

Consistency and Accuracy of CelebA Attribute Values

Haiyu Wu¹, Grace Bezdol¹, Manuel Günther², Terrance Boulton³,
Michael C. King⁴, Kevin W. Bowyer¹

¹University of Notre Dame, ²University of Zurich,

³University of Colorado Colorado Springs, ⁴Florida Institute of Technology

Abstract

We report the first systematic analysis of the experimental foundations of facial attribute classification. Two annotators independently assigning attribute values shows that only 12 of 40 common attributes are assigned values with $\geq 95\%$ consistency, and three (high cheekbones, pointed nose, oval face) have essentially random consistency. Of 5,068 duplicate face appearances in CelebA, attributes have contradicting values on from 10 to 860 of the 5,068 duplicates. Manual audit of a subset of CelebA estimates error rates as high as 40% for (no beard=false), even though the labeling consistency experiment indicates that no beard could be assigned with $\geq 95\%$ consistency. Selecting the mouth slightly open (MSO) for deeper analysis, we estimate the error rate for (MSO=true) at about 20% and (MSO=false) at about 2%. A corrected version of the MSO attribute values enables learning a model that achieves higher accuracy than previously reported for MSO. Corrected values for CelebA MSO are available at <https://github.com/HaiyuWu/CelebAMS0>.

1. Introduction

Facial attributes have potential uses in face matching/recognition [3, 4, 12, 13, 17, 22], face image retrieval [15, 18], re-identification [21, 23, 24], training GANs [5, 6, 10, 14] for generation of synthetic images, analyzing AI biases [29, 30] and other areas. CelebA [16] is the largest and most widely used dataset in this research area. However, recent papers have described cleaning the identity groups in CelebA [27], and suggested that the facial attribute annotations [8, 26] need some “cleaning”. This paper provides the first analysis of the consistency with which the commonly-used face attributes can be manually marked, and also of the quality of the attribute values distributed with the CelebA images. We also propose an auditing workflow to clean existing annotations, and demonstrate that using a corrected set of




Figure 1. Which Set of Attribute Values Enables Learning a Better Model? Lower left quadrant contains images with original (MSO=false) corrected to (MSO=true); Upper right quadrant contains images with original (MSO=true) corrected to (MSO=false).

attribute values enables learning a substantially different and more accurate model. Contributions of this work include:

- Analysis of independent manual annotation of the 40 commonly used face attributes shows that only 12 are labeled with $\geq 95\%$ consistency and 3 have random (50%) consistency. (See Section 3.)
- For the 12 attributes that we determined can be labeled with $\geq 95\%$ consistency across annotators, we audit the attribute values provided with the CelebA images and find that (1) the error rate is

often asymmetric between true / false, and (2) the error rate is as high as 40% for some attribute values. (See Section 3.)

- We propose a semi-automated workflow to clean existing annotations, and use it to create corrected MSO attribute values for CelebA. In part of this correction/cleaning, we identify (1) a small number of images that are unusable and we propose should be dropped from CelebA, and (2) identify images for which true/false cannot be assigned to a particular attribute and so we propose an “information not visible” value must be introduced. (See Sections 4.1 and 4.2.)
- Comparing models learned using the original MSO values versus our cleaned values, we show that the models are substantially different, and that our cleaned values enable a model that achieves state-of-the-art accuracy on MSO. (See Section 4.3.)

2. Related work

There is a large literature in facial attribute analysis, and several surveys give a broad coverage of the field [1, 2, 26, 31]. We cover only a few of the most relevant works here.

CelebA was introduced by Liu et al. [16] in 2015 specifically to support research in deep learning for facial attributes. CelebA has 202,599 images grouped into 10,177 identities. Each image has 40 true/false attributes – pointy nose, oval face, gray hair, wearing hat, etc. – and five landmark locations. CelebA also has a recommended subject-disjoint split into train, validation and test. The creation of the attribute values is described only as – “Each image in CelebA and LFWA is annotated with forty facial attributes and five key points by a professional labeling company” [16]. No description of how the attribute values are created, or estimate of their consistency or accuracy, is given.

Thom and Hand stated [26] that, “CelebA and LFWA were the first (and only to date) large-scale datasets introduced for the problem of facial attribute recognition from images. Prior to CelebA and LFWA, no dataset labeled with attributes was large enough to effectively train deep neural networks.” In number of identities and images, CelebA is substantially larger than LFWA, and is the most-used research dataset in this area. In addition, Thom and Hand speculate that noise in the attribute values may lead to an apparent plateau in research progress [26] – “There is a recent plateau in facial attribute recognition performance, which could be due to poor labeling of data. . . . While crowdsourcing such tasks can be very useful and result in large quantities of reasonably labeled data, there are some tasks which may

be consistently labeled incorrectly ...”.

The only paper we are aware of to discuss cleaning CelebA is [27]. They deal specifically with errors in the identity groupings, and do *not* consider errors in the attribute values. Compared to the original 202,599 images of 10,177 identities, their identity-cleaned version has 197,477 images of 9,996 identities. In a simple manual check of 100 identities in CelebA, we found a few additional instances of identity errors in the identity-cleaned version of [27]. Because our work is focused on attribute classification and not face matching, we start with the original CelebA rather than the version of [27].

Terhörst et al. [25] firstly estimated the quality of annotations in CelebA by letting three human evaluators manually evaluate randomly selected 50 positively-annotated and 50 negatively-annotated images for each attribute. They claimed that “Similar to LFW, there is a tendency that most of the wrong annotations are within the positives”. In this paper, we provide a more statistically and systematically analysis on each attribute, and we, furthermore, provide a possible solution to clean the dataset.

Motivated by comments in [8, 26] on the importance of correct labels for machine learning and errors encountered in CelebA attribute values, we present the first detailed analysis of the accuracy of CelebA attribute values. We create a cleaned version of the MSO values, and perform experiments to assess the impact of the original versus cleaned MSO attribute, using AFFACT [8], MOON [20], DenseNet [11] and ResNet [9]. We show that the cleaned attribute values result in learning a more coherent model that achieves higher accuracy.

3. Accuracy of Attributes In Training Data

We first examine the general consistency of manual annotations of face attributes. Then we use duplicate faces in CelebA to examine the consistency of the attribute values distributed with CelebA. Then we manually audit a sample of CelebA to estimate the accuracy of its attribute values.

3.1. Consistency of Manual Annotations?

To estimate the level of consistency that can be expected in manual annotations of commonly used face attributes, two annotators independently assigned values for each of 40 attributes of 1,000 images. The images were randomly selected from CelebA and should be representative of web-scraped, in-the-wild celebrity images. The annotators viewed the cropped, normalized face images, with no knowledge of the other annotator’s results.

Table 1 lists, for each attribute, the number of images for which the two annotations disagree. The least disagreement was for eyeglasses, at just 3 images. The 3



Figure 2. Example duplicate and near-duplicate pairs. The first three pairs have duplicate facial appearance, although the overall images have some differences from editing. The fourth (rightmost) pair is only a near-duplicate, with highly similar but different face appearance; e.g., mouth closed in the left image, open in right image.

Attribute	N_d	Attribute	N_d
High Cheekbones	512	Smiling	141
Pointy Nose	490	Wearing Lipstick	133
Oval Face	466	Blurry	118
Arched Eyebrows	406	5 o Clock Shadow	110
Narrow Eyes	394	Chubby	110
Attractive	377	Sideburns	105
Straight Hair	369	Blond Hair	101
Wavy Hair	299	Double Chin	58
Big Nose	286	Wearing Earrings	50
Bags Under Eyes	284	Wearing Necklace	49
Brown Hair	284	Gray Hair	37
Pale Skin	237	No Beard	34
Black Hair	232	Wearing Necktie	27
Rosy Cheeks	217	Male	23
Big Lips	206	Mustache	22
Heavy Makeup	181	Bald	18
Receding Hairline	178	MSO	17
Young	169	Goatee	16
Bushy Eyebrows	158	Wearing Hat	14
Bangs	150	Eyeglasses	3

Table 1. Number of differences (N_d) between manual annotations of 1,000 images. N_d grouped into three subjectivity levels, from pink (high) to green (low).

images present an “edge case” for the attribute definition. There are eyeglasses in the image, but positioned on the person’s head rather than in front of their eyes. One annotator marked these as eyeglasses=true and the other marked these as eyeglasses=false. Similarly, the 14 disagreements on the wearing_hat attribute were primarily instances of a literal hat versus a more general head covering (scarf, visor, ...). Inconsistencies arising from this type of ambiguity in the attribute definition should be reduced if annotators are provided a sufficiently detailed definition. However, it seems that the face attribute research community has so far relied in formal definitions being sufficient. At the other extreme, the two annotators had essentially random agreement on high cheekbones, pointy nose and oval face, disagreeing on 512, 490 and 466 of 1,000 images, respectively. Based on these results, we divide the 40 attributes into three ranges of inherent consistency. Attributes labeled

consistently for $\geq 95\%$ of the 1,000 images have *green* background in Table 1. Attributes with consistency $> 85\%$ and $< 95\%$ have *yellow* background and attributes with $\leq 85\%$ consistently have *pink* background. **The results show that the ambiguity of definition can bring a large inconsistency on annotations.**

3.2. Consistency Across Duplicate Faces?

Web-scraped datasets naturally contain duplicate and near-duplicate images. The same face image may appear on multiple websites with different brightness or other edits; these are duplicates. Also, images may be taken at slightly different times or from different points of view; these are near duplicates. To identify duplicate face appearances in CelebA, we start with image pairs from the same identity group that have ArcFace (ResNet101 backbone) similarity ≥ 0.9 , and manually inspect these pairs to discard near-duplicates. This resulted in 5,068 duplicate pairs from 3,094 identities. Additional duplicates could be found by extending to lower ArcFace similarity, but 5K+ pairs is more than sufficient for a useful estimate of the consistency of CelebA attribute values.

If the CelebA attribute values are assigned in a perfectly consistent manner, the attribute values for each duplicate image pair should be identical. An attribute whose values are assigned at random with a 50/50 true/false split is expected to have consistent values on 50% of the duplicate pairs. An attribute whose values are assigned at random with a 90/10 true/false split is expected to have consistent attribute values on 82% (81% + 1%) of the duplicate pairs. In general, the level of inconsistency observed in a particular attribute’s values across the duplicate pairs can be computed as:

$$\mathcal{P}_{in} = \frac{\mathcal{N}_{differ}}{(\mathcal{P}(n|p) + \mathcal{P}(p|n)) \times \mathcal{N}_{total}} \quad (1)$$

where \mathcal{N}_{total} is total number of duplicate image pairs, \mathcal{N}_{differ} is number of duplicate pairs with attribute values that differ, $\mathcal{P}(n|p) + \mathcal{P}(p|n)$ is the number of image pairs expected to agree if the values are assigned at random with the overall relative frequency. The estimated inconsistency level ranges from 0, representing perfectly

Attribute	\mathcal{N}_{differ}	N_n	N_p	\mathcal{P}_{in}	Attribute	\mathcal{N}_{differ}	N_n	N_p	\mathcal{P}_{in}
Blurry	154	9,836	300	0.529	5 o Clock Shadow	192	8,616	1,520	0.149
Pointy Nose	860	7,338	2,798	0.425	Mustache	79	9,566	570	0.147
Pale Skin	144	9,644	492	0.308	Black Hair	270	7,644	2,492	0.144
Rosy Cheeks	226	9,294	842	0.293	High Cheekbones	343	5,396	4,740	0.136
Oval Face	535	7,710	2,426	0.290	Sideburns	90	9,418	718	0.135
Narrow Eyes	280	8,936	1,200	0.265	Goatee	79	9,484	652	0.130
Straight Hair	447	7,920	2,216	0.258	Big Lips	275	7,092	3,044	0.129
Chubby	159	9,424	712	0.240	Heavy Makeup	249	6,128	4,008	0.103
Wearing Necklace	316	8,524	1,612	0.233	Young	190	2,640	7,496	0.097
Bags Under Eyes	454	7,456	2,680	0.230	Gray Hair	49	9,594	542	0.096
Brown Hair	376	8,038	2,098	0.226	Bald	25	9,830	306	0.084
Receding Hairline	201	9,140	996	0.224	Bangs	111	8,560	1,576	0.083
Double Chin	139	9,446	690	0.216	Blond Hair	112	8,530	1,606	0.083
Wearing Earrings	333	8,096	2,040	0.204	Smiling	196	5,072	5,064	0.077
Bushy Eyebrows	280	8,418	1,718	0.196	No Beard	119	1,920	8,216	0.077
Wavy Hair	441	6,700	3,436	0.194	MSO	178	5,214	4,922	0.070
Attractive	482	4,486	5,650	0.193	Wearing Lipstick	159	5,254	4,882	0.063
Arched Eyebrows	430	6,758	3,378	0.191	Wearing Hat	12	9,860	276	0.045
Big Nose	343	6,956	3,180	0.157	Eyeglasses	10	9,662	474	0.022
Wearing Necktie	140	9,110	1,026	0.152	Male	12	5,648	4,488	0.005

Table 2. Level of inconsistency in CelebA attributes based on analysis of duplicate face appearances.

consistent, to 1, representing a level of inconsistency indicating values assigned at random.

Table 2 summarizes, for each of the 40 attributes, the number of duplicate pairs whose attribute values differ, the number of total negative (attribute=false) samples, the number of total positive (attribute=true) samples, and the level of inconsistency computed using the above equation. The attributes are listed from high to low level of inconsistency. The highest level of inconsistency is for “blurry”, which is actually an image attribute rather than a face attribute. The same source image appearing online with, for example, two different levels of compression could appear different in blur. For this reason, we focus on the other 39 attributes for analyzing the consistency of attribute values.

There are 12 attributes with inconsistency less than 0.1, 14 with inconsistency between 0.1 and 0.2, and 14 with inconsistency above 0.2. The facial attribute pointy-nose has the highest inconsistency, at 0.425 and the attribute male has the lowest, at 0.005, for a difference of two orders of magnitude. *It is clear that the level of inherent inconsistency varies greatly across attributes for duplicate image pairs.*

3.3. Accuracy of CelebA Attribute Values?

Above experiments show that some attributes can be annotated with high consistency across annotators and across different versions of the same image. Restricting our attention to the 12 attributes with $\geq 95\%$ consistent

Attribute	N_n	N_p	Err_n	Err_p
Bald	972	500	0.41%	43.40%
Eyeglasses	925	500	0.43%	3.00%
Goatee	938	500	0.96%	57.40%
Gray Hair	953	500	3.78%	17.00%
Male	563	500	0.18%	4.20%
MSO	513	500	20.86%	1.60%
No Beard	500	837	39.20%	1.19%
Earrings*	793	500	8.07%	17.00%
Hat*	954	500	1.05%	5.80%
Necklace*	877	500	5.36%	55.80%
Necktie*	928	500	0.97%	30.00%
Mustache	959	500	7.09%	9.60%

Table 3. Error rate estimation of objective attributes. N is the number of random samples. n stands for negative in the original annotation, p stands for positive in the original annotation. * are the wearings.

manual annotations, we now estimate the error rate in the attribute values distributed with CelebA.

We increase the number of images annotated to a minimum of 500 with each original CelebA value for each of the 12 attributes. Two persons independently assign values for each image. For this experiment, attribute values that are initially inconsistent are considered again to arrive at a consensus correct value. The consensus values are used to mark the original CelebA

attribute values as correct or incorrect. Table 3 summarizes the error rate, broken out for true/false original CelebA value, and color-coded with green background for $\leq 5\%$ error and red for $> 15\%$ error.

Error rates for original CelebA attribute values vary greatly across attributes, and between true/false values for each attribute. For instance, the error rate for (bald=true) is over 43% and for (bald=false) it is only 0.41%, for over two orders of magnitude difference. High error rates and highly imbalanced error rates can present problems for machine learning algorithms.

Beyond plain error, attribute values in red can be categorized into three causes. One, there is increased ambiguity for facial hair attributes such as goatee and no_beard occur when hair length is short. Two, “wearing” attributes (necklace, necktie, earrings) may have errors if the original values were assigned from viewing full images whereas the cleaned values were assigned from viewing cropped images. For example, a necklace that is visible in the full image may not be visible in the cropped image. We are choosing for the “correct” value to be what is visible in the cropped face image. Three, errors for other attributes (gray_hair, MSO, bald) may arise from ambiguous definitions. (Examples of each type are shown in Figures 1, 2, and 3 of the Supplementary Material). The cause of errors is speculative, since, again, we do not know how the original CelebA attribute values were assigned. In general, even for attributes that in principle can have consistent annotations, the error rate in the original CelebA annotations seems subjectively “high” for use as training data for machine learning models.

4. Models From Cleaned Versus Original

In principle, higher quality attribute values should enable learning a better model. To investigate this, we select one attribute for cleaning/correction and compare the quality of the models learned from corrected versus original values. We select the MSO attribute because the original values are relatively evenly distributed between true/false, its estimated error rates for true/false are representative of the asymmetric error rate problem, and corrected values can be assigned with high consistency.

4.1. Audited Data-Clean Workflow

To reduce the error rate of the 0.2M MSO values, we propose an efficient quality-audit workflow, shown in Figure 3. The workflow starts with a small subset of manually cleaned data, so that values are consistent and with low error. Then, we randomly select k subsets from this cleaned subset, and use each subset to train a classifier (i.e. k trained classifiers in total). The

not-yet-cleaned data is classified by each of the k classifiers, where the results are used to calculate the agreement level. Based on agreement of the k classifiers, the not-yet-cleaned data is divided into $k + 1$ subsets. The error rates in the subsets with high agreement between classifiers are estimated by a manual audit, and if it is below a target error rate, the subset is kept as passing the audit. Subsets that have a low level of agreement among classifiers, but are small enough, are manually labeled. Subsets that are too large to be manually labeled and whose estimated error rate exceeds the target are moved on to the next round. The workflow ends when the overall estimated error is below the target rate.

Table 3 estimates the MSO error rate as 20.9% for negative (false) values and 1.6% for positive. A target error rate of below 5% suggests that cleaning should focus on the images with original false values. For the initial subset of cleaned data, we manually labeled the images in the original test set, leaving the original train and validation sets (90% of the data) untouched. *Consequently, we manually cleaned 60,858 (30%) of images from the original dataset.*

4.2. Definition of MSO

Figure 4 shows problems encountered in a detailed definition of MSO. A small number of images do not contain a visible human face, and so are dropped. A slightly larger but still small number of images do not contain visible information to assign a value for some attributes, and so we introduce a third possible attribute value, *information not visible*. For example, if a person is holding a microphone in front of their face, the information may not be visible to assign attribute values for mouth, lips or nose.

To assign MSO values consistently, we make a detailed definition, including some edge cases. For example, in Figure 4c, there is something - e.g. glass, cigarette, straw - between the lips but the other part of lips is touching. For our definition, these edge cases are defined as part of (MSO=true). However, it is inevitable that there are still edge cases where it is uncertain if the lips are touching or not (see Figure 4d). These cases should dominate the images with label “errors” in the audited/cleaned attributes.

Table 4 summarizes the original and cleaned dataset. The cleaned dataset has 166 images dropped as unusable for facial attribute classification, and an additional 797 images with (MSO=info_not_visible). Note that most of the images in these two categories are not in the test set. A substantial fraction of the images have not been manually examined in our quality-audit workflow, so it is possible that more images that could be dropped and more attribute values could be (MSO=info_not_visible) if all images were manually reviewed. The MSO-cleaned

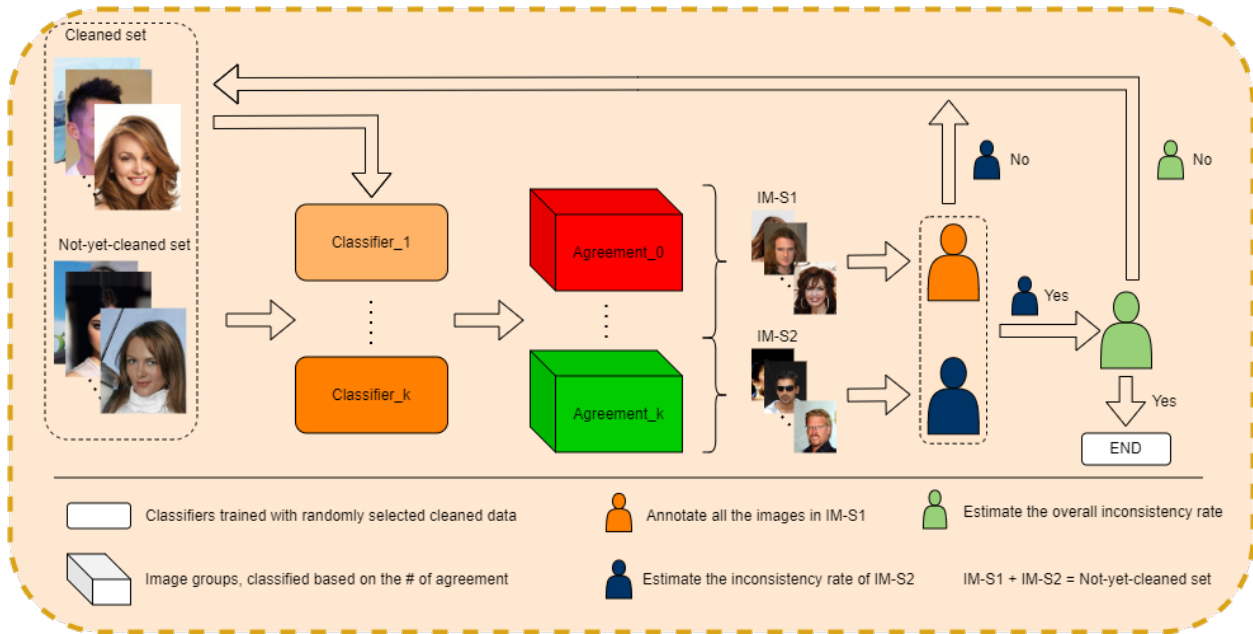


Figure 3. Quality-audit workflow for cleaning the annotations. Our annotation cleaning process is ended when the estimated inconsistency rate of all annotations is lower than 5%.




Figure 4. Examples of three annotations options, unusable images, ambiguous images, and edge images of MSO attribute (More examples are shown in Figures 4, 5, 6 from the supplemental material).

dataset has 161,982 training samples, 19,741 validation samples, and 19,913 testing samples. The number of positive and negative samples are not as balanced as in the original dataset.

4.3. Higher Quality, Higher Accuracy

To investigate the impact of learning models using the cleaned attribute values, we train four algorithms with different combinations of cleaned versus original

values. The four algorithms are AFFACT [8] (implementation obtained from original authors), MOON [20] (re-implemented), ResNet50 [9] (from Pytorch [19]) and DenseNet121 [11] (also PyTorch). We are experimenting with one attribute, MSO, so the last fully connected layer goes to one output. Input images are resized to 224×224 . MSE loss is used for training MOON, and binary cross-entropy loss for the other three. Each is trained for 50 epochs with 128 batch size. The learning

	Train	Val	Test	Info_not_vis	Unusable	NMSO	MSO
Original	162,770	19,867	19,962	—	—	104,657 (51.7%)	97,942 (48.3%)
Cleaned	161,982	19,741	19,913	797	166	73,704 (36.6%)	127,932 (63.4%)

Table 4. Comparison between original dataset and cleaned dataset on both annotation and image. Info_not_vis means the information of mouth is not visible in the image.

Train/Val/Test	AFF	MOON	RN50	DN121
Or/Or/Or	94.16	94.09	93.95	94.10
Or/Or/Cl	85.17	85.94	85.24	85.98
Or/Cl/Cl	86.17	86.49	86.54	85.69
Cl/Cl/Or	86.13	85.40	85.90	85.27
Cl/Cl/Cl	95.18	95.49	95.33	95.15

Table 5. Model accuracy on different combinations of original and cleaned MSO values used for train/validation/test set. Or = original, Cl = cleaned, AFF = AFFACT, RN = ResNet, DN = DenseNet.

rate is 0.00001 for MOON, and 0.0001 for the others. Horizontal flipping is used for data augmentation.

Table 5 summarizes the accuracies obtained by the four algorithms. For train/val/test with all original data, AFFACT obtains the highest accuracy, 94.16%. This is in keeping with what would be expected from the literature [8, 26]. For train/val/test with clean data, each algorithm obtains higher accuracy than with original data, *all four obtain an accuracy that exceeds the state-of-the-art obtained for this attribute with the original annotations*, and MOON obtains the highest, 95.49%.

The various combinations of mixed original and cleaned for train/val/test all suffer a dramatic drop in accuracy. This supports the premise that higher quality training data enables better models, and partial measures of cleaning only the test data, or only the validation and test, are counter-productive. If the collective values of the training data are not a quality approximation to the real-world concept, then the networks learn a model of some different concept. The starkness of the difference can be appreciated from the Score-CAM [7, 28] visualization of the average heat maps of the models. Figure 5 compares the average Score-CAM heatmap for MOON across the test set images for the model learned using our cleaned attribute values and the original attribute values. Face location is normalized in CelebA images, so the mouth appears on average in the same area of the image. The heatmap for the model learned using our cleaned values is clearly focused on the mouth region, as it should be to judge mouth open/closed. (The examples of heatmap comparison of individuals can be found in Figure 7 from the supplementary material). The heatmap for the model learned from original attribute values has no discernible focus on the mouth region.

This is visual evidence that the models learned from the cleaned and the original data are fundamentally different, and that the model learned from the cleaned data is better focused on the desired real-world concept.

5. Discussion and Conclusions

We seek to better understand the quality of models learned for facial attribute classification. We begin with an experiment to assess the consistency of human annotations for the set of 40 commonly-used binary facial attributes. We find that only 12 of the 40 are readily annotated with $\geq 95\%$ consistency between annotators. This result (a) highlights that most research datasets are created and distributed with no assessment of the consistency or reproducibility of the “ground truth labels”, (b) points to the need for more careful curation of meta-data, and (c) suggests that some face attributed may be too subjective to be useful.

Another experiment focuses specifically on the consistency of attribute values distributed with CelebA, the most widely used research dataset in this area. We identify 5,068 image pairs in CelebA with duplicate facial appearances, compare the attribute values across these pairs, and find that every attribute has contradicting values for some duplicate pairs, with the level of disagreement ranging up to 860 of the 5,068 pairs (for “pointed nose”).

The level of attribute value disagreement on CelebA duplicate face appearances motivates a more careful assessment of the correctness of the CelebA attribute values. This assessment focuses on the 12 attributes that we found can be annotated with $\geq 95\%$ consistency. Estimated error rates are highly asymmetric between the two values of an attribute; for example, the error rate for (MSO=false) is estimated at 20.9% and for (MSO=true) at 1.6%. Asymmetric error rates this high raise questions of what concept a model learns from this training data.

To investigate further, we create a cleaned version of the CelebA attribute values for MSO. The first step is manually-assigned correct MSO values for a small subset of CelebA. We use the initial subset of manually-assigned correct MSO values to bootstrap a hybrid automated/manual process that results in a cleaned version of the CelebA MSO values. (This cleaned version of the CelebA MSO values will be made available with the fi-




Figure 5. Average Score-CAM of the MOON created using original annotations and our cleaned annotations. The left side images of each pair is the average CAM generated by MOON trained with original dataset. The right side image of each pair is the average CAM generated by MOON trained with cleaned dataset.

nal paper.) The intended level of quality in our cleaned version of the CelebA MSO values is that a new independent manual labeling of the MSO attribute would agree at $\geq 95\%$ level with for both the positive and negative values of the attribute. As a side effect of creating the cleaned version of the CelebA MSO attribute, a small number of original CelebA images are marked as unusable in general due to not containing a human face image, and a small number are marked as “info not visible” due to the mouth being occluded in the image.

To assess the impact of our cleaned attribute values for learning a model, we compare accuracy of four algorithms trained using our cleaned versus the original attribute values. For all four algorithms, the model learned using cleaned values achieves higher accuracy than the state-of-the-art for models using the original data. Our implementation of the MOON algorithm achieves 95.49% accuracy using the cleaned dataset. Even more than the difference in accuracy levels achieved, comparing the Score-CAM heatmaps convincingly shows that the model learned using cleaned data is fundamentally better aligned with the attribute being classified.

Our results suggest that facial attribute classification research will benefit from datasets with higher quality meta-data. At the least, research datasets should come with a description of the level accuracy that can be expected in its attribute values. It may also be important, for multi-attribute classification, to allow for different subsets of attributes to be used from different images, in order to accommodate occlusion of different parts of the face in different images. The CelebA dataset has been foundational for research in face attribute classification, and now potentially spurs new directions.

6. Acknowledgement

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2022-21102100003. The views and conclusions contained herein are those of the authors

and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Olasimbo Ayodeji Arigbabu, Sharifah Mumtazah Syed Ahmad, Wan Azizun Wan Adnan, and Salman Yusoff. Recent advances in facial soft biometrics. *The Visual Computer*, 31:513–525, 2015. 2
- [2] Fabiola Becerra-Riera, Annette Morales-González, and Heydi Méndez-Vázquez. A survey on facial soft biometrics for video surveillance and forensic applications. *Artificial Intelligence*, 52:1155–1187, 2019. 2
- [3] Thomas Berg and Peter N Belhumeur. Poof: Part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 955–962, 2013. 1
- [4] Jui-Shan Chan, Gee-Sern Jison Hsu, Hung-Cheng Shie, and Yan-Xiang Chen. Face recognition by facial attribute assisted network. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3825–3829. IEEE, 2017. 1
- [5] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018. 1
- [6] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. 1
- [7] Jacob Gildenblat and contributors. Pytorch library for cam methods. <https://github.com/jacobgil/pytorch-grad-cam>, 2021. 7
- [8] Manuel Günther, Andras Rozsa, and Terranee E Boulton. Affact: Alignment-free facial attribute classification

- technique. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 90–99. IEEE, 2017. 1, 2, 6, 7
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 6
- [10] Zhenliang He, Wangmeng Zuo, Meina Kan, Shiguang Shan, and Xilin Chen. Attgan: Facial attribute editing by only changing what you want. *IEEE transactions on image processing*, 28(11):5464–5478, 2019. 1
- [11] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 2, 6
- [12] Neeraj Kumar, Alexander Berg, Peter N Belhumeur, and Shree Nayar. Describable visual attributes for face verification and image search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):1962–1977, 2011. 1
- [13] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar. Attribute and simile classifiers for face verification. In *2009 IEEE 12th international conference on computer vision*, pages 365–372. IEEE, 2009. 1
- [14] Daiqing Li, Junlin Yang, Karsten Kreis, Antonio Torralba, and Sanja Fidler. Semantic segmentation with generative models: Semi-supervised learning and strong out-of-domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8300–8311, 2021. 1
- [15] Yan Li, Ruiping Wang, Haomiao Liu, Huajie Jiang, Shiguang Shan, and Xilin Chen. Two birds, one stone: Jointly learning binary code for large-scale face image retrieval and attributes prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3819–3827, 2015. 1
- [16] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 1, 2
- [17] Ohil K Manyam, Neeraj Kumar, Peter Belhumeur, and David Kriegman. Two faces are better than one: Face recognition in group photographs. In *2011 International Joint Conference on Biometrics (IJCB)*, pages 1–8. IEEE, 2011. 1
- [18] Hung M Nguyen, Ngoc Q Ly, and Trang TT Phung. Large-scale face image retrieval system at attribute level based on facial attribute ontology and deep neuron network. In *Asian conference on intelligent information and database systems*, pages 539–549. Springer, 2018. 1
- [19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [20] Ethan M Rudd, Manuel Günther, and Terrance E Boulton. A mixed objective optimization network for the recognition of facial attributes. In *European Conference on Computer Vision*, pages 19–35. Springer, 2016. 2, 6
- [21] Zhiyuan Shi, Timothy M Hospedales, and Tao Xiang. Transferring a semantic representation for person re-identification and search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4184–4193, 2015. 1
- [22] Fengyi Song, Xiaoyang Tan, and Songcan Chen. Exploiting relationship between attributes for improved face verification. *Computer Vision and Image Understanding*, 122:143–154, 2014. 1
- [23] Chi Su, Fan Yang, Shiliang Zhang, Qi Tian, Larry Steven Davis, and Wen Gao. Multi-task learning with low rank attribute embedding for multi-camera person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1167–1181, 2017. 1
- [24] Chi Su, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. Deep attributes driven multi-camera person re-identification. In *European conference on computer vision*, pages 475–491. Springer, 2016. 1
- [25] Philipp Terhörst, Daniel Fährmann, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Maad-face: a massively annotated attribute dataset for face images. *IEEE Transactions on Information Forensics and Security*, 16:3942–3957, 2021. 2
- [26] Nathan Thom and Emily M Hand. Facial attribute recognition: A survey. *Computer Vision: A Reference Guide*, pages 1–13, 2020. 1, 2, 7
- [27] Viktor Varkarakis and Peter Corcoran. Dataset cleaning—a cross validation methodology for large facial datasets using face recognition. In *2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–6. IEEE, 2020. 1, 2
- [28] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 24–25, 2020. 7
- [29] Haiyu Wu, Grace Bezdold, Aman Bhatta, and Kevin W Bowyer. Logical consistency and greater descriptive power for facial hair attribute learning. *CVPR*, 2023. 1
- [30] Ying Xu, Philipp Terhörst, Kiran Raja, and Marius Pedersen. A comprehensive analysis of ai biases in deep-fake detection with massively annotated databases. *arXiv preprint arXiv:2208.05845*, 2022. 1
- [31] Xin Zheng, Yanqing Guo, Huaibo Huang, Yi Li, and Ran He. A survey of deep facial attribute analysis. *International Journal of Computer Vision*, 128:2002–2034, 2020. 2