



A Step Toward More Inclusive People Annotations for Fairness

Candice Schumann
Google
United States
cschumann@google.com

Susanna Ricco
Google
United States
ricco@google.com

Utsav Prabhu
Google
United States
utsavprabhu@google.com

Vittorio Ferrari
Google
Switzerland
vittoferrari@google.com

Caroline Pantofaru
Google
United States
cpantofaru@google.com

ABSTRACT

The Open Images Dataset [16] contains approximately 9 million images and is a widely accepted dataset for computer vision research. As is common practice for large datasets, the annotations are not exhaustive, with bounding boxes and attribute labels for only a subset of the classes in each image. In this paper, we present a new set of annotations on a subset of the Open Images dataset called the *MIAP (More Inclusive Annotations for People)* subset, containing bounding boxes and attributes for all of the people visible in those images. The attributes and labeling methodology for the *MIAP* subset were designed to enable research into model fairness. In addition, we analyze the original annotation methodology for the *person* class and its subclasses, discussing the resulting patterns in order to inform future annotation efforts. By considering both the original and exhaustive annotation sets, researchers can also now study how systematic patterns in training annotations affect modeling.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision**; Machine learning; • **Social and professional topics** → *User characteristics*.

KEYWORDS

datasets, fairness, computer vision

ACM Reference Format:

Candice Schumann, Susanna Ricco, Utsav Prabhu, Vittorio Ferrari, and Caroline Pantofaru. 2021. A Step Toward More Inclusive People Annotations for Fairness. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3461702.3462594>



This work is licensed under a Creative Commons Attribution International 4.0 License.

AIES '21, May 19–21, 2021, Virtual Event, USA.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8473-5/21/05.

<https://doi.org/10.1145/3461702.3462594>

1 INTRODUCTION

The Open Images Dataset [16] contains approximately 9 million images with annotations supporting research into image classification, object detection and instance segmentation, visual relationship detection, and spatially localized voice and text descriptions of the image contents. It is widely accepted as a key dataset to both measure and advance the state of the art in image and scene understanding.

In this paper, we introduce a new set of annotations, released as the *MIAP (More Inclusive Annotations for People)* collection in Open Images Extended, which improve the bounding box annotations for the five original classes in the *person* subtree. The annotations use a new labeling protocol designed to produce annotations that enable fairness analysis. With the increasing focus on reducing unfair bias as part of responsible AI research, we hope these annotations will encourage researchers already leveraging Open Images to incorporate fairness analysis in their research.

Open Images contains image-level labels for more than 19,000 classes, and bounding box annotations for a subset of 600 of those classes. Even for those 600 classes, it is infeasible to exhaustively label each of them as present (positive) or absent (negative) in every image. Hence, Open Images annotations are not exhaustive, as is common practice for large scale computer vision datasets, e.g. [7, 17]. Each image is associated with a small set of positive labels and a small set of negative labels. These labels are proposed by an ensemble of image classifiers (e.g., labels produced with confidence above a threshold by some algorithm) and then verified as either positive or negative by a human annotator. The presence/absence of any label not proposed by the initial image classifier is simply unknown.

After this stage of producing image-level labels, annotators then draw bounding boxes for the classes at the most specific level of the 600-class hierarchy (leaves) that are included in the image-level positive label set.

Of the 600 classes with bounding-boxes, five classes are part of the *person* hierarchy: *person*, *man*, *woman*, *boy*, and *girl*.

Here, Open Images departs from other large-scale object detection datasets [17, 22] by including these subclasses. Although most applications require annotations for only the encompassing *person* class, the availability of gender- and age range-specific labels would in theory allow for use of these labels for fairness analysis and bias mitigation. However, we found the specific labeling procedure used impacted the accuracy of these subgroup labels and resulted in

non-exhaustive bounding box annotations for the *person* superclass for some images.

Our new *MIAP* annotations collapse the gender- and age range-specific subclasses into a single *person* superclass at annotation time. We ask annotators to draw bounding boxes around *all* people appearing in an image with any positive image-level label of either *person*, *man*, *woman*, *boy*, or *girl*. *Person* boxes are then annotated with additional attributes corresponding to the perceived age range and perceived gender presentation of the individual. This procedure adds a significant number of boxes that were previously missing due to complex interactions between the scene content, the characteristics of the individual as interpreted by the initial image classifier, and the perception of two sets of human annotators. Our boxes give more complete ground truth for training a person detector as well as more accurate subgroup labels for fairness analysis and bias mitigation.

1.1 Summary of Contributions

The main contributions of this paper are:

- An analysis of the original Open Images annotation pipeline’s impact on the completeness of person-related box annotations.
- An updated annotation procedure to produce more complete annotations for the *person* class.
- The resulting annotations on a subset of images and a comparative analysis of the differences between the two annotation types. We publicly release our new annotations.

1.2 Intended Use

We include annotations for perceived age range and gender presentation for *person* bounding boxes because we believe these annotations are necessary to advance our field’s ability to better understand and work to mitigate and eliminate unfair bias or disparate performance across protected subgroups. We note that the labels capture the gender and age range presentation as assessed by a third party based on visual cues alone, rather than an individual’s self-identified gender or actual age. We do not support or condone building or deploying gender and/or age presentation classifiers trained from these annotations as we believe the risks associated with the use of these technologies outside fairness research outweigh any potential benefits. We refer readers to the excellent work detailing these potential risks (e.g., [12, 15]) for more information.

2 RELATED WORK

Research into techniques to improve fairness in computer vision algorithms is of increasing importance, as these systems become widely used. Existing approaches span the model development cycle, including (a) careful formulation of the task, (b) identifying and reducing potential biases in training datasets, (c) inventing techniques for training fairer models in spite of biased datasets, (d) defining and evaluating whether the resulting system achieves some desired fairness property. This work is mainly focused on (b) and (d).

2.1 Identifying and reducing bias in datasets

Bias can manifest in a dataset as unintended consequence of design choices made during dataset construction. Datasets containing people are particularly sensitive to these biases, and can have harmful downstream implications - such as resulting in widely-used systems that have unequal performance across population demographics. For example, datasets that sample images from internet sources may over-represent some geographical regions compared to others. Both Open Images and its predecessor ImageNet have been shown to contain a disproportionate number of images from North America and Europe [6, 25]. Classifiers trained on this distribution were more likely to misclassify images from the developing world for both person-related classes (e.g., *bridegroom*) [25] and object classes (e.g., *spices*) [6].

To partially address this limitation and encourage further research [2], Open Images now includes the Crowdsourced Extension¹, containing new images from countries underrepresented in the original set. This Open Images extension is focused on image-level labels for non-human classes. The faces of people appearing in this set are blurred, making it less suitable for analysis of person detectors.

Of course, the potential for bias in large web-scale datasets goes far beyond skewed geographic distributions. The exceptional work updating the person subtree of ImageNet [31] provides detailed analysis of sources of bias, from the definitions of class labels to the distribution of people representing given classes. Their work focuses on correcting image-level labels while we chose to focus on the bounding box annotations for people in Open Images. With our updated annotations, we provide the attribute labels similar to those annotated at the image level in ImageNet. We hope future research can leverage both datasets to build on this work, improving these and future large-scale computer vision datasets.

2.2 Identifying and reducing bias in trained models

Creating better balanced and less biased datasets are steps toward the eventual goal of ML models that work well for everyone, whose predictions do not unfairly disadvantage any segment of the population. Most definitions of model fairness (e.g., equality of opportunity [13]) require partitioning an evaluation set by the attributes to which model performance should be agnostic. Proposed standards for model transparency reporting [21], which complement proposed transparency reports for datasets [10], advocate for always reporting performance metrics broken down by relevant population groups, independent of specific fairness goals. Finally, many bias mitigation strategies require attribute labels on at least a portion of the model’s training set [3, 19].

There is a rich history of datasets containing images of people that include attribute annotations such as perceived gender and/or age. For some, these annotations were added for the explicit purpose of furthering ML fairness research. For example, the Pilot Parliaments Benchmark from Gender Shades [5] used gender and skin type annotations to highlight disparate performance of available gender classifiers across intersectional groups. Similarly, datasets

¹<https://storage.googleapis.com/openimages/web/extended.html>

for training or evaluating face recognition algorithms may include annotations for fairness analysis [11, 26]. In other cases, these annotations were initially curated as training data for attribute classifiers (e.g., Celeb-A [18], FairFace [14], UTKFace [32]) but may also be useful for studying bias [1, 23].

The attribute labels we provide differ from previous datasets in two key areas. First, our annotations do not rely on machine learned models for annotating either perceived gender or age range. Other large-scale annotation efforts exploring bias in ImageNet representation [9] or attempting to quantify the diversity in appearance across the population [20] resort to attribute labels produced by trained attribute classifiers. We prefer to reduce the influence of potentially biased attribute classifiers, especially when using these annotations for fairness applications. Second, and importantly, most other datasets with similar annotations for fairness applications are targeted toward either face attribute estimation or face recognition as their primary applications. As a result, they contain a single face of interest per image. In contrast, we augment a portion of the existing Open Images dataset, thus adding fairness annotations to images with multiple people per image, including for people whose faces are not clearly visible.

To our knowledge, the only other dataset containing multiple people in complex scenes that has been augmented with attribute annotations for the purpose of fairness research is BDD100k [29]. This work annotates Fitzpatrick skin type for localized pedestrians of sufficient size in the images. Skin tone is an important attribute to consider for fairness, and one that our work does not address. While both datasets can be used for evaluating whether a person detector performs consistently across population slices, BDD100k is limited to self-driving applications. In contrast, our annotations also provide data for fairness analysis of a broad set of scene understanding tasks.

3 ANALYSIS OF THE ORIGINAL OPEN IMAGES

In this section we take a deep dive into the effects of the partially annotated ground truth label generation process on the annotation of people boxes. Prior to this deep dive we do want to mention that the original Open Images dataset functions correctly as a standard partial ground truth dataset. However, due to the non-exhaustive nature of the annotations, fairness evaluations or bias mitigation using this dataset could provide potentially misleading results.

3.1 The Annotation Pipeline

Open Images employs a two-stage process, first annotating image-level labels and then object bounding-boxes based on them (Figure 1). In total, while Open Images considers a very large set of 19,794 possible classes at the image-level, object bounding-boxes are annotated for 600 of them. We focus on these from now on.

Given an image, an image classifier is used to generate potential positive and negative labels. These potential labels are then sent to annotators to verify. Each potential label is either marked as positive (i.e. the named object class is present in the image) or negative (i.e. the named object class does not appear in the image). The presence or absence of any label that is not proposed by the

image classifier is considered simply *unknown* (i.e. it is neither positive nor negative).

Annotators are then asked to draw bounding-boxes around objects for the positive image labels, focusing on the most specific labels in the hierarchy (Figure 2). For example, if an image has positive labels of “car”, “limousine”, and “screwdriver”, the annotators box all limousines and screwdrivers (since limousine is more specific than car in the hierarchy). This means that if the image contains a car that is not a limousine, there will not be a box drawn around it. Moreover, there will also be no boxes for classes that are unknown to be present or absent in the image (i.e. not proposed by the initial classifier, see above).

Note that this partial annotation protocol is working as intended – it would be a huge undertaking to box everything in every image and partial labeling is standard practice [7, 16, 17]. The Open Images protocol enables users to reconstruct the list of classes for which boxes might be missing, and take this into account when training or evaluating detectors (i.e. all unknown labels for an image, plus all positive labels which are not the most specific for that image). That being said, partial annotation can have unintended consequences, especially when annotating people.

3.2 Annotation Pipeline Mixing with Societal Norms

Many standard academic datasets for object detection (such as MSCOCO [17] and ImageNet [22]) contain a single *person* class. Instead, the Open Images hierarchy contains five classes that correspond to localized people: *person*, *man*, *woman*, *boy*, and *girl*. While these gender- and age-specific labels are relevant and useful for some applications (such as fairness evaluations or bias mitigation), for others, the additional specificity is unnecessary and even undesirable.

The most obvious approach for training a detector with a single gender- and age-agnostic *person* class would be to simply drop *man*, *woman*, *boy*, *girl* from the label set. However, because the hierarchical labeling process preferred labeling for a subclass when both superclass and subclass appeared in the positive image-level labels, dropping annotations for these four subclasses will result in dropping many ground truth boxes. For example, a training image containing a single man, with positive image-level labels for both *person* and *man* will contain a box labeled as *man* but no corresponding box labeled as *person* will exist.

Another option is to collapse the hierarchy to a single class, treating boxes labeled as *man*, *woman*, *boy*, *girl* as *person* boxes. This remapping could happen as a preprocessing step before training or as postprocessing at inference time. For the moment, we focus on remapping prior to training and defer a discussion of postprocessing to Section 3.3.

The previous example is resolved correctly with label preprocessing, with the *man* box remapped to *person* providing correct ground truth for that image. Unfortunately, this procedure does not always result in a complete set of bounding boxes at training time, which requires localizing every person in the image. For a person appearing in an image to have a corresponding bounding box after label remapping given the original annotation process, one of two things had to be true:




Figure 1: The Open Images annotation pipeline from raw images to labeled object bounding boxes.




Figure 2: Visualization of the Open Images hierarchy of classes that are boxed [16] to demonstrate its scope and complexity.




Figure 4: Image with true image-level labels *person*, *man*, *suit*, *clothing*.



Figure 3: Image with true image-level labels *person*, *woman*, *suit*.

- The image classifier proposed *person* as a potential image-level label for the image without proposing any gender- and age-specific subclass that the annotators marked as positive, or
- The image classifier must have proposed the subclass that both the image-level annotator and box-level annotators agreed described the specific person.

We show a few examples below where these conditions are not satisfied.

Gender Presentation. After-the-fact label manipulation can cause individuals of different genders to be treated differently in the resulting annotations.

Consider the case of the image in Figure 3. The following image-level labels were proposed and marked positive: *person*, *woman*, and *suit*. Using the hierarchical labeling strategy, annotators were asked to draw boxes around objects of the classes *woman* and *suit*, but not *person*. As a result, the annotations for this image are missing the bounding boxes around the men wearing the suits.

Note that this problem can also happen in the other direction. See Figure 4, where the image-level label set includes *man* and *person*, but not *woman*.

The cause of the missing boxes in these examples is obvious: the imbalance of the specificity of the image-level labels proposed for inclusion in the positive set, coupled with the decision to prefer subclasses for box annotations. But subtle omissions can occur even in cases where the positive image-level labels include all five *person*-related classes.

Consider another image containing five people: one man, one woman, one girl, one boy, and one person with unidentifiable or nonconforming gender presentation. Assume (unlike the previous example), that the image classifier successfully proposed all relevant labels: *person*, *man*, *woman*, *boy*, and *girl*. Bounding-box annotators would be instructed to draw boxes around individuals corresponding to the four subclasses. Depending on the specific decisions made by the annotators for this image, this process could potentially miss annotations for people that do not fit into a gender-stereotyped role.




Figure 5: Image with a negative label for *girl* and no *woman* replacement

Age and language. Similar problems arise in instances where the image classifier and annotators perceive the age of the image subjects differently. In Figure 5, we show an example where an age-dependent image-level label (*girl*) is labeled as negative with no replacement of the age-appropriate label (*woman*). The result is similar to Figure 4, with a bounding box annotation for only one of the two people in the image.

In Figure 6, the *girl* image-level label is considered positive, but an annotator only perceived that label as applying to two people in the image. As a result, several people of similar ages do not have boxes drawn around them.

We observed that phenomena of this type were more likely to cause missed bounding boxes for images with boxes for *girl* or *woman* than for images with *man* or *boy*. This may be the result of the difference in how these terms are used and understood by society. The word *girl* can be used to describe a female of any age while *boy* typically describes only a child or adolescent.²

Potential correlation with image context. Because the image content influences which labels are proposed in the classification step of the process (step 2 in Figure 1), it affects which boxes are included in the annotations. Partial annotations may be more prevalent, and also potentially more concerning, in scenes that correspond to existing societal stereotypes. Take, for example, the three images in Figure 7. In traditional heterosexual European and US-style weddings the focus of the day is generally placed on the bride. This emphasis is reinforced in the annotations, with a bounding box for the woman in Figure 7a but not for the man. As another example, in the US the majority of military personnel are male [8]; a prior that is reinforced

²For example, Wikipedia defines *girl* as “A girl is a young female human, usually a child or an adolescent. When she becomes an adult, she is described as a woman. The term girl may also be used to mean a young woman, and is sometimes used as a synonym for daughter. Girl may also be a term of endearment used by an adult, usually a woman, to designate adult female friends.” [28] while *boy* is defined as “A boy is a young male human. The term is usually used for a child or an adolescent. When a male human reaches adulthood, he is described as a man.” [27]



Figure 6: Several people of similar age, only two have bounding boxes drawn around them.

in the missing bounding box annotation for the woman in uniform in Figure 7b.

3.3 Potential risks of direct use of non-exhaustive annotations

A person detector trained on these partial annotations may learn to imitate the non-exhaustive patterns in the training data. As a result, a trained detector may be less likely to correctly detect individuals from certain subpopulations in certain contexts. This is also true for detectors that rely on postprocessing to collapse predictions of *man*, *woman*, *boy*, and *girl* to *person*, where correct output would rest on the ability of a specific classifier (e.g., *woman*) to generalize to scene contexts that may not have been sufficiently represented in the non-exhaustive training data.

The severity of such errors depends on the use case in question, ranging from annoyances like a less efficient experience searching through personal image collections to highly consequential failures in life safety systems (e.g., pedestrian detection in assistive driving features).

The degree to which trained models are resilient to noisy annotations remains an active area of research. In particular, we do not yet understand the effects that *structured patterns* of the type we observe here have on model performance; typically studies of robustness in the presence of noise make simplifying assumptions and experiment with annotations that are modified or missing uniformly at random [30]. Especially concerning here is that the ground truth in the validation and test sets are also non-exhaustive, so it may be difficult to detect disparate performance of a model inheriting bias from these annotations. A detailed study of these important questions warrants significant future work that our *MIAP* annotations are designed to support.

4 MORE INCLUSIVE ANNOTATIONS FOR PEOPLE

To generate a subset with complete ground truth for the *person* class and enable related fairness research, we relabel 100K randomly sampled images (70,000 from training and 30,000 from validation/test)



Figure 7: Image context and gender norms sometimes dictate which image-level labels are flagged.

that contained at least one of the five *person* classes (*man*, *woman*, *boy*, *girl*, or *person*) in the positive image-level set. Annotators were asked to draw boxes around all people in these images. They were then asked to provide two different attributes for each person.

Perceived Gender Presentation. Annotators selected either *predominantly feminine*, *predominantly masculine*, or *unknown* to describe the human-perceived gender presentation of an individual based on the visual cues in the image.

We recognize that gender is not binary and that an individual’s gender identity may not match their perceived or intended gender presentation. We chose to use *predominantly feminine* and *predominantly masculine* rather than the historical labels *female* and *male* to acknowledge that an individual’s gender presentation is complex while maintaining the ability to analyze these labels in the context of the existing image-level and box labels of *man*, *woman*, *boy*, and *girl*. When used in aggregate, these annotations are useful to assess how model performance may differ for people who present gender differently.

We did not include *non-binary* as a class label as it is not possible to label gender identity from images. Gender identity should only be used in situations where participants are able to self-report gender [24].

In an effort to mitigate the effects of unconscious bias on the annotations, we reminded annotators that norms around gender expression vary across cultures and have changed over time. Annotation guidelines included statements like:

No single aspect of a person’s appearance “defines” their gender expression. For example, a person may still present as predominantly masculine while wearing jewelry. Another may present as predominantly feminine while having short hair. Consider the entirety of a person’s appearance in the image and avoid rigid definitions based on cultural stereotypes.

In addition to the written instructions, we provided annotators with visual examples that included people across a variety of ages and from different cultures. These sets included visual examples to counteract common stereotypes (e.g., including images of women with short hair).

Finally, annotators were instructed to select the *unknown* label in situations where there was not enough information in the image to make a determination. Note that annotators could decide to select

Datasets	Train	Val	Test	Total
images	70,000	7,410	22,590	100,000
samples				
images with missing boxes	23,241	1,836	5,397	30,474

Table 1: Distribution of the images selected for relabeling. The second row counts the number of images where the *MIAP* annotations add boxes that do not exist in the original annotations.

either *predominantly feminine* or *predominantly masculine* even if an individual’s face was not visible.

Perceived Age Range. Annotators selected either *young*, *middle*, *older*, or *unknown* to describe the perceived age range of an individual based on their appearance in the image. Annotators were instructed to prefer the older of two categories in situations where there was enough information to form an impression but were unsure of a boundary case. Someone who appears old enough to possibly belong to *middle* should be assigned that attribute label. Someone who appears old enough to possibly belong to *older* should be assigned that attribute label.

Annotators could select *unknown* for perceived age range while selecting one of the other three labels for perceived gender presentation. Similarly they could select *young*, *middle*, or *older* even if they annotated the same bounding box as *unknown* for perceived gender presentation.

5 ANALYSIS

It is important to note the differences between the original Open Images dataset and the new *MIAP* dataset to understand what has improved and what still needs work.

Raw counts. In Table 1, we see that approximately 30% of the 100,000 images had additional boxes added during the relabeling effort. This is quite a large percentage of the images, underscoring how often the phenomena described in Section 3.2 did, in fact, occur.

Table 2 compares the total number of boxes in the original annotations to the *MIAP* set, both in aggregate and broken down by attribute label. This analysis assigns a bounding box in the original annotation the two attribute labels according to the class labels

Datasets	Train		Validation		Test		Total	
	original	MIAP	original	MIAP	original	MIAP	original	MIAP
boxes	294,695	363,165	15,398	22,060	47,777	69,106	357,870	454,331
<i>predominantly feminine</i>	67,979	83,004	2,039	4,453	6,265	13,215	76,283	100,672
<i>predominantly masculine</i>	127,468	141,300	3,779	7,898	12,073	24,849	143,320	174,047
<i>unknown</i>	99,248	138,861	9,580	9,709	29,439	31,042	138,267	179,612
gender presentation								
<i>young</i>	16,634	22,621	1,205	1,508	3,709	4,677	21,548	28,806
<i>middle</i>	178,813	190,179	4,613	10,591	14,629	32,904	198,055	233,674
<i>older</i>	-	7,176	-	471	-	1,376	-	9,023
<i>unknown</i> age range	99,248	143,189	9580	9490	29,439	30,149	138,267	182,828

Table 2: Comparison counts of boxes across the 100,000 samples from the original Open Images dataset and the MIAP dataset. Counts are split between the training set, the validation set, and the testing set, with the totals appearing in the last two columns. Gender presentation and age range presentation for the original set was assumed via the class label (for example, the class *woman* would have gender presentation *predominantly feminine* and age range presentation *middle*). There was no class for *older* in the original Open Images dataset.

(e.g., *woman* maps to attribute labels *predominantly feminine* and *middle*).

There is a large increase in *unknown* labels (for both gender presentation and age range) in the training set but only a small increase in the validation and test sets. This is explained by the difference in annotation procedure between training and validation/test in the original protocol. The training set used the process described in Section 3.1. For the validation and test sets, annotators drew boxes for all positive image-level labels. This means that if the label *person* appeared in the positive image-level labels for a validation or testing image, the annotator would be asked to draw generic *person* boxes, regardless of the presence of subclass labels. Boxes for which annotators were not able to assign attribute labels are likely to have been drawn in response to a prompt for *person* boxes but not drawn when only a more specific class was requested.

We see a significant increase for both *predominantly feminine* and *predominantly masculine* labels across the training, validation, and test splits. Although this labeling effort has decreased the gap between *predominantly masculine* and *predominantly feminine* slightly, the dataset is still unbalanced with the majority of gender presentation labels, after *unknown*, belonging to *predominantly masculine*. *Middle* is by far the most common of the age range attribute labels.

Statistical analysis. In order to measure the association of attribute labels with missing boxes we use normalized pointwise mutual information (*nPMI*) [4]. Pointwise mutual information is defined as

$$PMI(x, y) = \ln \frac{p(x, y)}{p(x)p(y)}$$

where x and y are discrete random variables. Normalized pointwise mutual information is therefore defined as

$$nPMI(x, y) = \frac{PMI(x, y)}{-\ln p(x, y)}.$$

We use *nPMI* to calculate the co-occurrences of images with missing boxes and a given label (for example the co-occurrence of images with missing boxes and gender presentation *predominantly feminine*), denoted as $nPMI(\text{Missing}, \text{label})$, as well as the co-occurrence of images without missing boxes and a given label (for example the co-occurrence of images without missing boxes and gender presentation *predominantly feminine*), denoted as $nPMI(\text{Not Missing}, \text{label})$. Using these two co-occurrences we can take the difference

$$nPMI(\text{Missing}, \text{label}) - nPMI(\text{Not Missing}, \text{label})$$

to find labels that are more likely to co-occur with images that are missing people boxes than images that are not missing people boxes in the original annotations.

Table 3 shows these results in order of significance. Unsurprisingly *unknown* labels for both age range and gender presentation are much more likely to be found in images with missing boxes. This is mostly likely due to the training set labeling process only annotating the most specific *person* classes. Additionally, the label *predominantly feminine* is much more likely to co-occur with missing boxes than the label *predominantly masculine*, possibly due to the syntactic confusion between *woman* and *girl*. The age range labels exhibit the least amount of skew between images with missing boxes and those without.

Empirical visualization. In addition to exploring raw counts and statistical analysis, we visualized the differences between the original Open Images and the MIAP annotations. Figure 8 shows four examples. In each image, the original Open Images had at least one person not annotated. The new MIAP dataset has bounding box annotations for all people.



Figure 8: Examples of new boxes in *MIAP*. In each subfigure the magenta boxes are from the original Open Images dataset, while the yellow boxes are new boxes from the *MIAP* dataset.

Label	nPMI difference
gender presentation <i>unknown</i>	0.458
age range <i>unknown</i>	0.453
<i>predominantly feminine</i>	0.365
<i>predominantly masculine</i>	0.251
<i>young</i>	0.236
<i>middle</i>	0.156
<i>older</i>	0.155

Table 3: nPMI differences between images with missing boxes and images without missing boxes ($nPMI(\text{Missing, label}) - nPMI(\text{Not Missing, label})$) in order of significance. A high nPMI difference indicates that images with missing boxes were more likely to have new boxes with the given labels.

6 DISCUSSION

Exhaustively annotated, large-scale visual data is particularly valuable, since it often enables new types of analyses that are ineffectively addressable by partial annotations. With the exhaustive box annotations for the *person* class, we hope to close this gap for fairness-related analysis of people-centric ML models. We appreciate and encourage its use for tasks such as disaggregating/stratifying metrics along age range and gender presentation slices, in the hope that such analyses will eventually lead to less biased models, classifiers, and workflows that work equitably for all.

It is important to note that the gender presentation and age range presentation labels, while useful for evaluation purposes, do not represent gender identity or actual age of the people. This data should not be used to create gender and/or age classifiers. As a further precaution, a gender presentation of *unknown* has been set for any box labeled as *young* in the released dataset³.

³Analysis included in this paper includes gender presentation annotations for people marked as *young*.

6.1 Limitations

The continually evolving and expanding nature of Internet imagery implies that the distribution of Flickr images included within the original Open Images dataset (and consequently, the subset included in this relabeling effort) may not accurately represent the uses of models trained or evaluated using this data. Additionally, recent works have demonstrated that the geographical sampling of Flickr images as well as the use of English as the primary language for dataset construction and taxonomy definition result in inherent cultural bias within the datasets [6].

While this effort makes initial strides towards relabeling the dataset, there is much work yet to be done. Several parts of the original annotation pipeline for the Open Images dataset were not addressed - including image selection, object class hierarchy, determination of classes that are boxed, and annotation of relationship triplets. Our image selection task rejected images which did not contain at least one *person*-subclass label, which may have propagated/amplified any biases that may have existed in the labeling workflow for the original dataset. While this effort has slightly aided in balancing the bounding box annotations, there remains an imbalance towards a gender presentation of *predominantly masculine*. This relabeling effort did not improve age range imbalance. Owing to the large amount of manual effort required to exhaustively label each image, this task was also restricted only to certain *person* labels (*man/woman/girl/boy/person*), but not all *person*-related labels that exist in the Open Images taxonomy (e.g. parts of the human body, activities, clothing, etc). Other, non-*person*-related classes in the Open Images taxonomy may also benefit from similar analysis and exhaustive relabeling, but were not addressed in this paper. Relationship triplets were not re-annotated, so they only exist for the subset of boxes that were included in the original dataset.

We acknowledge that there are trade-offs between providing the capability to do fairness analysis based on societal concepts of perceived gender, and perpetuating binary definitions of gender. Future work could include analysis of objective attributes that can confound computer vision models, such as makeup or facial hair.

7 CONCLUSION

Dataset construction, annotation procedures, fairness analysis and techniques for modeling under noisy ground truth conditions are all continually evolving areas of research.

In this paper, we have presented a new annotation procedure and the resulting annotations for a subset of the Open Images dataset, titled *MIAP (More Inclusive Annotations for People)*. The new annotation procedure results in more annotations for the *person* class, as well as annotations for gender presentation and age range presentation. The annotation procedure was additionally designed to enable fairness analysis.

Our analyses of the two annotation pipelines and resulting annotation patterns demonstrate the importance of understanding annotation design decisions and their consequences. We have discussed how annotation procedures can be seen to interact with societal norms, language, and image context, as well as other data patterns, here resulting in differences for the *person* class and subclass labels. The original annotations and new annotations for Open Images are both useful in different situations since they vary in scope and coverage, or they can be used as complementary information. We hope that the findings in this paper can inform usage of the current dataset as well as future dataset annotation efforts.

We encourage researchers to use our new *MIAP* annotations to incorporate fairness analysis into their research. Given the exhaustive nature of our updated annotation procedure, researchers can expect more robust fairness evaluations.

IMAGE CITATIONS

All images used in this paper are licensed under Creative Commons By 2.0 from Open Images. References for each image are as follows:

- Figure 3: thalisantunes, <https://www.flickr.com/photos/thalisantunes/21322592792/>
- Figure 4: Gobierno de Chile, <https://www.flickr.com/photos/gobiernodechile/16822006848/>
- Figure 5: Waifer X, <https://www.flickr.com/photos/waiferx/85365663/>
- Figure 6: Alex Erde, <https://www.flickr.com/photos/alexerde/9658256603/>
- Figure 7a: sapex, <https://www.flickr.com/photos/sapexoxo/4884905970/>
- Figure 7b: 29th Combat Aviation Brigade Public Affairs, <https://www.flickr.com/photos/29thcab/6155989139/>
- Figure 8a: Grigory Kravchenko, https://www.flickr.com/photos/grisha_21/3464190158/
- Figure 8b: Quincena Musical, <https://www.flickr.com/photos/quincenamusal/15026832341>
- Figure 8c: Simon Fraser University - Communications & Marketing, <https://www.flickr.com/photos/sfupamr/6238487408/>
- Figure 8d: Emma Racila, <https://www.flickr.com/photos/94605075@N05/8617525388/>

REFERENCES

- [1] M. Alvi, A. Zisserman, and C. Nellaker. 2018. Turning a Blind Eye: Explicit Removal of Biases and Variation from Deep Neural Network Embeddings. In *ECCV Workshop on Bias Estimation in Face Analytics*.
- [2] James Atwood, Yoni Halpern, Pallavi Baljekar, Eric Breck, D. Sculley, Pavel Ostyakov, Sergey I. Nikolenko, Igor Ivanov, Roman Solovyev, Weimin Wang, and Miha Skalic. 2020. The Inclusive Images Competition. In *The NeurIPS '18 Competition*, Sergio Escalera and Ralf Herbrich (Eds.). Springer International Publishing, Cham, 155–186.
- [3] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Allison Woodruff, Christine Luu, Pierre Kreitmann, Jonathan Bischof, and Ed H Chi. 2019. Putting fairness principles into practice: Challenges, metrics, and improvements. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 453–459.
- [4] Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL* (2009), 31–40.
- [5] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Fairness, Accountability, and Transparency (FAT*)*. PMLR, 77–91.
- [6] Terrance de Vries, Ishan Misra, Changhan Wang, and Laurens van der Maaten. 2019. Does Object Recognition Work for Everyone?. In *CVPR Workshop on Fairness Accountability Transparency and Ethics in Computer Vision (FATE CV)*.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 248–255.
- [8] DoD. 2018. *2018 Demographics Report: Profile of the Military Community*. Technical Report. Department of Defence, Office of the Deputy Assistant Secretary of Defense for Military Community and Family Policy (ODASD (MC&FP)).
- [9] Chris Dulhanty and Alexander Wong. 2019. Auditing ImageNet: Towards a Model-driven Framework for Annotating Demographic Attributes of Large-Scale Image Datasets. In *CVPR Workshop on Fairness Accountability Transparency and Ethics in Computer Vision (FATE CV)*.
- [10] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [11] Patrick Grother, Mei Ngan, and Kayee Hanaoka. 2018. Ongoing Face Recognition Vendor Test (FRVT) Part 1: Verification. *National Institute of Standards and Technology* (2018).
- [12] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. 2018. Gender Recognition or Gender Reductionism? The Social Implications of Embedded Gender Recognition Systems. In *CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada). 1–13. <https://doi.org/10.1145/3173574.3173582>
- [13] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc., 3315–3323. <https://proceedings.neurips.cc/paper/2016/file/9d2682367c3935defcb1f9e247a97c0d-Paper.pdf>
- [14] Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 1548–1558.
- [15] Os Keyes. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proceedings of the ACM on Human-Computer Interaction*, Article 88 (Nov. 2018), 22 pages. <https://doi.org/10.1145/3274357>
- [16] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision (IJCV)* (2020), 1–26.
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*. Springer, 740–755.
- [18] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In *International Conference on Computer Vision (ICCV)*.
- [19] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2018. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*. PMLR, 3384–3393.
- [20] Michele Merler, Nalini Ratha, Rogerio S Feris, and John R Smith. 2019. Diversity in faces. *arXiv preprint arXiv:1901.10436* (2019).
- [21] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Fairness, Accountability, and Transparency (FAT*)*. 220–229. <https://doi.org/10.1145/3287560.3287596>
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115, 3 (2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [23] Hee Jung Ryu, Hartwig Adam, and Margaret Mitchell. 2017. Inclusivefacenet: Improving face attribute detection with race and gender diversity. *arXiv preprint arXiv:1712.00193* (2017).
- [24] Morgan Klaus Scheuerman, Katta Speil, Oliver L. Haimson, Foad Hamidi, and Stacy M. Branham. 2020. HCI Guidelines for Gender Equity and Inclusivity. <http://morgan-klaus.com/gender-guidelines.html>. Version 1.1.

- [25] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. 2017. No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World. In *NeurIPS Workshop: Machine Learning for the Developing World*.
- [26] Craig I Watson. 2016. NIST Special Database 18. NIST Mugshot Identification Database (MID).
- [27] Wikipedia contributors. 2020. Boy — Wikipedia, The Free Encyclopedia. <https://en.wikipedia.org/w/index.php?title=Boy&oldid=992668770>. [Online; accessed 29-January-2021].
- [28] Wikipedia contributors. 2021. Girl — Wikipedia, The Free Encyclopedia. <https://en.wikipedia.org/w/index.php?title=Girl&oldid=1000129400>. [Online; accessed 29-January-2021].
- [29] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. 2019. Predictive Inequity in Object Detection. In *CVPR Workshop on Fairness Accountability Transparency and Ethics in Computer Vision (FATE CV)*.
- [30] Zhe Wu, Navaneeth Bodla, Bharat Singh, Mahyar Najibi, Rama Chellappa, and Larry S Davis. 2018. Soft sampling for robust object detection. *arXiv preprint arXiv:1806.06986* (2018).
- [31] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. 2020. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. In *Fairness, Accountability, and Transparency (FAT*)*.
- [32] Zhifei Zhang, Yang Song, and Hairong Qi. 2017. Age Progression/Regression by Conditional Adversarial Autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.