

# InclusiveFaceNet: Improving Face Attribute Detection with Race and Gender Diversity

Hee Jung Ryu<sup>1</sup> Hartwig Adam<sup>\*1</sup> Margaret Mitchell<sup>\*1</sup>

## Abstract

We demonstrate an approach to face attribute detection that retains or improves attribute detection accuracy across gender and race subgroups by learning demographic information prior to learning the attribute detection task. The system, which we call InclusiveFaceNet, detects face attributes by transferring race and gender representations learned from a held-out dataset of public race and gender identities. Leveraging learned demographic representations while withholding demographic inference from the downstream face attribute detection task preserves potential users’ demographic privacy while resulting in some of the best reported numbers to date on attribute detection in the Faces of the World and CelebA datasets.

## 1. Introduction

Detecting and recognizing different face attributes has become an increasingly feasible machine learning task due to the rise in deep learning techniques and large people-focused datasets (Niu et al., 2016; Liu et al., 2015). However, this rise has come at a cost: As public technology incorporates modern computer vision techniques, we are observing a troubling gap between how well some demographics are recognized compared with others (EOP, 2016). For example, a female may be classified as a male if she has short hair (Rodden, 2017), and a black face may not be recognized unless wearing a white mask (Buolamwini, 2016).

In this paper, we examine the problem of unequal performance across different race and gender subgroups in the context of face attribute detection. We focus on **inclusion**, which we use in this paper to mean improving accuracy across minority subgroups; and **demographic privacy**, where our approach to embracing coarse-grained demographic characteristics that affect the visual characteristics

<sup>\*</sup>Equally contributing co-supervisors. <sup>1</sup>Google, USA. Correspondence to: Hee Jung Ryu <imiheej@gmail.com>, Margaret Mitchell <mmitchellai@google.com>.




Figure 1. End-to-end Architecture of InclusiveFaceNet. Layers from the face recognition model are transferred to our diversity classifier and multi-head face attribute detector InclusiveFaceNet. Yellow denotes the decoupled race layer. Pink denotes layers trained from scratch for face attribute detection.

of face attributes does not require demographic inference on users of the downstream face attribute task. Inclusion is related to but not identical to notions of fairness, where the goal is to obtain equal performance across all subgroups (even if that means lower accuracy). Our approach is similar in spirit to the recently introduced proposal for *decoupled classifiers* (Dwork et al., 2018), where the learning of sensitive attributes can be separated from a downstream task in order to maximize both fairness and accuracy. As in (Dwork et al., 2018), we use transfer learning. While (Dwork et al., 2018) uses its sub-type, domain adaptation, we use its different sub-type, *task transfer learning*. In our setting, both gender and race are learned as independent, coupled classifiers, and are not inferred at run time.

InclusiveFaceNet matches or improves over a baseline without transfer learning, defining a new state-of-the-art performance for face attribute detection across gender boundaries in both the CelebA (Liu et al., 2015) and FotW (Escalera et al., 2016) datasets.

## 2. Ethical Considerations

**Intent.** The intent of this work is to demonstrate the utility of reasoning about demographics – in this case, race and gender – in an ethical and legal setting.

**Race and Gender categories.** Race and gender categories

are imperfect for many reasons. For example, there is no gold-standard for ‘race’ categories, and it is unclear how many race and gender categories should be stipulated (or whether they should be treated as discrete categories at all). Throughout our experiments and in this paper, we obscure race and gender names in order to avoid introducing unconscious societal bias associated with existing labels, and anticipate more nuanced treatment of the sensitive categories in future work.

Race and gender identities in this work are learned from a held-out set of public identifications, and the learned representations are then transferred for learning the face attribute models. This permits face attribute detection to leverage race and gender representation without ever having to uniquely predict these characteristics on another individual.

### 3. Background & Related Work

Research in computer vision has included work on issues that have direct social impact, such as security and privacy, however, research on the related issue of diversity and inclusion in vision is surprisingly lacking (but see, e.g., (Buolamwini, 2016)). There are several other strands of research background relevant to the work in this paper: face attribute detection and machine learning fairness.

**Face Attribute Detection.** The ChaLearn “Looking at People” challenge from 2016 (Escalera et al., 2016) provides the Faces of the World (FotW) dataset, which annotates gender and the presence of smiling on faces. (Zhang et al., 2016) won first place in this challenge, utilizing multi-task learning (MTL) and fine-tuning on top of a model trained for face recognition (Parkhi et al., 2015). (Rajeev Ranjan & Chellappa, 2017) later published an out-performing result for the same task on FotW utilizing MTL and transfer learning from a face recognition model (Sankaranarayanan et al., 2016). (Wang et al., 2016) defined state-of-the-art performance on CelebA face attributes (Liu et al., 2015), employing both MTL and transfer-learning from a face recognition model based on FaceNet (Schroff et al., 2015). This work is similar to the approach taken here, but includes additional processing and data, e.g., geo-location.

(Escalera et al., 2016) points out the importance of demographic diversity of the source material in a vision dataset, and postulates that analysis made on skewed datasets cannot be representative enough for benchmarking progress. The authors provide Faces of the World (FotW), collected with the aim to achieve a uniform distribution across two genders and four ethnic groups, and we use this dataset in the current paper (but note that we still find demographic skew).

**Fairness in Machine Learning.** (Dwork et al., 2012) demonstrates the importance of “Fairness through Awareness”, understanding sensitive characteristics like gender and race in order to build demographically inclusive models.

Proposals for fairness have included *parity*, such as Demographic Parity (Dwork et al., 2012; Hardt et al., 2016; Beutel et al., 2017), and equality measures such as that of (Hardt et al., 2016) Equality of Odds, which requires equal false negative rates and false positive rates across subgroups.

Very recent work from (Dwork et al., 2018) discusses theory related to the approach here, where the authors demonstrate how to define the most accurate classifier as the classifier that simultaneously minimizes false positives and false negatives. They introduce the notion of *decoupled classifiers*, in which a separate classifier is trained on each sensitive subgroup, and positive instances of one sensitive subgroup are paired against negative instances across all subgroups. The approach in this paper is similar in that sensitive subgroups are learned separately and transferred, but the classifiers in this work are coupled and learned independently. We anticipate that a decoupled approach would lead to further gains in future work.

### 4. Learning Diversity

**Hypothesis.** At the core of this work lies the idea that faces look different across different races and genders (S et al., 2014). If such a relationship holds, then face attribute detection could be improved by learning about race and gender characteristics.

**Evaluation Metrics.** In addition to overall accuracy, we also provide accuracy per demographic subgroup to examine the effect of our experiments on each race and gender subgroup. Inspired by recent work in machine learning fairness (see Section 3), we also evaluate with a metric that averages the false positive rate and false negative rate, Average False Rate (AFR), which is robust to label and subgroup imbalances in the test data.

#### 4.1. Twofold Transfer Learning

Transfer learning is useful in the case where there is a feature space  $\mathcal{X}$  and a marginal distribution over the feature space  $P(X_S)$ , where  $X_S = x_1, x_2 \dots x_n \in \mathcal{X}$ , as well as a marginal distribution over the related feature space  $P(X_D)$ , where  $X_D \in \mathcal{X}$ . Here,  $X_S$  and  $X_D$  are completely disjoint sets, respectively representing face images for face attribute detection and face images for race classification. Formally, given:

- A face recognition domain  $F_X$
- A face domain  $F_R$  for the task of race classification  $T_R$  with labels  $Y_R$  and images  $X_R$
- The target face domain  $F_S$  for face attribute detection  $T_S$  with labels  $Y_S$  and images  $X_S$

We optimize for  $T_R$  by learning  $P(Y_R|X_R)$  via transfer learning from a network trained in the face recognition space  $F_X$  with the pre-trained network frozen, and then freeze the learned representations from  $F_X$  and  $F_R$  in the target domain  $F_S$  when learning  $P(Y_S|X_S)$ . We refer to the use

	Total	Estimated Race					Gender			
		S1	S2	S3	S4	Other	G1	G2	Other	
FOTW	Train	6171	8%	48%	11%	15%	19%	48%	54%	2%
	Valid	3086	8%	53%	11%	17%	11%	55%	44%	1%
CELEBA	Train	162687	7%	75%	7%	8%	4%	42%	58%	-
	Valid	19863	7%	76%	6%	6%	6%	42%	58%	-
	Test	19955	9%	69%	9%	9%	6%	39%	61%	-

Table 1. Estimated Race & Groundtruth Gender Distribution across Datasets. Duplicated images in CelebA are counted only once in this table.

of transfer learning twice in this way as **twofold transfer learning**. The second fold of the transfer learning permits the model to leverage demographic representations without inferring demographic characteristics when deployed.

### 4.2. Learning Demographic Diversity

Given face crops similar to those from Picasa, we train a FaceNet model (Schroff et al., 2015) and extract features from the layer called *avgpool*. We add fully connected layers and then fine-tune the model for the race and gender classification tasks. Race and gender are learned from a held-out dataset with a uniform distribution across race and gender intersections. Our gender model’s performance on the publicly available FotW and CelebA are 93.87% and 99%, respectively, setting the new state-of-the-art gender classification accuracy.

**Race Classification.** We explore four race subgroups for which we were able to scrape over 100,000 web images of famous identities from (Wikidata; Rothe et al., 2016; Somandepalli, 2017), and train to 98% or greater AUC across all subgroups. We do not name the race subgroups here, as our intention is to not to label individuals, but to understand whether learning demographic categories can provide further gains across demographic categories (see Section 2).

**InclusiveFaceNet Architecture.** We train a multihead face attribute detector which we call InclusiveFaceNet. Figure 1 provides a high-level overview of the architecture used for learning to detect face attributes, denoted with AttrA-\* and AttrB-\*.

## 5. Face Attribute Detection Datasets

**Faces of the World (FotW).** FotW (Escalera et al., 2016) is originally introduced by the ChaLearn Looking at People<sup>1</sup> Smile and Gender Challenge (Escalera et al., 2016) with a goal of demographic diversity. In the dataset, each image displays a single face labeled with smile, and some with gender. See Table 1 for details. We estimate Subgroup 2 to make up half the data, with the rest as minorities.

<sup>1</sup><http://chalearnlap.cvc.uab.es/challenge/13/track/20/description/>

**Large-scale CelebFaces Attributes (CelebA).** CelebA (Liu et al., 2015) has each image labeled with presence of 40 face attributes (see Table 5), including smiling/not-smiling, and male/not-male. We estimate Subgroup 2 to be the majority, with relatively little data for other race subgroups. See Table 1 for details.

## 6. Experiments

### 6.1. Analyzing Inclusion

Table 1 details the racial inclusion in FotW and CelebA. For FotW, the race model estimates a high variance among the subgroups, with the plurality of the data for Subgroup 2. Subgroup 1 is the smallest of the subgroups, with only 523 training data instances. For CelebA, there is less variance across races, with Subgroup 2 the majority.

### 6.2. Evaluating

Results on Faces of the World are in Table 2. We achieve an accuracy higher than that of the prior art (Rajeev Ranjan & Chellappa, 2017), using both race and gender SMILING+R+G, and set a new state-of-the-art performance of **90.96%** with the model that includes a learned race representation, SMILING+R.

On FotW, our smile model SMILING+R+G reaches an accuracy 1.62% (absolute) higher than the 2016 ChaLearn Looking at People Challenge winner SIAT\_MMLAB (Zhang et al., 2016), and 0.13% (absolute) higher than the published state-of-the-art UMIACS (Rajeev Ranjan & Chellappa, 2017) on the FotW validation dataset. Our model outperforms prior arts (Li et al., 2016; Uřičář et al., 2016; Ehrlich et al., 2016) with an increased performance gap as large as 11.66% absolute. The only point in the evaluation where transferred knowledge *hurts* performance is from the addition of gender for Subgroup 4 in FotW, where accuracy is less than that of the baseline SMILING. The addition of race improves over the baseline.

The Faces of the World experiments suggest that learning race subgroups may be particularly useful for inclusive models. We next turn to evaluate the effect of learning race on all 40 face attributes found in the widely used dataset CelebA. We find that the proposed system sets new record accuracies across 23 face attributes, defining the state of the art on the 10 attributes in Table 5. Race representations also further improve accuracy for individual subgroups: For Gender 1, *bangs, big lips, blurry, chubby, double chin, narrow eyes, pale skin, pointy nose, wearing necklace, and young* improve with transfer learned race; For Gender 2, *attractive, bangs, big lips, big nose, black hair, bushy eyebrows, goatee, wearing necktie, and young* all improve with race representations (see Table 3). Improvements are marginal but consistent across the board.

## InclusiveFaceNet: Improving Face Attribute Detection with Race and Gender Diversity

Train: FotW Eval: FotW	Subgroup 1		Subgroup 2		Subgroup 3		Subgroup 4		Other		Gender 1		Gender 2		Total	
	Acc.%	AFR%	Acc.%	AFR%	Acc.%	AFR%	Acc.%	AFR%	Acc.%	AFR%	Acc.%	AFR%	Acc.%	AFR%	Acc.%	MAFR% R/G
SMILING	88.5	12.2	90.9	10.7	90.0	12.2	94.3	6.2	88.2	18.4	90.2	10.6	91.0	12.5	90.57	12.0 / 11.6
+G	88.5	<b>12.1</b>	91.2	10.1	<b>90.9</b>	<b>10.9</b>	94.1	6.1	88.2	18.4	90.3	10.4	<b>91.3</b>	<b>11.6</b>	90.76	<b>11.5</b> / 11.0
+R	88.5	<b>12.1</b>	<b>91.6</b>	<b>9.7</b>	90.6	11.3	<b>94.5</b>	<b>5.9</b>	88.2	18.4	<b>90.7</b>	<b>10.0</b>	<b>91.3</b>	11.8	<b>90.96</b>	<b>11.5</b> / <b>10.9</b>
+1FC(2048)	88.1	12.5	91.3	10.2	90.3	11.8	<b>94.5</b>	<b>5.9</b>	88.2	18.7	90.6	10.2	91.1	12.3	90.80	11.8 / 11.2

Table 2. Accuracy (ACC.), Average False Rate (AFR), and mean AFR (MAFR) by race and gender subgroup for smiling detection with proposed systems, on FotW. +G denotes transferred representations from the gender head of the diversity classifier, +R from the race head of the diversity classifier. The last row serves as a baseline: A model with the same number of parameters as the +G, +R, but without the transferred representations. We see improvements across minority subgroups, as well as the plurality subgroups, for both race and gender. The transferred knowledge does not impact the race ‘Other’ category, which is a racially diverse set of people whose representative racial features our diversity classifier is not trained to capture and extract. We exclude the handful of unlabeled FotW gender instances in reporting the Gender MAFR.

	Big Lips				Big Nose				Young			
	Gender 1		Gender 2		Gender 1		Gender 2		Gender 1		Gender 2	
	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR
VAN.	63.7	45.0	82.7	35.8	90.4	36.7	77.2	24.9	90.7	24.2	86.3	15.2
+R	<b>64.6</b>	<b>43.1</b>	<b>82.8</b>	<b>34.9</b>	90.4	<b>36.5</b>	<b>77.8</b>	<b>24.6</b>	<b>90.8</b>	<b>22.8</b>	<b>86.8</b>	<b>14.4</b>

Table 3. Example Improved Face Attributes in CelebA Evaluated by Gender. Examples of some of the most improved attributes, evaluated by gender. Accuracy (Acc.) and average false rate (AFR) are in percentage (%). VAN. refers to *vanilla*. +R refers to our model with race input. We see an improvement up to .9% absolute.

	Oval Face				Brown Hair				Wavy Hair			
	Gender 1		Gender 2		Gender 1		Gender 2		Gender 1		Gender 2	
	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR	Acc.	AFR
VAN.	<b>74.6</b>	<b>33.4</b>	<b>77.2</b>	<b>44.2</b>	<b>86.8</b>	<b>18.8</b>	<b>90.1</b>	<b>22.9</b>	<b>80.6</b>	<b>19.4</b>	<b>88.7</b>	<b>32.2</b>
+R	74.2	34.3	77.1	45.5	86.5	18.2	89.9	22.7	80.4	19.7	88.7	32.6

Table 4. Example Face Attributes with Loss in CelebA Evaluated by Gender. Examples of some of the attributes with the most loss, evaluated by gender. Accuracy (Acc.) and average false rate (AFR) are in percentage (%). VAN. refers to *vanilla*. +R refers to our model with race input. We see a loss up to .4% absolute.

	Black Hair	Blurry	Bushy Eyebrows	Chubby	Goatee	Male	Narrow Eyes	Pale Skin	Wear Necklace	Young
	LA	88	84	90	91	95	98	81	91	71
WL	84	91	93	89	92	96	79	85	77	86
VAN.	89.7	96.1	92.7	95.8	97.3	98.8	86.9	96.7	86.3	89.2
+R	<b>89.8</b>	<b>96.2</b>	<b>92.8</b>	<b>95.9</b>	<b>97.4</b>	<b>99.0</b>	<b>87.0</b>	<b>96.8</b>	<b>86.4</b>	<b>89.3</b>

Table 5. Example Improved Face Attributes in CelebA. Transfer learning with race representations matches or improves face attribute detection accuracy for 35 out of 40 attributes, with marginal but consistent gains improving 13 attributes over the baseline model, and setting new state-of-the-art on 10 face attributes (presented here). Prior art compared are LA=LNets+ANet (Zhang et al., 2013) and WL=Walk&Learn (Wang et al., 2016).

Among the 10 attributes with accuracy improvements with race input, *big lips* improved by as large as .9% absolute for Gender 1. On the other hand, five out of 40 attributes have accuracy losses with the addition of race input: *oval face*, *brown hair*, *heavy makeup*, *high cheekbones*, and *wavy hair*. *oval face* has the largest accuracy loss delta of .4% absolute for Gender 1. Qualitatively looking at randomly selected 20 example images of *oval face*, faces with light skin tone

are predicted correctly more often by the vanilla model than the +R model. Meanwhile, some of the examples that the vanilla model misses are correctly predicted by the +R model. Similar patterns are observed in randomly selected example images of *big lips* with accuracy improvements too.

## 7. Discussion

Building an inclusive model and building a state-of-the-art model can go hand-in-hand. This does not have to come at the expense of predicting characteristics that may be sensitive to predict at run-time.

Transferring race knowledge increases the accuracy in almost all cases in FotW and for 13 of the CelebA attributes. Notably, learning race improves performance for smiling across both genders in FotW. Transferring gender knowledge similarly leads to accuracy and AFR improvements, although it is not as impactful as race. In FotW, the addition of gender features increases the accuracy of the gender groups as well as race subgroups 2 and 3, in addition to improving accuracy overall.

Qualitatively, the addition of race and gender representations seem to improve smiling detection on neutral faces, faces with an open mouth, black-and-white photos, and in faces turned away from the camera. This work establishes the feasibility of using transfer learning with demographics to improve performance across demographic categories, which we refer to as algorithmic *inclusion*. Future work could examine the effects of joint training and decoupled classification.

## 8. Acknowledgements

Special thanks to Google, Blaise Agüera y Arcas and Corinna Cortes for their support and sponsorship. Thanks to Marco Andreetto, Timnit Gebru, Sergio Guadarrama, Moritz Hardt, Gautam Kamath, David Karam, Shrikanth Narayanan, Caroline Pantofaru, Florian Schroff, Krishna Somandepalli, Rahul Sukthankar and Jason Yosinski for helpful conversations. Thanks to Ben Hutchinson, Yongsun Lee, Divya Tyam and Andrew Zaldivar for donating their face photos for our research.

## References

- Beutel, Alex, Chen, Jilin, Zhao, Zhe, and Chi, Ed H. Data decisions and theoretical implications when adversarially learning fair representations. *CoRR*, abs/1707.00075, 2017. URL <http://arxiv.org/abs/1707.00075>.
- Buolamwini, Joy. How i'm fighting bias in algorithms. [https://www.ted.com/talks/joy\\_buolamwini\\_how\\_i\\_m\\_fighting\\_bias\\_in\\_algorithms#t-63664](https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms#t-63664), 2016.
- Dwork, Cynthia, Hardt, Moritz, Pitassi, Toniann, Reingold, Omer, and Zemel, Richard. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1115-1.
- Dwork, Cynthia, Immorlica, Nicole, Kalai, Adam Tauman, and Leiserson, Max. Decoupled classifiers for group-fair and efficient machine learning. *Proceedings of FAT\**, 2018.
- Ehrlich, Max, Shields, Timothy J., Almaev, Timur, and Amer, Mohamed R. Facial attributes classification using multi-task representation learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2016.
- EOP, Executive Office of the President. Big data: A report on algorithmic systems, opportunity, and civil rights. [https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016\\_0504\\_data\\_discrimination.pdf](https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf), 2016.
- Escalera, Sergio, Torres Torres, Mercedes, Martinez, Brais, Baro, Xavier, Jair Escalante, Hugo, Guyon, Isabelle, Tzimiropoulos, Georgios, Corneou, Ciprian, Oliu, Marc, Ali Bagheri, Mohammad, and Valstar, Michel. Chalearn looking at people and faces of the world: Face analysis workshop and challenge 2016. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2016.
- Hardt, Moritz, Price, Eric, , and Srebro, Nati. Equality of opportunity in supervised learning. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., 2016.
- Li, Chenghua, Kang, Qi, and Ge, Quojing. Deepbe: Learning deep binary encoding for multi-label classification. In *Computer Vision and Pattern Recognition Workshops CVPRW, 2016 IEEE Conference on*. IEEE, 2016. doi: 10.1109/CVPRW.2016.98.
- Liu, Ziwei, Luo, Ping, Wang, Xiaogang, and Tang, Xiaoou. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Niu, Zhenxing, Zhou, Mo, Wang, Le, Gao, Xinbo, and Hua, Gang. Ordinal regression with multiple output cnn for age estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Parkhi, Omkar M., Vedaldi, Andrea, and Zisserman, Andrew. Deep face recognition. In *British Machine Vision Conference*, 2015.
- Rajeev Ranjan, Swami Sankaranarayanan, Carlos D. Castillo and Chellappa, Rama. An all-in-one convolutional neural network for face analysis. In *2017 IEEE 12th International Conference on Automatic Face & Gesture Recognition*. IEEE, 2017.
- Rodden, Kerry. Is that a boy or a girl? exploring a neural networks construction of gender. <https://medium.com/@kerryrodden/is-that-a-boy-or-a-girl-cb93abbae6da>, 2017.
- Rothe, Rasmus, Timofte, Radu, and Gool, Luc Van. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision (IJCV)*, July 2016.
- S, Fu, H, He, and ZG, Hou. Learning race from face: A survey. *IEEE Trans Pattern Anal Mach Intell.*, 36, Dec 2014.
- Sankaranarayanan, S., Alavi, A., Castillo, C. D., and Chellappa, R. Triplet probabilistic embedding for face verification and clustering. In *2016 IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pp. 1–8, Sept 2016. doi: 10.1109/BTAS.2016.7791205.
- Schroff, Florian, Kalenichenko, Dmitry, and Philbin, James. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- Somandepalli, Krishna. Prediction race from face for movie data. <https://github.com/usc-sail/mica-race-from-face/wiki>, 2017.
- Uříčář, Michal, Timofte, Radu, Rothe, Rasmus, Matas, Jiří, and Gool, Luc Van. Structured output SVM prediction of apparent age, gender and smile from deep features. In *Proceedings of IEEE conference on Computer Vision and Pattern Recognition Workshops*, Las Vegas, USA, 2016.
- Wang, Jing, Cheng, Yu, and Feris, Rogerio Schmidt. Walk and learn: Facial attribute representation learning from egocentric video and contextual data. In *CVPR*, 2016.

Wikidata. Wikidata. <http://www.wikidata.org/>, 2018.

Zhang, Kaipeng, Tan, Lianzhi, Li, Zhifeng, and Qiao, Yu. Gender and smile classification using deep convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 34–38, 2016.

Zhang, Ning, Paluri, Manohar, Ranzato, Marc’Aurelio, Darrell, Trevor, and Bourdev, Lubomir D. PANDA: pose aligned networks for deep attribute modeling. *CoRR*, 2013.