

Bootstrapping Fine-Grained Classifiers: Active Learning with a Crowd in the Loop

Human-in-the-loop Discriminative Patch Discovery: New Algorithm

Input: Dataset \mathcal{D} of image patches, set of negative images from the wild \mathcal{N}
Output: Classifiers C for attributes A

```


1  $A \leftarrow$  attributes ▷ acquired through consultation with experts
2 for  $A_i \in A$  do
3    $S_{ij} \leftarrow$  seed exemplars of  $A_i$  in pose  $j$ 
4 end
5 for  $A_i \in A$  do
6   for  $S_{ij} \in S_i$  do
7      $C_{ij} = \text{svmTrain}(S_{ij}, \mathcal{N})$ 
8      $N_{ij} = \emptyset \quad \forall i, j$  ▷ set of hard negatives is initially empty
9     repeat
10       $\mathcal{D} = \text{orderDetections}(C_{ij}, \mathcal{D} - N_{ij})$  ▷ crowdsourced method
11       $N_{ij} = N_{ij} \cup \text{hasNegatives}(\mathcal{D})$ 
12       $C_{ij} = \text{svmTrain}(S_{ij}, N_{ij} \cup \mathcal{N})$ 
13    until convergence()
14  end
15 end
16 return  $C, A$ 
    
```

Mechanical Turk User Interface:

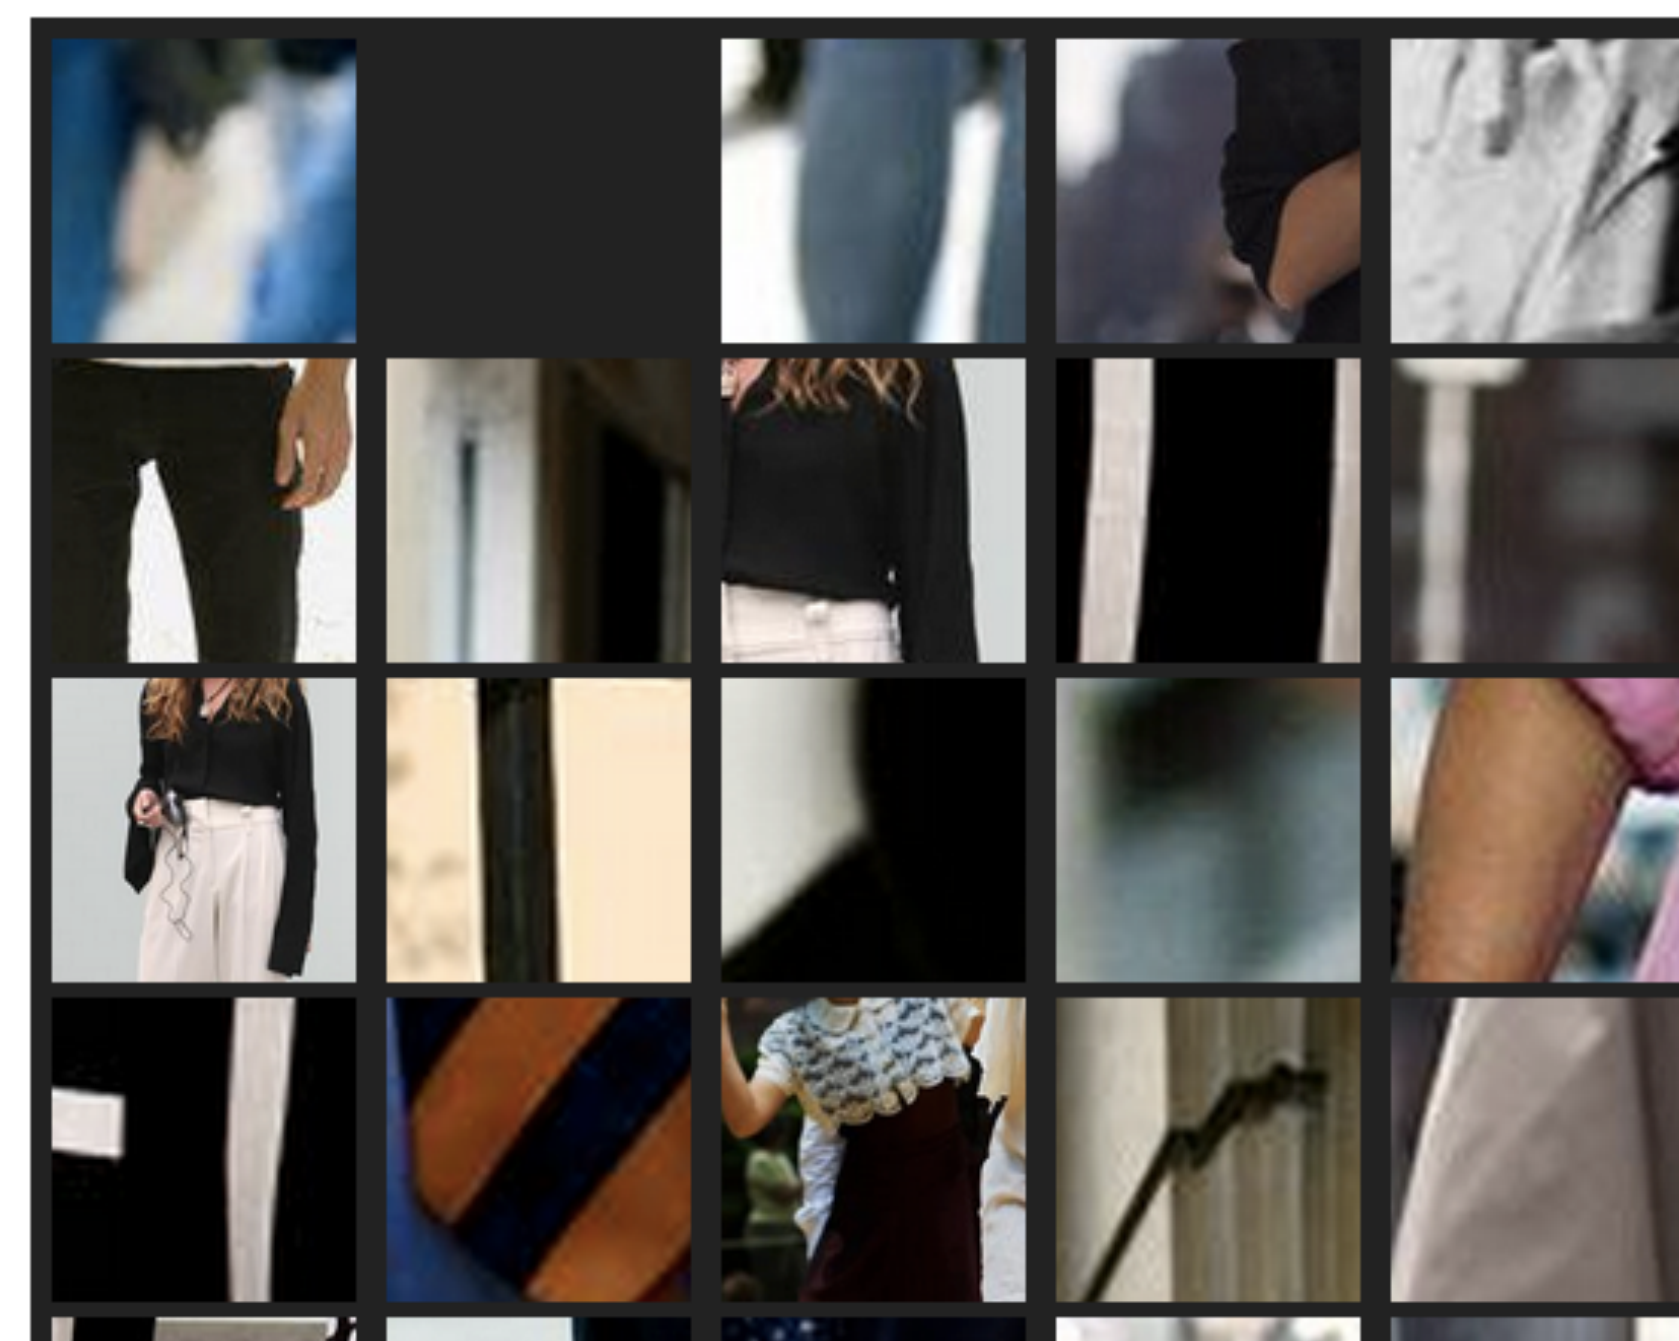
Click on all the image patches that contain: **a blazer viewed from the front**

Click submit when you are finished.
 Drag or click images to select or deselect.

Example Positive Images



Candidate Images



Yes Images



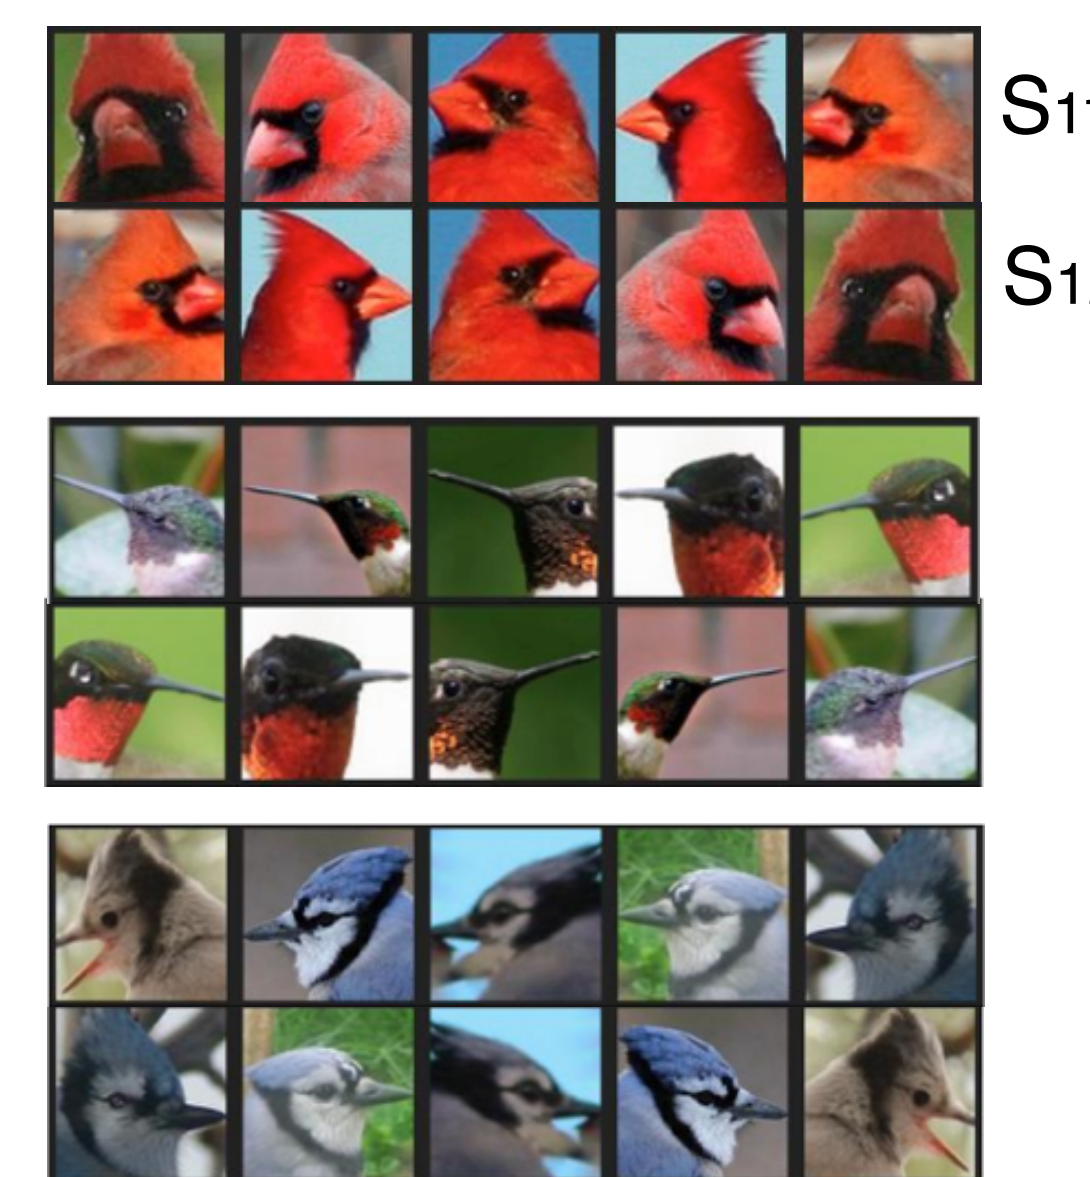
Experts begin by making a list of parts and/or attributes for their taxon and selecting exemplars for the canonical poses of each attribute or part.

This set depends on the Taxon

Part/ Attribute List

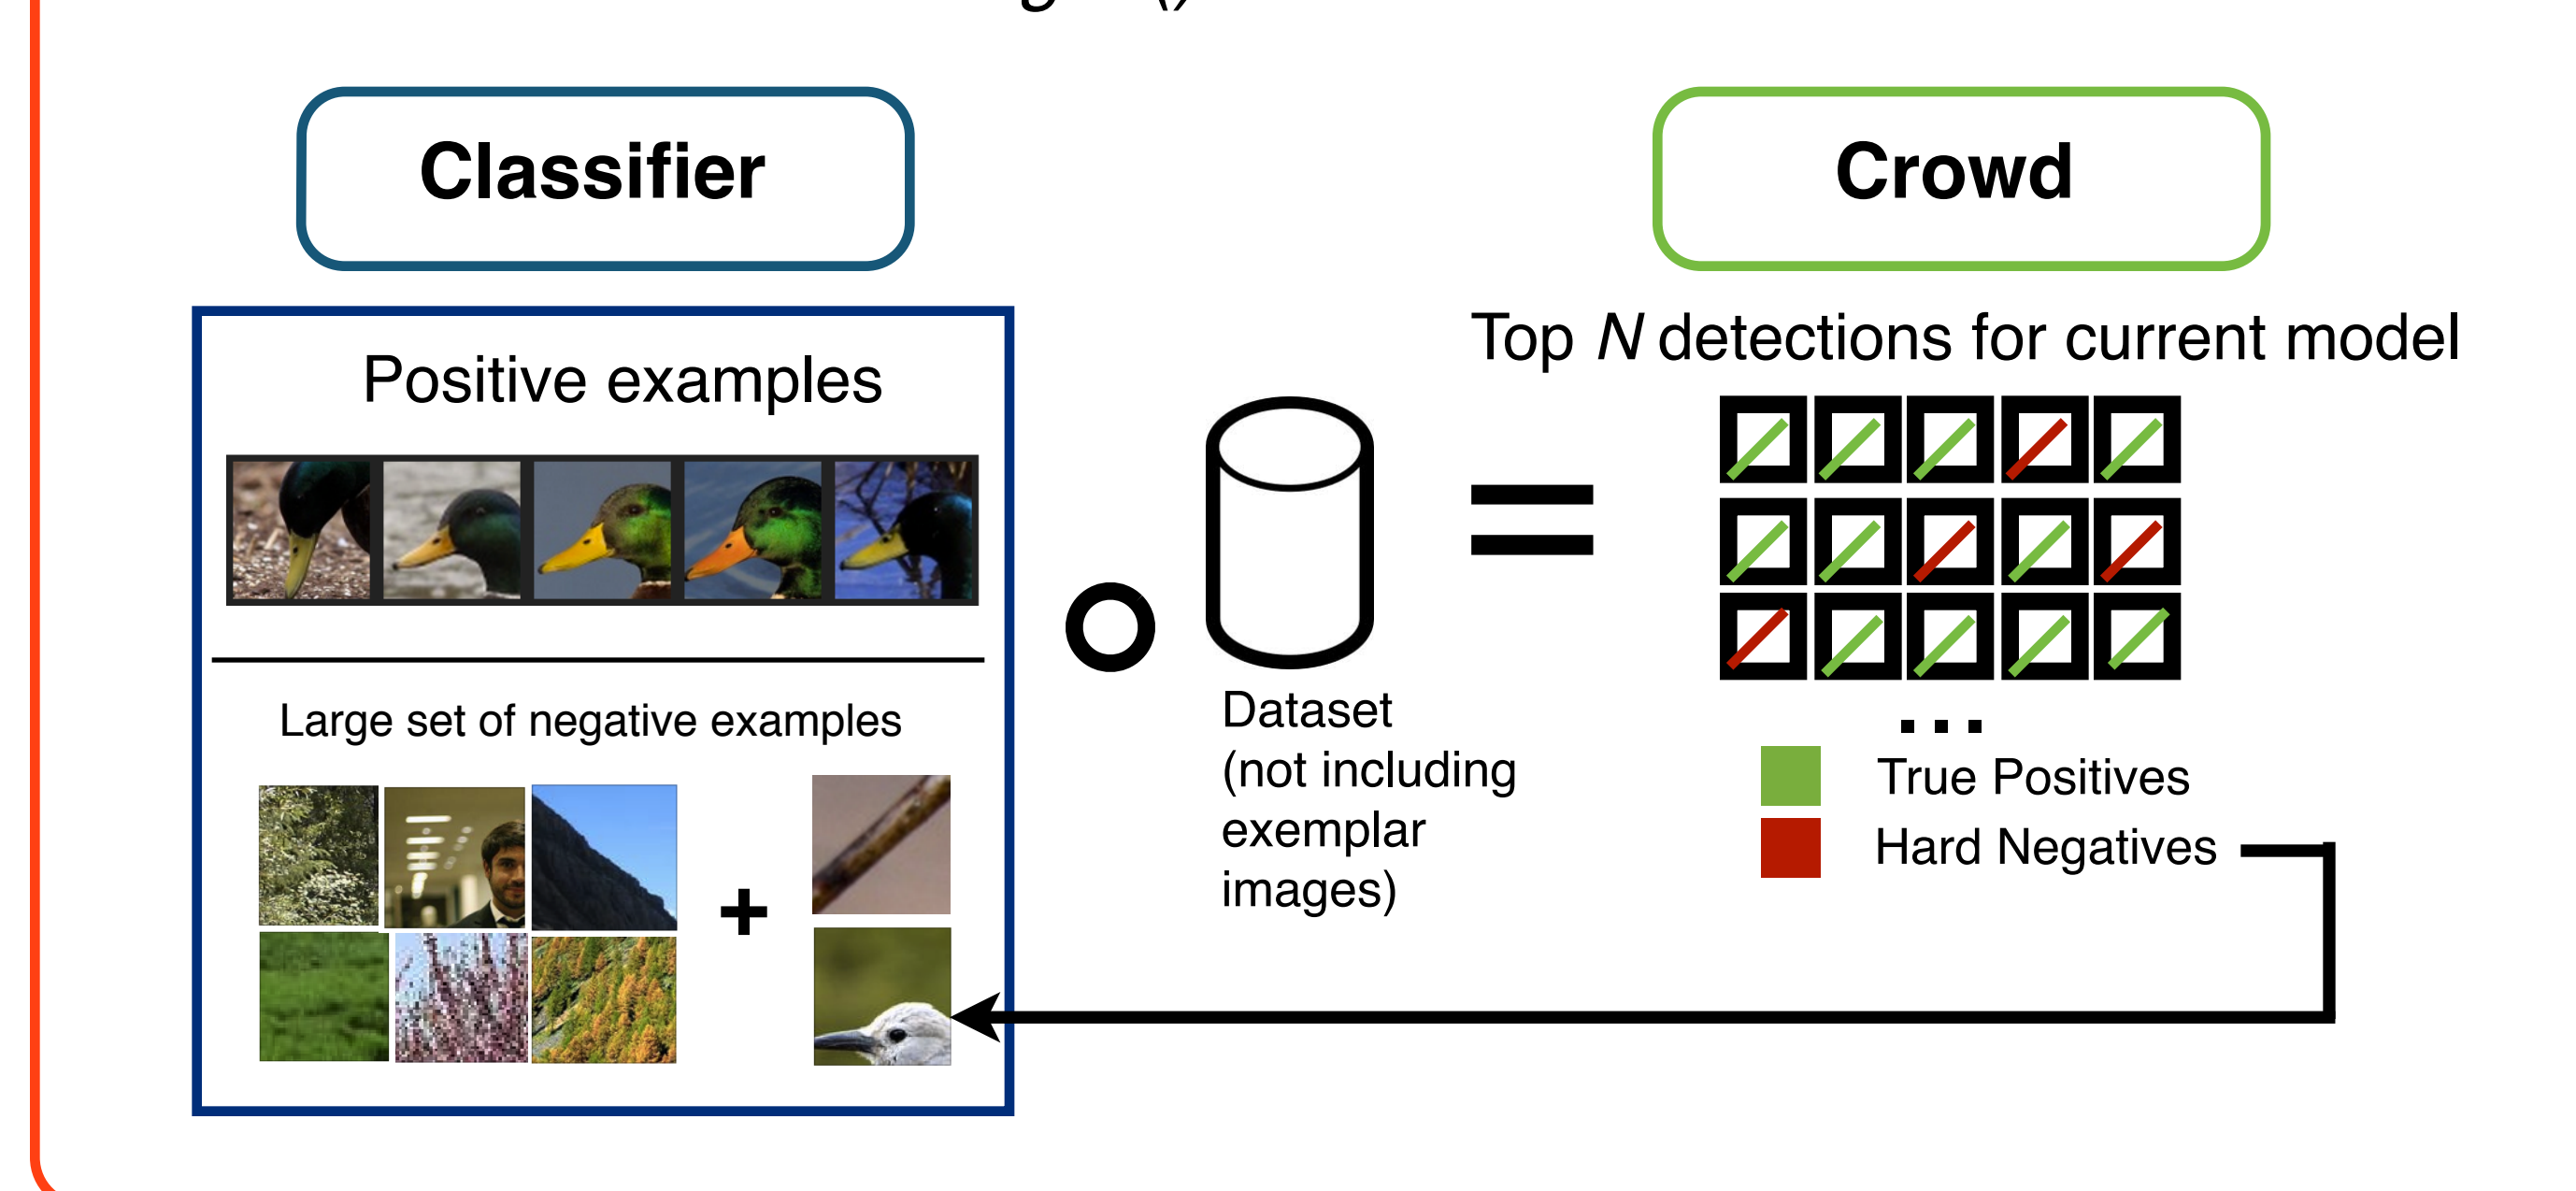
- Cardinal Head Pattern
- Humming Bird Ruby Throat
- Blue Jay Wing
- Mallard Beak
- Eagle Head Pattern
- Crested Head Pattern
- Mask Pattern
- Mottled Breast Pattern
- Ear Tufts
- ...

Exemplars



For all Exemplars

While model not converged()



Output:
 Final Set of Classifiers

Expert Review:

The initial experts can review detections for each attribute model, discarding models that are low-quality, e.g. low visual coherence, does not capture an aspect of the attribute.

Concept Drift:

When non-expert crowd members select hard negatives, sometimes the concept being learned can drift from the original exemplars. The example of 'shorts' to the right shows the classifier drifting to a more general 'leg-exposing garment' classifier. This problem does not occur if experts answer the active learning queries.

