```

---

supervised training for ProHOC: labels are remapped to every height of the hierarchy. These remapped labels are then used with a standard cross-entropy loss to train the corresponding depth-specific networks.

For the unlabeled data, the teacher model generates predictions using ProHOC. These predictions are reduced to subtree pseudo-labels following (2). To support age-gating, we maintain an SPL log as a mapping  $(c, g) \rightarrow e$ , where  $c$  is a class,  $g$  is a unique sample identifier, and  $e$  is the epoch of assignment. New sample-class pairs are added on the first assignment and removed if not reassigned in subsequent epochs (*i.e.*, if confidence decreases or the prediction switches branch).

We then remove any subtree pseudo-labels assigned *after* the cutoff epoch, as determined by the age-gating procedure. Note that a sample is not necessarily excluded entirely from training. Rather, one or more of its deepest subtree pseudo-labels are blocked and ignored in loss computation.

The remaining subtree pseudo-labels are used to train the student model. If a sample’s subtree pseudo-labels extend to depth  $d$  in the hierarchy, the corresponding top  $d$  networks contribute to its loss via standard cross-entropy, while deeper networks are unaffected.

At the end of each epoch (*i.e.*, a full iteration through the unlabeled data), we apply the cutoff-detection method in Algorithm 1 to all hierarchy nodes to determine new cutoff epochs. The flow of data in SemiHOC is visualized in Fig. 5 and a detailed training algorithm is provided in the supplementary material.

We use DinoV2 features [22], a widely used image backbone, as input to our depth-specific models. DinoV2 is trained fully self-supervised, ensuring that our approach does not exploit any label information beyond the small labeled training set. The teacher model is obtained by taking an exponential moving average of the student model parameters. We also apply dropout [31] to the input features and the intermediate features of the student model, analogous to how strong data augmentations are applied to student inputs in, *e.g.*, the seminal SSL method FixMatch [30].

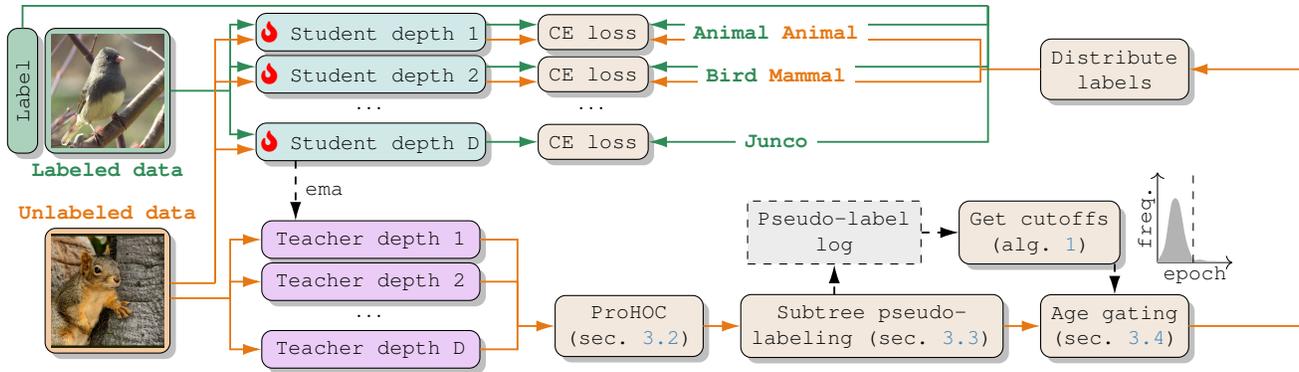


Figure 5. The flow of labeled and unlabeled data in SemiHOC. Labeled data are used to train depth-specific student models with cross-entropy losses using labels mapped to the specified depths. Unlabeled data are predicted by teacher models via ProHOC, producing hierarchical distributions that form subtree pseudo-labels. After age-gating, these subtree pseudo-labels are mapped to the appropriate depths and used to train the students with cross-entropy losses.

## 4. Experiments and results

### 4.1. Datasets

We evaluate our method on three large-scale image classification benchmarks with class hierarchies: iNaturalist19 [33], iNaturalist21-Aves [7], and SimpleHierImageNet [35], the latter constructed from ImageNet [28] and the WordNet graph [21]. The full iNaturalist21 dataset contains 10K classes and 2.7M images. To make it more accessible for experimentation, we use the subtree starting from *Class: Aves*, *i.e.*, all bird classes. This bird subset contains 1.5K classes and 500K images. For iNaturalist19 and SimpleHierImageNet, we adopt the ID/OOD splits defined in [35], while for iNaturalist21-Aves, we construct a random ID/OOD split (available in the linked code). We create labeled subsets by sampling up to  $n$  images per ID class, using all available samples if fewer than  $n$  exist. In our experiments, we use 10, 20, and 50 labels per class.

For iNaturalist19 and iNaturalist21-Aves, we append both the unlabeled training set and the test set with 300 randomly selected images from Office31 [29], providing out-of-hierarchy OOD data, *i.e.*, OOD samples belonging at the root of the hierarchy. For SimpleHierImageNet, we do not introduce root-level OOD samples, since its top-level nodes correspond to abstract and broad concepts such as *physical object* and *matter*. More dataset details are available in the supplementary material.

### 4.2. Training details

Training of neural networks involves many configurable settings; for most of these, we follow standard configurations from the literature. Our depth-specific classification heads for the ProHOC predictions consist of four layers with a hidden dimension of 512. The backbone is a DinoV2 ViT-L/14 with registers [3]. All models are trained for 400 epochs using stochastic gradient descent with momentum 0.9 and weight decay of 0.001. The labeled batch

size is 128, and the unlabeled batch size is set to four times this value. The teacher model is updated as an exponential moving average of the student parameters with momentum 0.999. For pseudo-labeling, we use the threshold  $\tau = 0.95$ .

The hyperparameters we find show sensitivity are the learning rate and dropout. In our experiments, we use a dropout rate of 0.3 throughout. We use a learning rate of 0.01 on iNat19 and iNat21-Aves, and 0.1 for SimpleHierImageNet. A more detailed discussion on the choice of learning rate and dropout is provided in Sec. 4.4.

### 4.3. Main results

To the best of our knowledge, no prior methods have been developed for semi-supervised hierarchical open-set classification. To show the effectiveness of SemiHOC, we compare against several intuitive baselines: 1) **Supervised only**: a standard pipeline of self-supervised pretraining followed by supervised training of classification heads; 2) **SSL (Node PLs)**: semi-supervised training from pseudo-labels assigned by thresholding ProHOC predictions per node; and 3) **SSL (per depth)**: training the depth-specific heads independently using FixMatch-style pseudo-labels per depth. For comparison, we also include **Sup. (all ID labels)**, which is equal to Supervised only but uses all ID labels, and **SPL oracle**, which follows SemiHOC without age-gating but with access to correct subtree pseudo-labels for all unlabeled data. All methods are trained on top of frozen DinoV2 features, which have proven to be strong feature representations for our evaluated benchmarks [22]. All compared methods use equal hyperparameters within each evaluated dataset. Detailed descriptions of the included baselines are provided in the supplementary material.

Results are shown in Tab. 1, evaluated using the balanced mean hierarchical distance (BMHD) [35] measured between predictions and ground-truths. Our primary performance metric is BMHD Mix, which averages over ID and OOD performance. We also report the ID and OOD compo-

Table 1. Balanced mean hierarchical distance (BMHD) on ID, OOD, and mixed ID+OOD test sets for SemiHOC and baseline methods. The best results under limited labels are **boldfaced**.

BMHD ↓	10 labels / class			20 labels / class			50 labels / class		
	ID	OOD	Mix	ID	OOD	Mix	ID	OOD	Mix
iNATURALIST19									
Sup. (all ID labels)	0.72	0.71	<b>0.71</b>	0.72	0.71	0.71	0.72	0.71	0.71
SSL (SPL oracle)	0.58	0.37	0.48	0.58	0.37	0.48	0.58	0.37	0.48
Supervised only	0.82	0.84	0.83	0.71	0.81	0.76	0.71	0.73	0.72
SSL (Node PLs)	0.81	0.94	0.87	0.71	0.88	0.79	0.64	0.81	0.72
SSL (per depth)	0.86	1.11	0.99	0.70	1.00	0.85	0.62	0.86	0.74
SemiHOC	0.81	0.73	<b>0.77</b>	0.70	0.69	<b>0.69</b>	0.64	0.62	<b>0.63</b>
iNATURALIST21-AVES									
Sup. (all ID labels)	0.77	0.73	<b>0.75</b>	0.77	0.73	0.75	0.77	0.73	0.75
SSL (SPL oracle)	0.66	0.53	0.59	0.66	0.53	0.59	0.66	0.53	0.59
Supervised only	0.89	1.00	0.94	0.79	0.95	0.87	0.74	0.85	0.79
SSL (Node PLs)	0.91	0.90	0.91	0.80	0.90	0.85	0.76	0.80	0.78
SSL (per depth)	0.93	1.28	1.10	0.82	1.19	1.01	0.73	1.02	0.88
SemiHOC	0.90	0.82	<b>0.86</b>	0.80	0.87	<b>0.84</b>	0.75	0.76	<b>0.76</b>
SIMPLEHIERIMAGENET									
Sup. (all ID labels)	0.75	1.36	1.05	0.75	1.36	1.05	0.75	1.36	1.05
SSL (SPL oracle)	0.57	0.59	0.58	0.57	0.59	0.58	0.57	0.59	0.58
Supervised only	0.82	1.88	1.35	0.72	1.73	1.23	0.69	1.65	1.17
SSL (Node PLs)	0.82	1.80	<b>1.31</b>	0.72	1.66	1.19	0.68	1.58	1.13
SSL (per depth)	0.93	2.36	1.65	0.86	2.07	1.47	0.67	1.70	1.19
SemiHOC	0.83	1.79	<b>1.31</b>	0.72	1.64	<b>1.18</b>	0.68	1.55	<b>1.12</b>

nents separately for completeness. Predictions are obtained from the teacher models via ProHOC.

SemiHOC consistently achieves the best performance among methods using limited labels. On iNat19, it even matches the model with full ID supervision when trained with only 20 labels per class, and surpasses it when using 50 labels per class. Notably, the SPL oracle outperforms supervised only with full ID supervision on ID test data, indicating that accurate subtree pseudo-labels for OOD samples provide a valuable learning signal that also benefits ID performance.

#### 4.4. Learning rate and dropout

Selecting hyperparameters in settings that involve OOD predictions is not straightforward since we cannot assume access to a test set containing OOD data. Moreover, HND inherently requires balancing performance between ID and OOD data. For example, maximizing ID performance can be achieved trivially by always predicting leaf classes, but this leads to poor OOD predictions.

For SemiHOC, we find that the key hyperparameters for controlling this trade-off are the learning rate and the dropout rate of the student model. Increasing either has a regularizing effect, making the model less confident: this benefits OOD performance but reduces ID performance by decreasing the number of samples predicted as leaf classes.

Figure 6 illustrates this trade-off when varying the learn-

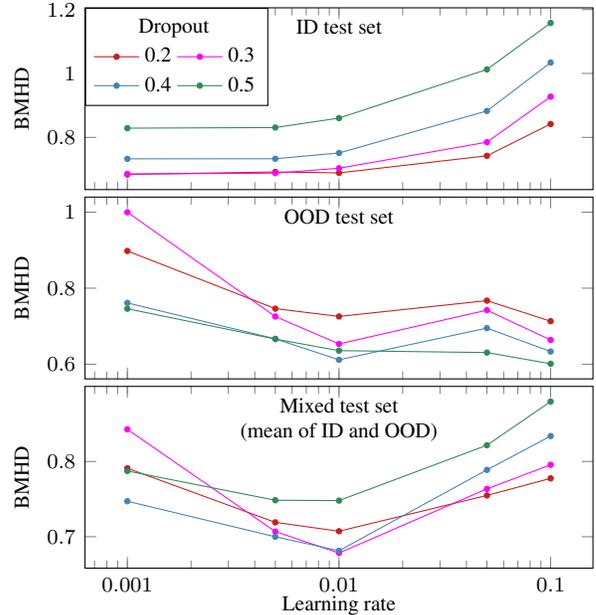


Figure 6. Increasing the learning or dropout rate has a regularizing effect and improves performance on OOD while decreasing that for ID. The ID and OOD curves share a shoulder point (at around dropout 0.3 and learning rate 0.01), which provides a good trade-off for performance on the mixed ID+OOD setting.

ing rate and dropout. Although OOD performance cannot be observed in practice, the curves are roughly symmetric around a region where both ID and OOD performance are strong. This symmetry provides a practical guideline: choose the largest learning rate and dropout value that do not significantly degrade ID performance. For this scenario, iNaturalist with 20 labels per class, a suitable choice is a learning rate of 0.01 and dropout of 0.3. The corresponding results for iNaturalist21-Aves and SimpleHierImagenet are available in the supplementary material.

#### 4.5. Age-gating

In Sec. 3.4, we introduced our age-gating strategy for preventing overconfident predictions on OOD data. The key idea is to block samples that are assigned subtree pseudo-labels unexpectedly late. As shown in Fig. 7, age-gating has a clear effect on the purity of subtree pseudo-labels for OOD data. With age-gating, the purity is maintained throughout training, whereas without it, the model gradually becomes overconfident, leading to a decline in purity.

To further analyze this strategy, we can view age-gating as a binary classifier at each node: correct subtree pseudo-label assignments are treated as positives, and incorrect assignments are negatives. With this perspective, our main objective is to achieve a low false positive rate (FPR), *i.e.*, to minimize the number of incorrect assignments that pass through the gate. However, trivially rejecting all samples would yield zero FPR. We therefore evaluate FPR together

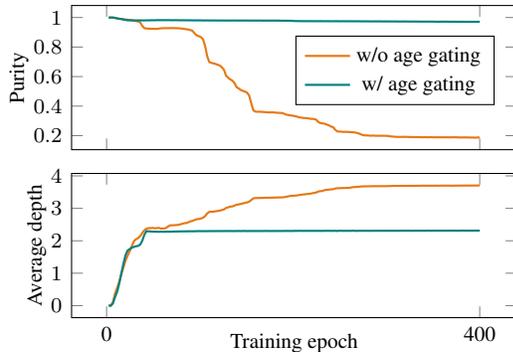


Figure 7. Purity and average depth of SPLs assigned to OOD data during training with and without age-gating. Age-gating prevents overconfidence and maintains purity for OOD data throughout training. Results are from iNat19 (10 labels per class).

with the *coverage*, defined as the fraction of assignments that are not gated, ensuring that a substantial amount of subtree pseudo-labels remains available for training.

The results of this analysis are summarized in Fig. 8, where we report FPR and coverage across all assignments for all hierarchy nodes as micro-averages. The results are obtained from a run on iNat19 with 10 labels per class. We vary the hyperparameters of the cutoff-detection in Algorithm 1 to study their effect. Increasing the bin width or decreasing the drop threshold allows more samples to pass, pushing both FPR and coverage towards one. Conversely, smaller bin widths and higher drop thresholds result in stricter gating, reducing FPR and coverage. We observe a trade-off region around 0.8 coverage and 0.1 FPR; beyond this point, FPR increases rapidly. Importantly, we note that results in this region are obtained for many parameter pairs, indicating that our age-gating procedure is robust to hyperparameter choices. Unless otherwise stated, we use  $w = 1$  and  $\gamma = 0.01$ . The supplementary material contains additional results on age-gating, including analyses of error types and results from other datasets.

#### 4.6. Ablation

To better understand the contributions of subtree pseudo-labels and age-gating, we analyze the performance over training epochs on iNat19 (20 labels per class). We compare SemiHOC with and without age-gating, as well as against baselines. The results in Fig. 9 show that subtree pseudo-labels provide a strong learning signal, enabling SemiHOC to quickly surpass the baselines. However, without age-gating, performance drops as training progresses due to overconfident predictions on OOD data.

### 5. Limitations

This work assumes that classes can be organized into a meaningful hierarchy, which does not hold for all classification problems. We further rely on the hierarchy reflect-

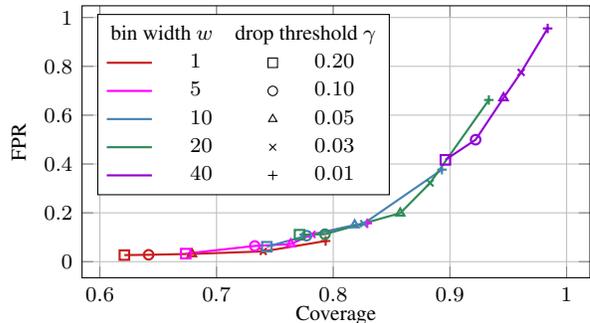


Figure 8. Analyzing our age-gating performance under varying hyperparameters. The goal is to block incorrect assignments (low FPR) while retaining sufficient data (high coverage). We obtain results in a favourable trade-off region around 0.8 coverage and 0.1 FPR for many parameter values.

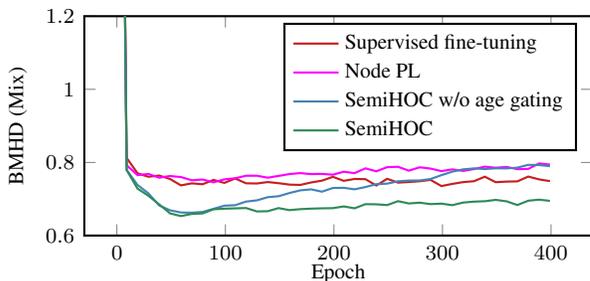


Figure 9. Performance on iNat19 with 20 labels per class, comparing SemiHOC to baselines with and without age-gating. Subtree pseudo-labels enable SemiHOC to learn efficiently and surpass the baselines, but without age-gating, performance drops over training due to overconfident OOD predictions.

ing visually observable concepts; otherwise, OOD samples cannot be expected to be correctly predicted. Finally, our evaluations use the same OOD classes in unlabeled training and test data. An interesting extension for the future is to consider test sets with previously unseen OOD classes.

### 6. Conclusion

We introduced semi-supervised hierarchical open-set classification, enabling the use of large unlabeled datasets in HND. Two key challenges were identified: 1) generating reliable pseudo-labels for OOD samples, and 2) preventing overconfident, overly specific predictions of OOD data. Our SemiHOC framework addresses these challenges and achieves strong results across the evaluated benchmarks. We believe there is still room for further work. For example, finding methods for accurately separating OOD data from uncertain ID data would be very valuable in this semi-supervised setting. Moreover, alternative strategies to our age-gating strategy can be explored to address the overconfidence issue. We hope our work can inspire the community to explore new ideas for hierarchical open-set classification.

## Acknowledgement

This work was supported by Saab AB and Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

## References

- [1] Luca Bertinetto, Romain Mueller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A Lord. Making better mistakes: Leveraging class hierarchies with deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [2] Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. In *International Conference on Learning Representations*, 2022. 3
- [3] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *International Conference on Learning Representations*, 2024. 6
- [4] Yue Fan, Anna Kukleva, Dengxin Dai, and Bernt Schiele. SSB: Simple but strong baseline for boosting performance of open-set semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 1, 3, 4
- [5] Ashima Garg, Depanshu Sani, and Saket Anand. Learning hierarchy aware features for reducing mistake severity. In *Proceedings of the European Conference on Computer Vision*, 2022. 2
- [6] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017. 1, 2
- [7] iNaturalist 2021 competition dataset. iNaturalist 2021 competition dataset. [https://github.com/visipedia/inat\\_comp/tree/master/2021](https://github.com/visipedia/inat_comp/tree/master/2021), 2021. 6
- [8] Kanishk Jain, Shyamgopal Karthik, and Vineet Gandhi. Test-time amendment with a coarse classifier for fine-grained classification. In *Advances in Neural Information Processing Systems*, 2024. 2
- [9] Shyamgopal Karthik, Ameya Prabhu, Puneet K. Dokania, and Vineet Gandhi. No cost likelihood manipulation at test time for making better mistakes in deep networks. In *International Conference on Learning Representations*, 2021.
- [10] Aris Kosmopoulos, Ioannis Partalas, Eric Gaussier, Georgios Paliouras, and Ion Androutsopoulos. Evaluation measures for hierarchical classification: a unified view and novel approaches. *Data Mining and Knowledge Discovery*, 29(3): 820–865, 2015. 2
- [11] Dong-Hyun Lee. Pseudo-Label: The simple and efficient semi-supervised learning method for deep neural networks. In *Proceedings of the ICML Workshop on Challenges in Representation Learning*, 2013. 3
- [12] Kibok Lee, Kimin Lee, Kyle Min, Yuting Zhang, Jinwoo Shin, and Honglak Lee. Hierarchical novelty detection for visual object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [13] Zekun Li, Lei Qi, Yinghuan Shi, and Yang Gao. IOMatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3
- [14] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *International Conference on Learning Representations*, 2018. 1, 2
- [15] Tong Liang and Jim Davis. Inducing neural collapse to a fixed hierarchy-aware frame for reducing mistake severity. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2
- [16] Randolph Linderman, Jingyang Zhang, Nathan Inkawich, Hai Li, and Yiran Chen. Fine-grain inference on out-of-distribution data with hierarchical classification. In *Proceedings of the Conference on Lifelong Learning Agents*, 2023. 2
- [17] Jiaming Liu, Yangqiming Wang, Tongze Zhang, Yulu Fan, Qinli Yang, and Junming Shao. Open-world semi-supervised novel class discovery. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2023. 3
- [18] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *Advances in Neural Information Processing Systems*, 2020. 2
- [19] Fan Lu, Kai Zhu, Wei Zhai, Kecheng Zheng, and Yang Cao. Uncertainty-aware optimal transport for semantically coherent out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1, 3
- [20] Qiankun Ma, Jiyao Gao, Bo Zhan, Yunpeng Guo, Jiliu Zhou, and Yan Wang. Rethinking safe semi-supervised learning: Transferring the open-set problem to a close-set one. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3
- [21] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 6
- [22] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. 5, 6
- [23] Spandan Pyakurel and Qi Yu. Hierarchical novelty detection via fine-grained evidence allocation. In *Proceedings of the International Conference on Machine Learning*, 2024. 2
- [24] Spandan Pyakurel and Qi Yu. Learning state-based node representations from a class hierarchy for fine-grained open-set detection. In *Proceedings of the International Conference on Machine Learning*, 2025. 1, 2

- [25] Mamshad Nayeem Rizve, Navid Kardan, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Openldn: Learning to discover novel classes for open-world semi-supervised learning. In *Proceedings of the European Conference on Computer Vision*, 2022. [3](#)
- [26] Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. Towards realistic semi-supervised learning. In *Proceedings of the European Conference on Computer Vision*, 2022. [3](#)
- [27] Idoia Ruiz and Joan Serrat. Hierarchical novelty detection for traffic sign recognition. *Sensors*, 22(12):4389, 2022. [2](#)
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. [6](#)
- [29] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proceedings of the European Conference on Computer Vision*, 2010. [6](#)
- [30] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. FixMatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems*, 2020. [3](#), [5](#)
- [31] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014. [5](#)
- [32] Yiyu Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *Proceedings of the International Conference on Machine Learning*, 2022. [1](#)
- [33] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. [6](#)
- [34] Erik Wallin, Lennart Svensson, Fredrik Kahl, and Lars Hammarstrand. Prosub: Probabilistic open-set semi-supervised learning with subspace-based out-of-distribution detection. In *Proceedings of the European Conference on Computer Vision*, 2024. [1](#), [3](#)
- [35] Erik Wallin, Fredrik Kahl, and Lars Hammarstrand. Prohoc: Probabilistic hierarchical out-of-distribution classification via multi-depth networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. [1](#), [2](#), [3](#), [6](#)
- [36] Jingkan Yang, Haoqi Wang, Litong Feng, Xiaopeng Yan, Huabin Zheng, Wayne Zhang, and Ziwei Liu. Semantically coherent out-of-distribution detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. [1](#), [3](#)
- [37] Jingkan Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *International Journal of Computer Vision*, 132(12):5635–5662, 2024. [2](#)