

A Study of Face Obfuscation in IMAGENET

Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky



PRINCETON
UNIVERSITY




Stanford
University



Large-Scale Visual Datasets



[Deng et al. CVPR 2019]
[Russakovsky and Deng, IJCV 2015]



[Kuznetsova et al. IJCV 2020]



[Zhou et al. TPAMI 2018]



[Krishna et al. IJCV 2017]




[Lin et al. ECCV 2014]

Large-Scale Visual Datasets



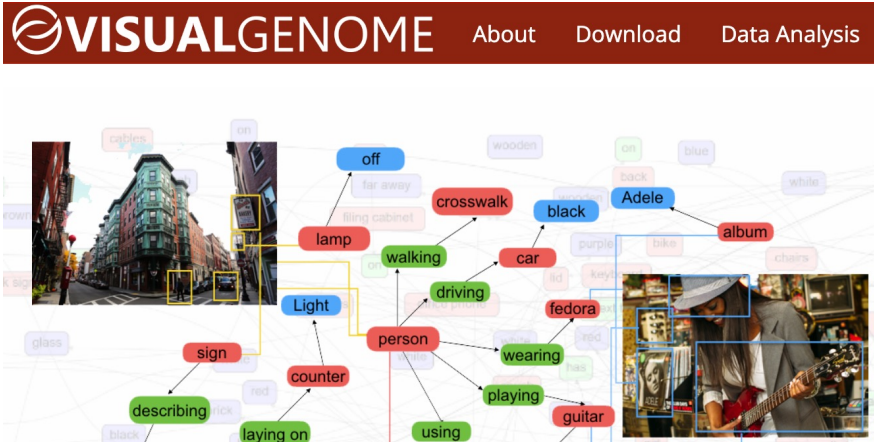
[Deng et al. CVPR 2019]
[Russakovsky and Deng, IJCV 2015]



[Kuznetsova et al. IJCV 2020]



[Zhou et al. TPAMI 2018]



[Krishna et al. IJCV 2017]



[Lin et al. ECCV 2014]


Privacy of Incidental People in Images

- Only 3 out of 1000 labels are people-centric: scuba diver, bridegroom, baseball player
- But many incidental people co-occur with other objects
- People's privacy is at risk since these datasets are freely available and widely used

[Prabhu and Birhane, WACV 2021]



Barber chair



Husky



Beer bottle



Volleyball



Military uniform


Privacy of Incidental People in Images

- Only 3 out of 1000 labels are people-centric: scuba diver, bridegroom, baseball player
- But many incidental people co-occur with other objects
- People's privacy is at risk since these datasets are freely available and widely used


[Prabhu and Birhane, WACV 2021]



Barber chair



Husky



Beer bottle



Volleyball




Military uniform

Contributions

- Annotate faces in ILSVRC, facilitating subsequent research on privacy protection
- Demonstrate that face obfuscation does not hurt large-scale visual recognition

Faces in the ImageNet Challenge Data


- Annotated face bounding boxes on **1,431,093** images
- **243,198** images have at least one face (**17%**)
- **562,626** faces in total



Face Obfuscation by Blurring and Overlaying

- Simple, widely used in practice
- Without any formal guarantee of privacy

[Chang et al. "People Identification with Limited Labels in Privacy-Protected Video" ICME 2006]
[Oh et al. "Faceless Person Recognition; Privacy Implications in Social Media" ECCV 2016]



Blurred




Overlaid

Face Obfuscation by Blurring and Overlaying


- Simple, widely used in practice
- Without any formal guarantee of privacy

[Chang et al. "People Identification with Limited Labels in Privacy-Protected Video" ICME 2006]

[Oh et al. "Faceless Person Recognition; Privacy Implications in Social Media" ECCV 2016]



Blurred



Overlaid

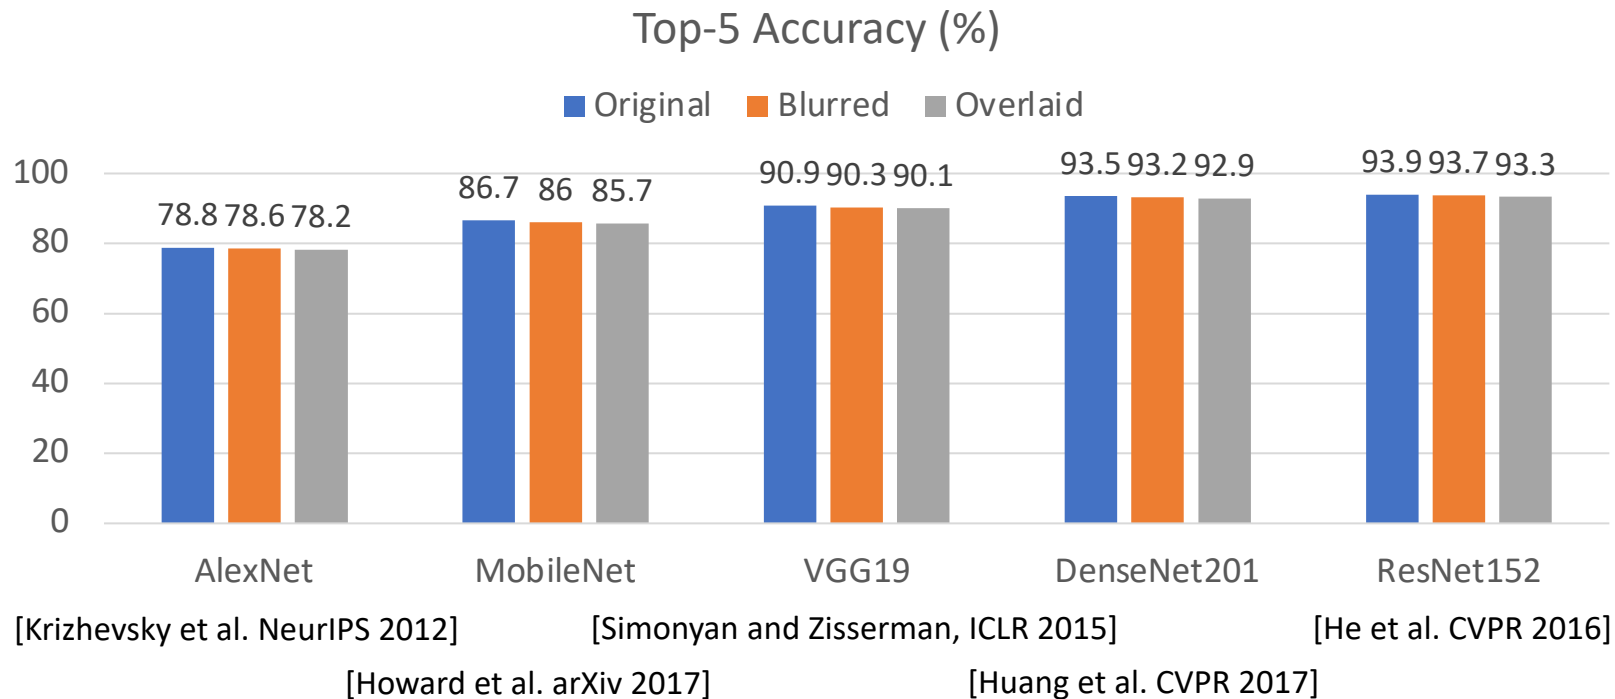
Does face obfuscation hurt ImageNet's utility?

Face Obfuscation Does Not Hurt Image Classification

- 15 convolutional model architectures: AlexNet, VGG, ResNet, etc.
- 3 settings: train and evaluate on **original/blurred/overlaid** images

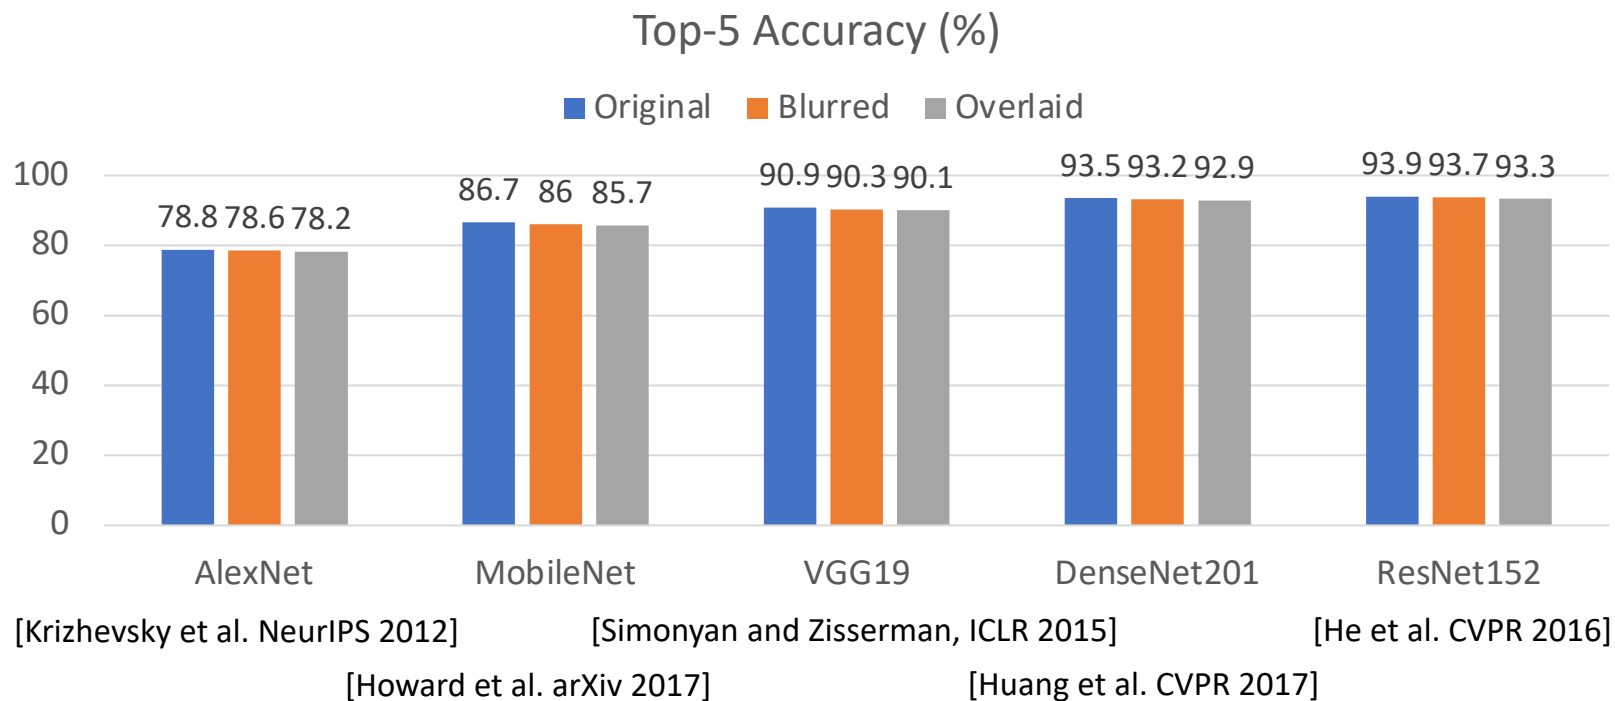
Face Obfuscation Does Not Hurt Image Classification

- 15 convolutional model architectures: AlexNet, VGG, ResNet, etc.
- 3 settings: train and evaluate on original/blurred/overlaid images



Face Obfuscation Does Not Hurt Image Classification

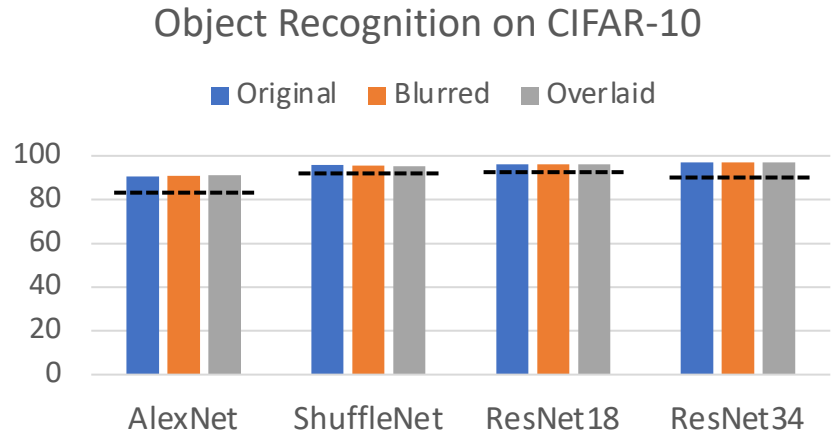
- 15 convolutional model architectures: AlexNet, VGG, ResNet, etc.
- 3 settings: train and evaluate on original/blurred/overlaid images
- **Marginal (< 1.0%) drop in top-5 accuracy**



Evaluating Face Obfuscation on Transfer Learning

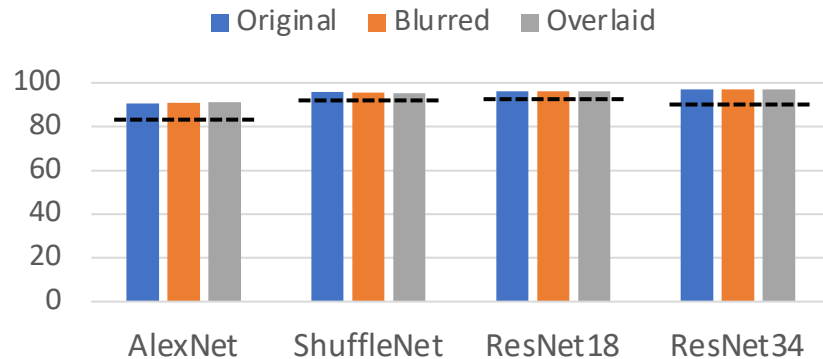
- 3 settings: Pretraining on original/blurred/overlaid ImageNet
- Finetune on 4 downstream tasks
 - Object recognition on CIFAR-10 [Krizhevsky and Hinton, 2009]
 - Scene recognition on SUN [Xiao et al. CVPR 2010]
 - Object detection on PASCAL VOC [Everingham et al. IJCV 2010]
 - Face attribute classification on CelebA [Liu et al. ICCV 2015]

Face Obfuscation Does Not Hurt Transfer Learning




Face Obfuscation Does Not Hurt Transfer Learning

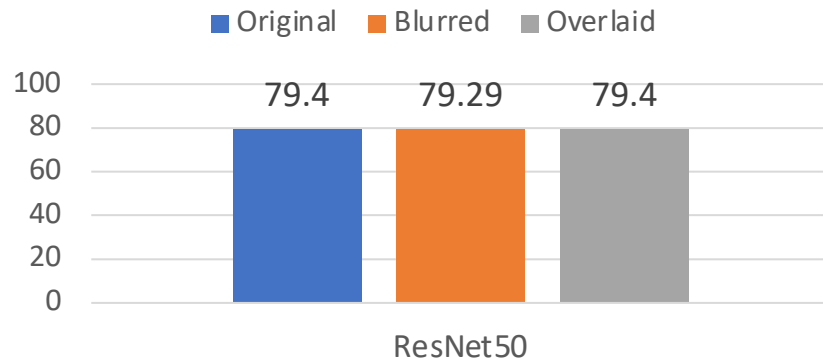
Object Recognition on CIFAR-10




Scene Recognition on SUN



Object Detection on PASCAL VOC



Face Attribute Classification on CelebA



No accuracy drop on transferring to downstream tasks

Key Takeaways

- Face Obfuscation in ImageNet
 - improves privacy
 - does NOT hurt image classification
 - does NOT hurt transfer learning

Related Work

- Obfuscating private information in datasets
 - Google Street View [Frome et al. ICCV 2009] [Uittenbogaard et al. CVPR 2019]
 - nuScenes [Caesar et al. CVPR 2020]
 - AViD. [Piergiovanni and Ryoo, NeurIPS 2020]
- Image datasets without people [Asano et al. NeurIPS 2021 Datasets and Benchmarks Track]
- Ethical concerns in constructing datasets

Hanley et al. Navigating the Broader Impacts of AI Research Workshop. “An Ethical Highlighter for People-Centric Dataset Creation”

Tomei et al. Workshop on Fair, Data Efficient and Trusted Computer Vision. “Estimating (and Fixing) the Effect of Face Obfuscation in Video Recognition”

Hanley et al. Navigating the Broader Impacts of AI Research Workshop. “An Ethical Highlighter for People-Centric Dataset Creation”

Tomei et al. Workshop on Fair, Data Efficient and Trusted Computer Vision. “Estimating (and Fixing) the Effect of Face Obfuscation in Video Recognition”

Birhane et al. arXiv 2021. “Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes”

Birhane et al. arXiv 2021. “The Values Encoded in Machine Learning Research”

Paullada et al. ML-RSA Workshop 2020. “Data and its (Dis)contents: A Survey of Dataset Development and Use in Machine Learning Research”

Scheuerman et al. CSCW 2021. “Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development”

Denton et al. Big Data & Society. “On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet”

Milagros et al. FAccT 2021. “Documenting Computer Vision Datasets: An Invitation to Reflexive Data Practices”

Hutchinson et al. FaccT 2021. “Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure”

A Study of Face Obfuscation in IMAGENET

Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky



<https://github.com/princetonvisualai/imagenet-face-obfuscation>

