

Seeing through the Human Reporting Bias: Visual Classifiers from Noisy Human-Centric Labels

Ishan Misra¹ * C. Lawrence Zitnick³ Margaret Mitchell² Ross Girshick³
¹Carnegie Mellon University ²Microsoft Research ³Facebook AI Research

Abstract

When human annotators are given a choice about what to label in an image, they apply their own subjective judgments on what to ignore and what to mention. We refer to these noisy “human-centric” annotations as exhibiting human reporting bias. Examples of such annotations include image tags and keywords found on photo sharing sites, or in datasets containing image captions. In this paper, we use these noisy annotations for learning visually correct image classifiers. Such annotations do not use consistent vocabulary, and miss a significant amount of the information present in an image; however, we demonstrate that the noise in these annotations exhibits structure and can be modeled. We propose an algorithm to decouple the human reporting bias from the correct visually grounded labels. Our results are highly interpretable for reporting “what’s in the image” versus “what’s worth saying.” We demonstrate the algorithm’s efficacy along a variety of metrics and datasets, including MS COCO and Yahoo Flickr 100M. We show significant improvements over traditional algorithms for both image classification and image captioning, doubling the performance of existing methods in some cases.

1. Introduction

Visual concept recognition is a fundamental computer vision task with a broad range of applications in science, medicine, and industry. Supervised learning of visual concept classifiers has been highly successful partly due to the use of large-scale, high-quality datasets (e.g., [8, 11, 29]). Depending on the complexity of the supported task, these datasets generally include annotations for 100s to 1000s of ‘typical’ concepts. To support an even broader range of applications, it is necessary to train classifiers for tens or even hundreds of thousands of visual concepts that may not be typical. Since supervised learning methods require exhaustive and clean annotations, one would require high quality datasets with orders of magnitude more annotations to train

*Work done during internship at Microsoft Research.

(a) A woman standing next to a **bicycle** with basket. (b) A city street filled with lots of people walking in the rain.



	Human Label	Visual Label
Bicycle	✓	✓



	Human Label	Visual Label
Bicycle	✗	✓

(c) A **yellow** Vespa parked in a lot with other cars.



	Human Label	Visual Label
Yellow	✓	✓

(d) A store display that has a lot of bananas on sale.



	Human Label	Visual Label
Yellow	✗	✓

Figure 1: Human descriptions capture only some of the visual concepts present in an image. For instance, the bicycle in (a) is described, while the bicycle in (b) is not mentioned. The Vespa in (c) is described as “yellow”, while the bananas in (d) are not, as being yellow is typical for bananas.

such methods. However, creating such datasets is expensive. An alternative approach is to relax this requirement of pristinely labeled data. The learning algorithm can be enabled to use readily-available sources of annotated data, such as user-generated image tags or captions from social media services like Flickr or Instagram. Such datasets easily scale to hundreds of millions of photos with hundreds of thousands of distinct tags [49].

Images annotated with human-written tags [49] or captions [6] focus on the most important or salient information in an image, as judged implicitly by the annotator. These annotations lack information on minor objects or information that may be deemed unimportant, a phenomenon known as *reporting bias* [16]. For example, Figure 1 illustrates two concepts (*bicycle*, *yellow*) that are each present in two images, but only mentioned in one. The bicycle may be considered irrelevant to the overall image in (b); and the bananas in (d) are not described as yellow because humans often omit an object’s typical properties when referring to it [31, 53]. Following [3], we refer to this type of labeling as *human-centric annotation*.

Training directly on human-centric annotations does not yield a credible visual concept classifier. Instead, it leads to a classifier that attempts to *mimic* the reporting bias of the annotators. To separate reporting bias from visual ground truth, we propose to train a model that explicitly factors human-centric label prediction into a *visual presence* classifier (*i.e.*, “Is this concept visually present in this image?”) and a *relevance* classifier (*i.e.*, “Is this concept worth mentioning in this image, given its visual presence?”). We train all these classifiers jointly and end-to-end as multiple “heads” branching from the same shared convolutional neural network (ConvNet) trunk [27, 46].

We demonstrate improved performance on several tasks and datasets. Our experiments on the MS COCO Captions dataset [6] show an improvement in mean average precision (mAP) for the learned visual classifiers when evaluated on both fully labeled data (using annotations from the MS COCO detection benchmark [29]) and on the human generated caption data. We also show that using such visual predictions improves image caption generation quality. Our results on the Yahoo Flickr 100M dataset [49] demonstrate the ability of our model to learn from “in the wild” data (noisy Flickr tags) and **double the performance** of the baseline classification model. Apart from just numerical improvements, our results are interpretable and consistent with research in psychology showing that humans tend not to mention typical attributes [31, 53] unless required for unique identification [44] or distinguishability [39, 48].

2. Related work

Label noise is ubiquitous in real world data. It can impact the training process of models and decrease their predictive accuracy [1, 21, 37]. Since there are vast amounts of cheaply available noisy data, learning good predictors despite the label noise is of great practical value.

The taxonomy of label noise presented in [14] differentiates between two broad categories of noise: *noise at random* and *statistically dependent noise*. The former does not depend on the data, while the latter does. In practice, one may encounter a combination of both types of noise.

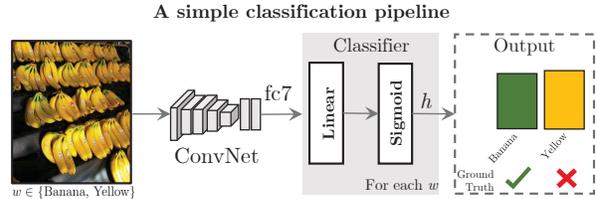


Figure 2: A simple classification model for learning from human-centric annotations. The noisy labels (*banana* is not annotated as *yellow*) impede the learning process.

Human-centric annotations [3] exhibit noise that is highly structured and shows statistical dependencies on the data [3, 14, 55]. It is structured in the sense that certain labels are preferentially omitted as opposed to others. Vision researchers have studied human-centric annotations in various settings, such as missing objects in image descriptions [3], scenes [4], and attributes [50] and show that these annotations are noisy [3]. Much of the work on learning from noisy labels focuses on robust algorithms [19, 32], voting methods [2], or statistical queries [22]. Some of these methods [19, 22] require access to clean oracle labels, which may not be readily available.

Explicitly modeling label noise has received increasing attention in recent years [35, 36, 47, 54]. Many of these methods operate under the “noise at random” assumption and treat noise as conditionally independent of the image. [26] models symmetric label noise (independent of the true label), which is a strong assumption for real world data. [36, 47] both model asymmetric label noise that is conditionally independent of the image. Such an assumption ignores the input image (and the objects therein) which directly affects the noisy annotations produced by humans [3].

Recently, Xiao *et al.* [54] introduced an image conditional noise model that attempts to predict what type of noise corrupts each training sample (no noise, noise at random, and structured label swapping noise). Unlike [54], our training algorithm does not require a small amount of cleanly labeled training data to bootstrap parameter estimation. Our model is also specifically designed to handle the noise found in human-centric annotations.

Bootstrapping [40], semi-supervised learning (SSL) [45, 57] and Positive Unlabeled (PU) learning [10, 28, 30, 34] are other ways of learning from noisy labeled data. However, they require access to clean oracle labels. SSL approaches are often computationally impractical [56, 58] or make strong independence assumptions [13] that do not hold in human-centric annotations. Our approach, which trains directly on noisy labels, can serve as a starting point for these approaches.

The work described here is also consistent with research in psycholinguistics on object reference and description. Such work demonstrates that humans store typical or

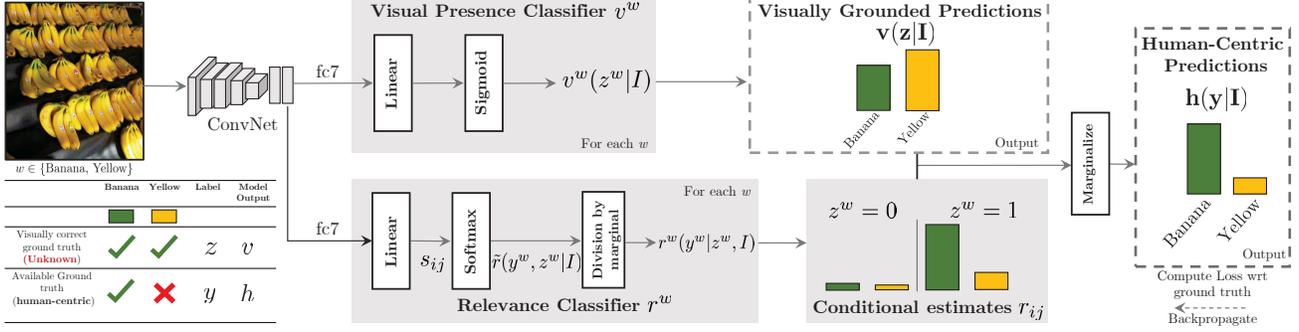


Figure 3: Our model uses noisy human-centric annotations y for learning visually grounded classifiers without access to the visually correct ground truth z . It uses two classifiers: a visual presence classifier v and a relevance classifier r . The visual presence classifier v predicts whether the visual concept w is visually present in an image. The relevance classifier r models the noise and predicts whether the concept should be mentioned or not. We combine these predictions to get the human-centric prediction h .

“prototypical” representations of objects and their properties [41]; and this background knowledge appears to have an effect on object description [31, 53]. People tend not to mention attributes that are obvious or typical for an object, preferring to name attributes required for conversational relevance [25], unique identification against alternatives [44, 51] and distinguishability [17, 23]. A similar separation between *what is observed* and *what is mentioned* falls out naturally from our proposed model.

3. Our Approach

Our goal is to train visually grounded image classifiers for a set of visual concepts $w \in \mathcal{W}$ (e.g., banana, yellow, zebra) using images and their human-centric annotations. The conventional approach to this problem, shown in Figure 2, is to naively apply a supervised learning algorithm: train a classifier h^w for each concept w , to predict its human-centric label $y^w \in \{0, 1\}$ (not mentioned vs. mentioned) as a conditional probability distribution $h^w(y^w|\mathcal{I})$, in which \mathcal{I} is an image.

The resulting classifier would attempt to mimic human reporting bias by predicting how a human would label the image \mathcal{I} regardless of whether w is visually present or absent. Thus, the predictions from each classifier h^w will not be visually grounded and do not meet our goal. How can we build classifiers that predict whether a visual concept w is present in an image and “see through” the noise in human-centric annotations?

3.1. Factor decoupling

We propose to structure each concept output $h^w(y^w|\mathcal{I})$ in terms of two classifiers v^w and r^w . The first classifier v^w models the conditional probability of the *visual presence* of the concept w in the image. The second classifier r^w models the conditional probability of the *relevance* of the

concept w , conditioned on the image and whether or not w is estimated to be visually present. The human-centric predictor is formed by marginalizing over the concept’s visual presence, as described next.

Let $z^w \in \{0, 1\}$ be a latent (or hidden) variable indicating whether the concept w is visually present in an image. Note that the training data only supplies human-centric labels y^w ; the true values of z^w are unknown during training. For instance, in Figure 1(a), $y^w = 1$ and $z^w = 1$ when the bicycle is present and mentioned, while $y^w = 0$ and $z^w = 1$ in Figure 1(b) when the bicycle is present but not mentioned. We refer to z^w as the *visual presence label*.

The conditional probability of the human-centric label given an image \mathcal{I} , $h^w(y^w|\mathcal{I})$, can now be computed by marginalizing over the latent visual presence label z^w :

$$h^w(y^w|\mathcal{I}) = \sum_{j \in \{0,1\}} r^w(y^w|z^w = j, \mathcal{I})v^w(z^w = j|\mathcal{I}). \quad (1)$$

An illustration of our model is shown in Figure 3. One important property of this formulation is that it allows the model to assign high confidence to unlabeled visual concepts: for an unmentioned concept ($y^w = 0$ and $z^w = 1$), the relevance classifier r^w allows the visual presence classifier v^w to assign a high probability to the true visual label ($z^w = 1$) while still making a prediction that matches the human-centric label ($y^w = 0$). This property enables the model to “see through” human reporting bias.

To simplify notation, we drop the concept index w from y, z, h, v and r when possible. We denote the probability values of r by:

$$r_{ij} = r(y = i|z = j, \mathcal{I}), \quad \forall (i, j) \in \{0, 1\}^2. \quad (2)$$

Another important property of the factorization in Equation (1) is that it provides a way to get two different predictions for the concept w in the same image. The model

can predict the visual presence of a visual concept; or predict how a human would annotate the image. Depending on the task at hand, one of these predictions may be more appropriate than the other, a point we demonstrate later via experimental results.

3.2. Model learning and parameterization

We estimate the model parameters by minimizing the regularized log loss of h^w , summed over all concepts w , on the training data annotated with human-centric labels. Since our model includes latent variables z for each concept and image, one approach could be to use Expectation Maximization (EM) [7]. We choose a direct optimization approach [5] over EM for simplicity, and thus both the model parameters and latent variable posteriors are updated and inferred online. In each SGD minibatch, the model predicts the conditional distributions r and v , marginalizes over the values of z , and uses the log loss of h to drive better estimates of r and v .

The conditional distributions r and v are modeled with a ConvNet [15, 27, 42]. As illustrated in Figure 3, the ConvNet trunk is shared between the two distributions (per concept) and then branches near the output into two sets of untied parameters. We jointly train one network for all visual concepts $w \in \mathcal{W}$ by treating learning as a multi-label classification problem. Further network architecture details are given in Section 4 with experiments.

The conditional probability distribution r models transition probabilities and thus its underlying joint distribution $\tilde{r}_{ij} = \tilde{r}(y = i, z = j | \mathcal{I})$ must be a valid probability distribution. To enforce this constraint, we directly estimate the joint distribution \tilde{r} with a softmax operation on a vector of unnormalized scores.

For each concept w , we first compute four scores s_{ij} using four linear models parameterized by weights m_{ij} and biases b_{ij} , and then normalize them using the softmax function to get a valid joint distribution \tilde{r}_{ij} :

$$s_{ij} = m_{ij}^T \phi(\mathcal{I}) + b_{ij}, \quad (3)$$

$$\tilde{r}_{ij} = \exp(s_{ij}) / \sum_{i'j'} \exp(s_{i'j'}). \quad (4)$$

For $\phi(\mathcal{I})$, we use global image features computed by the shared ConvNet trunk (e.g., fc7 layer activations from VGG16 [46]). These features capture the global image context, which is helpful in estimating r . Each r_{ij} can then be computed from \tilde{r} by dividing by the marginal $\tilde{r}(z = j | \mathcal{I})$:

$$r_{ij} = \tilde{r}_{ij} / \sum_{i'} \tilde{r}_{i'j}. \quad (5)$$

Figure 3 illustrates our process of computing r_{ij} from the image. Since our operations for estimating \tilde{r} (and thus r) are differentiable, we can backpropagate their errors to the ConvNet, allowing the full model to be trained end-to-end.

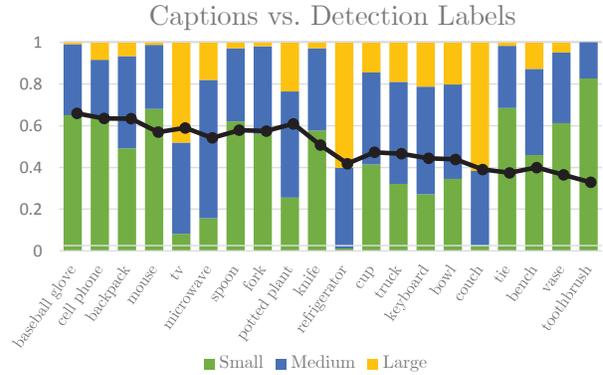


Figure 4: We use the MS COCO dataset to display the objects with the highest reporting bias. We use the detection labels for objects and see whether they are mentioned in the caption. The black line shows the probability of an object not being mentioned in the caption. The distribution of sizes for objects that are missed are shown by the color bars. Note that for many categories most unreported objects are small or medium in size (green and blue).

4. Experiments

We evaluate our proposed model on two datasets: Microsoft COCO [29] and a random subset of the Yahoo Flickr Creative Commons 100M (YFCC100M) dataset [49]. The YFCC100M dataset includes user-generated image tags, which we take as our source of human-centric annotations. For MS COCO, we take the supplied image captions [6] as human-centric annotations. We use the MS COCO object detection labels for dataset analysis and algorithm evaluation. These labels allow us to verify the accuracy of trained visual presence classifiers, v^w ; these labels are never used for training our model.

4.1. Experiments on MS COCO 1k visual concepts

Our first set of experiments use the MS COCO 1000 visual concepts from [12]. The visual concepts are the 1000 most common words in the MS COCO captions dataset [6] and include nouns, verbs, adjectives, and other parts of speech (see Table 1 for a breakdown).

For training, we generate image labels as 1000-dimensional binary vectors indicating which of the 1000 target visual concepts are present in the any of the 5 reference captions for each training image. The training set includes approximately 80k images. For evaluation, we follow [12] and split the val set into equally sized val and test sets of $\sim 20k$ images each; we use the same splits as in [12]. We report results on this 20k image test set.

4.1.1 Human reporting bias in image descriptions

We first analyze the annotation mismatch between the caption labels and the detection labels. We obtain labels for the

Table 1: mAP and PHR values on MS COCO captions ground truth (20k test images). We add our latent model to each baseline to make predictions that are visually grounded (v) or conform to human (h) labels. POS tags are as follows: Nouns (NN), Verbs (VB), Adjectives (JJ), Determiners (DT), Pronouns (PRP), Prepositions (IN).

		Mean Average Precision									Precision at Human Recall								
		Prob	NN	VB	JJ	DT	PRP	IN	Others	All	← Count	NN	VB	JJ	DT	PRP	IN	Others	All
			616	176	119	10	11	38	30	1000									
VGG16	MILVC [12]	-	41.6	20.7	23.9	33.4	20.4	22.5	16.3	34.0	52.7	32.8	40.5	40.3	32.2	33.0	24.6	45.8	
	MILVC + Multiple- fc8	-	41.1	20.9	23.7	33.6	21.1	22.8	16.8	33.8	51.2	32.6	40.8	41.1	31.7	33.5	27.3	45.0	
	MILVC + Latent (Ours)	v	42.9	21.7	24.9	33.1	19.6	23.0	16.2	35.1	53.6	35.4	43.3	41.3	28.0	36.0	24.4	47.2	
	MILVC + Latent (Ours)	h	44.3	22.3	25.8	34.4	21.8	23.6	17.3	36.3	55.5	36.3	44.7	42.9	32.1	37.3	26.4	48.9	
AlexNet	MILVC [12]	-	33.2	16.2	20.1	30.9	16.4	19.9	14.6	27.4	40.0	26.4	36.0	38.2	24.2	27.5	21.9	35.9	
	MILVC + Latent (Ours)	v	35.6	17.7	21.9	32.4	16.9	20.7	15.2	29.4	43.9	28.3	37.5	41.2	29.2	29.9	23.3	39.0	
	MILVC + Latent (Ours)	h	36.5	18.0	22.4	32.9	17.8	21.4	15.6	30.1	45.1	28.7	38.0	41.2	32.2	31.0	24.0	40.0	
VGG16	Classif.	-	34.9	18.1	20.5	32.8	19.2	21.8	16.3	29.0	42.5	30.4	33.9	40.5	30.4	30.7	23.8	38.2	
	Classif. + Multiple- fc8	-	34.2	17.7	19.9	32.6	19.0	21.5	15.9	28.4	41.3	27.9	32.3	39.6	29.6	31.2	22.6	36.8	
	Classif. + Latent (Ours)	v	37.7	19.6	22.0	32.6	20.2	22.0	16.3	31.2	46.3	32.9	36.8	38.9	32.3	33.1	27.0	41.5	
	Classif. + Latent (Ours)	h	38.7	20.1	22.6	33.8	21.2	23.0	17.5	32.0	47.8	33.7	37.9	42.5	34.2	34.4	29.0	42.9	

73 objects common in both the caption and object detection labels (see Section 4.1.3 for details). We use the notation from Section 3.1, and measure the human reporting bias as r_{01}^* : the probability of an object not being mentioned in the caption ground truth ($y = 0$) and being present in the detection ground truth ($z = 1$), over all the training images. To account for object size as a factor in an object not being mentioned, we split these measurements based on the size of the bounding box (sizes as defined in [29]). Figure 4 shows this mismatch for the top 20 objects with the highest r_{01}^* for values. A high r_{01}^* value indicates that there is a large mismatch between the caption ground truth and the detection ground truth – objects that are visually present but not mentioned in the captions.

As observed, there is a high degree of human labeling noise in the image descriptions, with the object with highest reporting bias mentioned roughly half as much as it appears.

4.1.2 Evaluating human-centric label prediction

We can evaluate our model in two ways: as a purely visual classifier (v) or as a predictor of human-centric (h) labels. We start with the latter and evaluate our model’s predictions against the human-centric MS COCO captions.

As a strong baseline, we use the recently proposed MILVC [12] approach. This method applies a ConvNet (VGG16 [46] or AlexNet [24])¹ in a fully-convolutional way to a large input image to generate a 12×12 grid of human-centric label predictions. It then uses a noisy-OR [52] to compute a single prediction from the 144 intermediate values. During training, the noisy-OR induces a form of multiple instance learning (MIL). Note that the baseline model estimates h labels directly without our decomposi-

tion into relevance and visual presence factors. As a second baseline (Classif.), we use a vanilla classification model akin to Figure 2, in which an ImageNet [43] pre-trained ConvNet is fine-tuned to directly predict human-centric labels.

We also include an additional baseline variant that adds extra parameters to control for the fact that our proposed model requires adding extra parameters. Specifically, we train a “Multiple- fc8 ” model for each method (MILVC and Classif.) that has the same number of parameters as our model. To train Multiple- fc8 , we add four extra randomly initialized fc8 (linear classification) layers, each with their own loss. At test time, we average the predictions of all the fc8 layers to get the final prediction.

We implement two variants of our model to parallel the MILVC and Classif. baselines. In the first variant (MILVC + Latent), v uses a noisy-OR over a 12×12 grid of visual presence predictions, and for r we average pool the 144 fc7 activation vectors to obtain a single 4096 dimensional $\phi(\mathcal{I})$ for Equation 3. The second variant (Classif. + Latent) generates a 1×1 output for v and fc7 , and therefore omits the noisy-OR and average pooling. In both cases, h predictions are obtained following Equation 1, using both r and v . Like the baselines, our model is trained to minimize the log loss (cross-entropy loss) over h .

To train our model, we set the joint noise distribution \tilde{r} to identity (i.e., $\tilde{r}_{11} = \tilde{r}_{00} = 0.5$) for the first two epochs and then update it for the last two epochs. Table 1 shows mean average precision (mAP) [11] and precision at human recall (PHR) [6] on the 20k test set. PHR is a metric proposed in [6], and measures precision based on human agreement. Briefly, this metric uses multiple references per image to compute a “human recall” value, an estimate of the probability that a human will use a particular word for an image.

¹Unless otherwise specified, we use VGG16 for all our experiments.

Table 2: Visual classification on 73 classes evaluated using fully labeled data from COCO.

	MILVC [12]	v	h	Using ground truth
mAP	63.7	66.8	66.5	76.3

Precision is then computed at this “human recall” value to get PHR. [6] shows that for the task of predicting visual concepts, PHR is a more stable metric than an AP metric, since it accounts for human agreement.

We report results for the 1000 visual concepts in aggregate, as well as grouped by their part-of-speech (POS) tags, on the MS COCO test split of 20k images (Table 1). Our latent variable model improves classification performance over all the baseline networks and architectures by **3 to 4 points** for both metrics (mAP and PHR). Interestingly, the Multiple- fc8 model, which has the same number of parameters as our latent model, does not show an improvement, even after extensive tuning of learning hyperparameters. This finding makes the contribution of the proposed model evident; the improvement is not simply due to adding extra parameters. It is worth noting that in Table 1, h is a better predictor of MS COCO caption labels than v , as h directly models the human-centric labels used for evaluation.

4.1.3 Evaluating visual presence prediction

The decoupling of visual presence (v) predictions and human-centric (h) label predictions allows our model to learn better visual predictors. To demonstrate this, we use the fully-labeled ground truth from the COCO detection annotations to evaluate the visually grounded v label predictions.

Since the 1000 visual concepts include many fine-grained visual categories (e.g., man, woman, child) and synonyms (e.g., bike, bicycle), we manually specify a mapping from the visual concepts to the 80 MS COCO detection categories, e.g., {bike, bicycle} \rightarrow bicycle. We find that 73 of the 80 detection categories are present in the 1000 visual concepts. We use this mapping only at evaluation time to compute the probability of a detection category as the maximum of the probabilities of its fine-grained/synonymous categories.

Table 2 shows the mean average precision (mAP) of our method, as well as the baseline on these 73 categories. As expected, using the human-centric model (h) for this task of visual prediction hurts performance (slightly). The performance drop is less dramatic when evaluating on these 73 classes because these classes have less label noise for large sized objects as compared to the 1000 visual concepts. We also train a noise-free reference model using the ground-truth visual labels from the detection dataset (i.e., the true values of the latent z labels).

Table 3: We show the importance of conditioning the relevance r on the input image. We measure the classification mAP on MS COCO 1k visual concepts using both the visually grounded (v) and the human-centric (h) predictions. Conditioning on the input image shows improvement over the baseline showing that human reporting bias statistically depends on the input image.

	MILVC [12]	w/o image		w/ image	
		v	h	v	h
mAP	34.0	34.2	34.3	35.1	36.3

4.1.4 Importance of conditioning on input images

A central point of this paper (and also in [3]) is that human-centric label noise is statistically dependent on image data. Here we demonstrate that our model is indeed improved by conditioning the noise (relevance) distribution on the input image, in contrast to previous work [36, 47] that estimates noise parameters *without* conditioning on the image.

To better understand the importance of this conditioning, we consider a model akin to [47]. We estimate the latent distribution r without conditioning on the input image, and compare it to our model that computes r conditioned on the image. Table 3 shows that mAP is significantly improved by conditioning on the image. When not conditioned on the image, only minor gains are achieved.

4.2. Experiments on Flickr image tagging

Datasets like MS COCO are curated by searching for images with specific objects [29]. In contrast, social media websites like Flickr contain much larger collections of images that are annotated with user-generated content such as tags, keywords, and descriptions. The words found in such data exhibit the human-centric annotation properties modeled by our approach.

We test our model on this “real world” data by using a random subset of $\sim 89k$ images from the YFCC100M dataset [49]. We ensure that these images have at least 5 and at most 30 human annotated tags that are present in the

Table 4: mAP values on a subset of YFCC100M. We add our latent model over the MILVC baseline to make predictions that are visually grounded (v) or that conform to human-centric (h) labels. POS tags: Nouns (NN), Verbs (VB), Adjectives (JJ), Pronouns (PRP), Prepositions (IN).

		Mean Average Precision							
		Count \rightarrow	Prob	NN	VB	JJ	PRP	IN	Others
VGG16	MILVC [12]	-	5.7	9.2	5.2	3.8	8.8	6.1	5.7
	MILVC + Multiple- fc8	-	4.6	6.2	3.8	2.7	7.3	3.1	4.5
	MILVC + Latent (Ours)	v	9.8	15.1	8.9	8.3	12.4	12.4	9.8
	MILVC + Latent (Ours)	h	11.2	15.4	9.9	8.2	16.3	12.5	11.2



Figure 5: Our model modifies visually correct detections to conform to human labeling. We show this modification for a few images of target visual concepts in the MS COCO Captions dataset. We first show the variation between h (y axis) and v values (x axis) for each concept in a 2D histogram. After thresholding at $v \geq 0.8$, we pick a representative image from each quantile of h (h increases from left to right). As you move from left to right, the model transitions from predicting that a human would not “speak” the word to predicting that a human would speak it. The human-centric h predictions of concepts depend on the image context, *e.g.*, `fence` at a soccer game vs. `fence` between a bear and a human (first row). Our model picks up such signals to not only learn a visually correct `fence` predictor, but also when a `fence` should be mentioned.

WordNet [33] lexicon. We split this dataset into 75k training images and 14k test images, and consider the top 1000 tags as the set of visual concepts. We train the baseline MILVC [12] model and our model for 4 epochs following the same hyperparameters used for MS COCO training.

Table 4 shows the numerical results of these models evaluated on the test set using the same human annotated tags. As explained in Section 4.1.2, we compare against the MILVC baseline, and a model with the same number of parameters as ours (denoted by Multiple-`fc8`). Our model has double the performance of the baseline MILVC model and **increases mAP by 5.5 points**.

4.3. Interpretability of the noise model

The relevance classifier r models human labeling noise conditioned on the image. Depending on the image, it can enhance or suppress the visual prediction for each concept. We show such modifications for a few visual concepts in

Table 5: LSTM captioning results on MS COCO

	Prob	BLEU-4	ROUGE	CIDEr
MILVC [12]	-	27.7	51.8	89.7
MILVC + Latent (Ours)	h	29.2	52.4	92.8

Figure 5. After thresholding at $v \geq 0.8$, we pick a representative image from each quantile of h (h increases from left to right). The variation in h values for these high confidence $v \geq 0.8$ images (shown in a 2D histogram in each row) indicates that h and v have been decoupled by our model. The images show that our model captures subtle nuances in the ground truth, *e.g.*, mention a hat worn by a cat, do not mention the color of a pumpkin, definitely mention pink sheep, *etc.* It automatically captures that context is important for certain objects like `fence` and `hat`, while certain attributes are worth mentioning to help distinguish objects like the `orange` pillow. Such connections have been shown in both vision research [3] and psychology [17, 44].

4.4. Correcting error modes by decoupling

Modeling latent noise in human-centric annotations allows us to learn clean visual classifiers. In Figure 6, we compare our model’s visual presence v predictions with the baseline (MILVC) and show a few error modes that it corrects. Our model is able to correct error modes like misspellings (`desert` vs. `dessert` in the first row), localizes objects correctly and out of context (`fridge` in the second row, `net` in the first row, *etc.*) and is better at counting (`zebra`, `banana` last row).

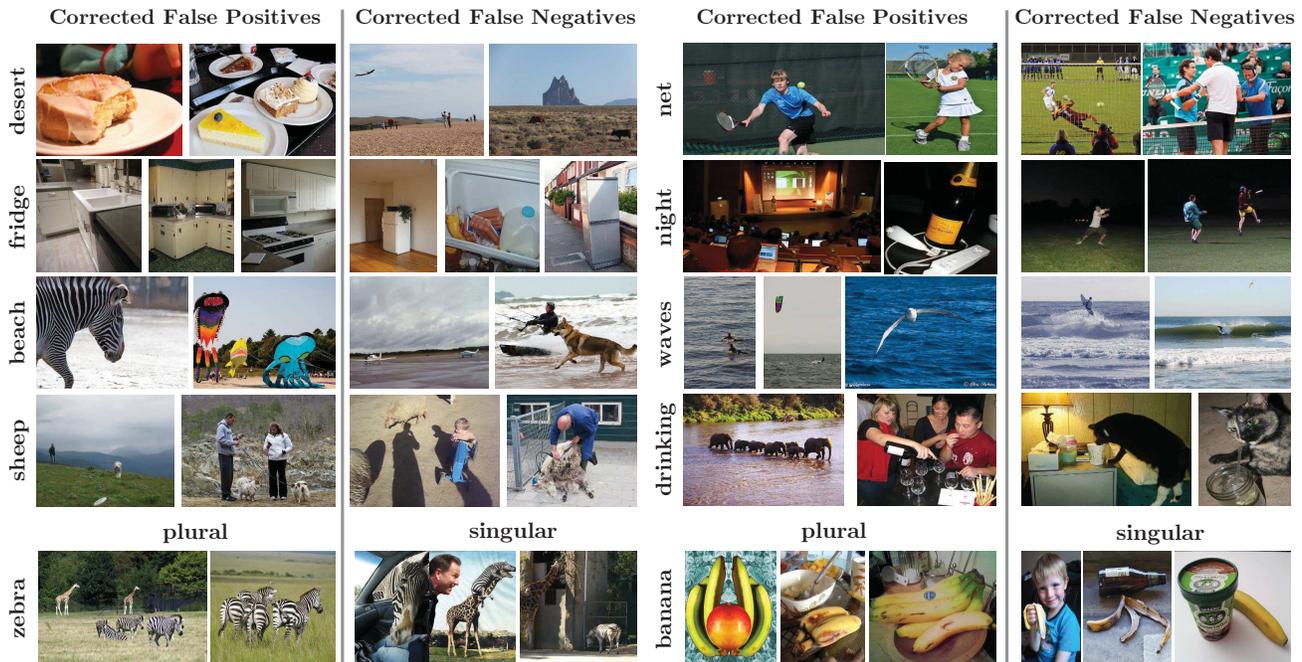


Figure 6: Our model learns clean visual predictors from noisy labels. Here we show corrected false positives: MILVC incorrectly reports a high probability ($h \geq 0.75$) for the concept, while our model correctly reports a low probability ($v \leq 0.3$); and corrected false negatives: MILVC incorrectly reports a low probability ($h \leq 0.3$) for the concept, while our model correctly reports a high probability ($v \geq 0.75$). For example, consider zebra vs. zebras, and banana vs. bananas in the last row, where our model correctly “counts” compared to the baseline. Images are from the MS COCO Captions dataset.

4.5. Using word detections for caption generation

We now look at the task of automatic image caption generation and show how our model can help improve the task. We consider a basic Long Short-Term Memory (LSTM) [18] network to generate captions. We use 1000 cells for the LSTM, and learn a 256 dimensional word embedding for the input words. Following [9], our vocabulary consists of words with frequency ≥ 5 in the input captions. The image features (1000 visual concept probabilities) are fed once to the LSTM as its first hidden input. We train this LSTM over all the captions in the MS COCO caption training data for 20 epochs using [20, 38]. We use beam size of 1 for decoding. Table 5 shows the evaluation of the automatically generated captions using standard captioning metrics. Using the probabilities from our model shows an improvement for all evaluation metrics. Thus, modeling the human-reporting bias can help downstream applications that require such human-centric predictions.

5. Discussion

We have introduced an algorithm that explicitly models *reporting bias* — the discrepancy between what exists and what people mention — for image labeling. By introducing a latent variable to capture “what is in an image” separate from “what is labeled in an image”, we leverage *human-*

centric annotations of images to their full potential, inferring visual concepts present in an image separately from the visual concepts worth mentioning. We demonstrate performance improvements over previous work on several tasks, including image classification and image captioning. Further, the proposed model is highly interpretable, capturing which concepts may be included or excluded based on the context and dependencies across visual concepts. Initial inspection of the model’s predictions suggests consistency with psycholinguistic research on object description, with typical properties noticed but not mentioned.

The algorithm and techniques discussed here pave the way for new deep learning methods that decouple human performance from algorithmic understanding, modeling both jointly in a network that can be trained end-to-end. Future work may explore different methods to incorporate constraints on the latent variables, or to estimate their posteriors (such as with EM). Finally, to fully exploit the enormous amounts of data which exist “in the wild”, algorithms that explicitly handle noisy data are essential.

Acknowledgments: We thank Jacob Devlin, Lucy Vanderwende, Frank Ferraro, Sean Bell, Abhinav Shrivastava, and Saurabh Gupta for helpful discussions. Devi Parikh and Dhruv Batra for their suggestions and organizing the fun ‘snack times’.

References

- [1] K. Barnard and D. Forsyth. Learning the semantics of words and pictures. In *ICCV*, 2001. 2
- [2] E. Beigman and B. B. Klebanov. Learning with annotation noise. In *ACL-IJCNLP*, 2009. 2
- [3] A. C. Berg, T. L. Berg, H. Daume III, J. Dodge, A. Goyal, X. Han, A. Mensch, M. Mitchell, A. Sood, K. Stratos, et al. Understanding and predicting importance in images. In *CVPR*, 2012. 2, 6, 7
- [4] A. Borji, D. N. Sihite, and L. Itti. What stands out in a scene? a study of human explicit saliency judgment. *Vision research*, 91, 2013. 2
- [5] O. Chapelle. Modeling delayed feedback in display advertising. In *KDD*, 2014. 4
- [6] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollar, and C. L. Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 2, 4, 5, 6
- [7] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977. 4
- [8] J. Deng, W. Dong, R. Socher, L. jia Li, K. Li, and L. Feifei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [9] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 8
- [10] C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. In *SIGKDD*, 2008. 2
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007). 1, 5
- [12] H. Fang, S. Gupta, F. N. Iandola, R. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, C. L. Zitnick, and G. Zweig. From captions to visual concepts and back. In *CVPR*, 2015. 4, 5, 6, 7
- [13] R. Fergus, Y. Weiss, and A. Torralba. Semi-supervised learning in gigantic image collections. In *NIPS*, 2009. 2
- [14] B. Fréney and M. Verleysen. Classification in the presence of label noise: a survey. *NNLS*, 25, 2014. 2
- [15] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202, 1980. 4
- [16] J. Gordon and B. V. Durme. Reporting bias and knowledge extraction. In *Automated Knowledge Base Construction (AKBC) 2013: The 3rd Workshop on Knowledge Extraction, at CIKM 2013, AKBC'13*, 2013. 2
- [17] R. L. Gregory. Eye and brain: The psychology of seeing. 1966. 3, 7
- [18] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 8
- [19] H. Izadinia, B. C. Russell, A. Farhadi, M. D. Hoffman, and A. Hertzmann. Deep classifiers from image tags in the wild. In *Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions*. ACM, 2015. 2
- [20] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACMM*, 2014. 8
- [21] A. Joulin, L. van der Maaten, A. Jabri, and N. Vasilache. Learning visual features from large weakly supervised data. *arXiv preprint arXiv:1511.02251*, 2015. 2
- [22] M. Kearns. Efficient noise-tolerant learning from statistical queries. *JACM*, 45, 1998. 2
- [23] R. Koolen, A. Gatt, M. Goudbeek, and E. Krahrmer. *Journal of Pragmatics*, 43(13):3231–3250, 2011. 3
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 5
- [25] A. Kronfeld. Conversationally relevant descriptions. *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, 1989. 3
- [26] J. Larsen, L. Nonboe, M. Hintz-Madsen, and L. K. Hansen. Design of robust neural network classifiers. In *Acoustics, Speech and Signal Processing*, volume 2, 1998. 2
- [27] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989. 2, 4
- [28] X. Li and B. Liu. Learning to classify texts using positive and unlabeled data. In *IJCAI*, volume 3, 2003. 2
- [29] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*. 2014. 1, 2, 4, 5, 6
- [30] B. Liu, Y. Dai, X. Li, W. S. Lee, and P. S. Yu. Building text classifiers using positive and unlabeled examples. In *ICDM*, 2003. 2
- [31] M. M., R. E., and van Deemter K. Typicality and object reference. 2013. 2, 3
- [32] N. Manwani and P. Sastry. Noise tolerance under risk minimization. *Cybernetics*, 43, 2013. 2
- [33] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 7
- [34] I. Misra, A. Shrivastava, and M. Hebert. Watch and learn: Semi-supervised learning of object detectors from videos. In *CVPR*, 2015. 2
- [35] V. Mnih and G. E. Hinton. Learning to label aerial images from noisy data. In *ICML*, 2012. 2
- [36] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari. Learning with noisy labels. In *NIPS*. 2013. 2, 6
- [37] D. F. Nettleton, A. Orriols-Puig, and A. Fornells. A study of the effect of different types of noise on the precision of supervised learning techniques. *Artificial intelligence review*, 33, 2010. 2
- [38] NLPcaffe. <http://github.com/Russell91/NLPCaffe>. 8
- [39] D. R. Olson. Language and thought: Aspects of a cognitive theory of semantics. *Psychological review*, 77, 1970. 2
- [40] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, and A. Rabinovich. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*, 2014. 2
- [41] E. Rosch. Principles of categorization. *Concepts: core readings*, pages 189–206, 1999. 3
- [42] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Technical report, DTIC Document, 1985. 4

- [43] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115, 2015. 5
- [44] J. C. Sedivy. Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of psycholinguistic research*, 32(1), 2003. 2, 3, 7
- [45] A. Shrivastava, S. Singh, and A. Gupta. Constrained semi-supervised learning using attributes and comparative attributes. In *ECCV*. 2012. 2
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 4, 5
- [47] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, and R. Fergus. Training convolutional networks with noisy labels. In *ICLR Workshop*, 2015. 2, 6
- [48] J. M. Tenenbaum. On locating objects by their distinguishing features in multisensory images. *Computer Graphics and Image Processing*, 2, 1973. 2
- [49] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. The new data and new challenges in multimedia research. *arXiv preprint arXiv:1503.01817*, 2015. 1, 2, 4, 6
- [50] N. Turakhia and D. Parikh. Attribute dominance: What pops out? In *ICCV*, 2013. 2
- [51] I. van der Sluis, A. Gatt, and K. van Deemter. Manual for the tuna corpus: Referring expressions in two domains. *Technical Report AUCS/TR0705*, 2006. 3
- [52] P. Viola, J. C. Platt, and C. Zhang. Multiple instance boosting for object detection. In *NIPS*, 2006. 5
- [53] H. Westerbeek, R. Koolen, and A. Maes. Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology*, 6, 2015. 2, 3
- [54] T. Xiao, T. Xia, Y. Yang, C. Huang, and X. Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015. 2
- [55] K. Yun, Y. Peng, D. Samaras, G. J. Zelinsky, and T. Berg. Studying relationships between human gaze, description, and computer vision. In *CVPR*, 2013. 2
- [56] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. *Advances in neural information processing systems*, 16(16):321–328, 2004. 2
- [57] X. Zhu. Semi-supervised learning. In *Encyclopedia of Machine Learning*, 2010. 2
- [58] X. Zhu, Z. Ghahramani, J. Lafferty, et al. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, volume 3, pages 912–919, 2003. 2