[RE] CNN-generated images are surprisingly easy to spot... for now

Anonymous Author(s)
Affiliation
Address
email

Reproducibility Summary

This work evaluates the reproducibility of the paper "CNN-generated images are surprisingly easy to spot... for now" by Wang et al. published at CVPR 2020. The paper addresses the challenge of detecting CNN-generated imagery, which has reached the potential to even fool humans. The authors propose two methods which help an image classifier to generalize from being trained on one specific CNN to detecting imagery produced by unseen architectures, training methods, or data sets.

Scope of Reproducibility

The paper proposes two methods to help a classifier generalize: (i) using a diverse data set and (ii) utilizing different kinds of data augmentations. This report focuses on assessing if these techniques indeed help the generalization process and if the results extend beyond the original contribution.

Methodology

We decided to implement the methods from scratch to highlight possible hurdles for practitioners who want to adapt the results. For our experiments, we utilized the training and test data provided by the authors, as well as, creating our own data set to analyze how the proposed techniques perform on different data sets. Note that overall we estimate the entire experiments to take around 380 GPU hours.

Results

In general, we were able to replicate the results reported in the paper but we also obtained results different from the paper. After further investigations, we discovered that we had misunderstood one of the proposed data augmentation methods based on the description in the paper. This misunderstanding lead to a different implementation and the discrepancies. However, the overall trends of our experiments matched the original publication, i.e., data diversity and augmentations help generalization. Thus, we believe both techniques to be effective. We also performed several additional experiments on other data sets, highlighting limitations and clarifying the methods.

What was Easy

The paper is very detailed and we could faithfully replicate the experiments (bar the aforementioned exception). We only utilized the authors' code sparingly, but the repository is very well documented and provides pre-trained models.

What was Difficult

The naming of one of the augmentations caused confusion on our part. Note that the description in the paper is correct, however, the naming suggest the method works different than actually implemented. This resulted in us implementing a variation of the proposed technique. Surprisingly, our method actually improved on the original results.

Communication with Original Authors

Most of our question could be resolved by comparing our implementation against the authors' code. We contacted the first author regarding our different results, he was very responsive and answered all of our questions.

Submitted to ML Reproducibility Challenge 2020. Do not distribute.

1 Introduction

High-quality implementations of state-of-the-art image synthesis methods are available to the general public [Karras et al., 2020]. This technology is astonishing and concerning at the same time. One the one hand, such algorithms can be used for practical use cases like image-to-image translation [Isola et al., 2017, Liu et al., 2017, Zhu et al., 2017] or image super-resolution [Lai et al., 2017, Ledig et al., 2017, Lim et al., 2017]. On the other hand, they can generate images which even trick human observers [Simonite, 2019]. Recently, fake news and tampered media have become practical problems which even have the potential to influence democratic processes [Thompson and Lapowsky, 2017, Hao, 2019].

Recognizing these problems, Wang et al. [2020] propose a "universal" detector for CNN-generated images. While creating a detector for know architectures has been done before [Zhang et al., 2019, Wang et al., 2019a], the authors of this paper ask whether it is conceivable to train a classifier which generalizes to new, unseen architectures. The authors demonstrated that with a diverse data set and with carefully chosen data augmentation methods, a standard image classifier can achieve this goal. More specifically, they train a ResNet50 [He et al., 2016] classifier on a data set consisting of images generated by 20 different instances of ProGAN [Karras et al., 2018]. Additionally, they augment the training process by utilizing aggressive image filtering techniques (i.e., Gaussian Blur and JPEG compression). Combining both techniques allows them to create an image classifier which even generalized to StyleGAN2 [Karras et al., 2020], which was published concurrently.

Scope of reproducibility: We opted to reimplement the methods from scratch, aiming at highlighting difficulties for practitioners who want to adapt these methods. While the authors released a very well documented code repository, we initially abstained from inspecting the source code to not prime ourselves. In the later stages of our analysis, we used their repository to uncover some small implementation differences, which lead to our implementation actually improving on the original results.

Using our implementation, we aim at verifying the two main claims of the paper:

1. More diverse data sets help the classifier generalize to unseen architectures.
2. The proposed data augmentations (usually) improve generalization.

We recreate parts of the original experiments which directly investigate these two claims (see Section 4.2, Section 4.3, Table 2, and the Figures 2 and 3 in the original paper [Wang et al., 2020]). Departing from their work, we perform several additional experiments: First, we further investigate the first claim, training multiple additional data set combinations to clarify differences between our results and the original paper. Second, we generate a new training set from a different CNN generator (StyleGAN2 [Karras et al., 2020]) and evaluate if the results transfer. Finally, we investigate if the method is exclusive to the specific image classifier used in the paper. To this end, we train two other classifiers and compared them to the original results: A different image classifier architecture [Simonyan and Zisserman, 2015] and a classifier which leverages spectral analysis [Frank et al., 2020].

2 Methodology

In this section, we provide an overview of our methodology. We start by presenting the different data sets used in our experiments. Then, we describe the overall experimental setup, including the models we studied, data augmentations performed, details on the training procedure, our evaluation criteria, and computational requirements. All of our reproducibility efforts are based on the arXiv version of the paper (v2 – submitted 04. April 2020) and the corresponding GitHub repository[1] (commit f692c13 – 26. October 2020). We also make our implementation, our pre-trained models and our data set publicly available [2].

2.1 Data Sets

For training our networks, we used two different data sets: The first one is provided by the original publication. The authors used 20 ProGAN [Karras et al., 2018] models, each trained on a different LSUN category [Yu et al., 2015], to sample 18K generated images per model. They also collected 18K training images for each of the corresponding classes, resulting in a total data set of $(18.000 + 18.000) * 20 = 720.000$ images.

Additionally, we created a new data set to study the proposed techniques on a different base generator. We downloaded the pre-trained models for StyleGAN2 [Karras et al., 2020]. Using these models, we generated a training set of 36K

[1] available at `https://peterwang512.github.io/CNNDetection/`

[2] available at `https://github.com/mlreprochallenge/CNNEasyToSpot`

train (18K fake and 18K real) images for the LSUN *cat and horse* data sets, following the pre-/post-processing steps outlined in the original publication. Since these categories are also contained in the data set by Wang et al., this allows for a one-to-one comparison.

We exclusively used the test data set provided by the authors. The data set consists of images from 11 different synthesis models:

- GANs: The data set contains samples from six different Generative Adversarial Networks (GANs): ProGAN [Karras et al., 2018], StyleGAN [Karras et al., 2019], BigGAN Brock et al. [2019], GauGAN [Park et al., 2019], CycleGAN [Zhu et al., 2017], and StarGAN [Choi et al., 2018].
- Perceptual loss: The authors also include samples from generative models which are directly trained to optimize a perceptual loss: Cascaded Refinement Networks (CRN) [Chen and Koltun, 2017] and Implicit Maximum Likelihood Estimation (IMLE) [Li et al., 2019].
- Image manipulation models: It includes samples from Seeing In The Dark (SITD) [Chen et al., 2018], which approximates long-exposure photography, and the Second-Order Attention Network (SAN) [Dai et al., 2019], which generates images at a higher resolution.
- Deep fakes: Finally, the authors include samples from the FaceForensics++ [Rössler et al., 2019] data set.

2.2 Experiment Setup

In the following, we provide an overview of the different experiments performed in the paper. We start by discussing the different models used for both reproducing the original results and further experiments. Then we present the data augmentations proposed by Wang et al. to help the model generalize. We proceed by presenting the training schedule used in all experiments and discuss how our evaluation is performed. Finally, we conclude with an overview of the hardware details and the computational requirements to reproduce this report.

Models used: Following the original paper, we use a ResNet50 [He et al., 2016] model as the image classifier for the majority of our experiments. The classifier is pre-trained on ImageNet [Russakovsky et al., 2015]. We use random crops (to 224×224 pixels) and horizontal flipping during training of all of our classifiers. Additionally, we scale the image to the range $[0, 1]$, remove the mean, and scale them to unit variance.

For the experiments presented in Section 3.3.2, we also train a VGG-11 [Simonyan and Zisserman, 2015] model and a classifier which utilizes spectral information [Frank et al., 2020]. The VGG model is also pre-trained on ImageNet and uses the same preprocessing as our ResNet models. The spectral classifier also uses random crops and horizontal flipping as outlined above. Afterwards, we follow Frank et al. and use Discrete Cosine Transfrom-II [Ahmed et al., 1974] (DCT-II) to transform the images from the spatial to the spectral domain. Afterwards, we use MinMax-scaling to scale the images to the range $[0, 1]$ and train a ResNet50 model from scratch on the transformed data.

Data augmentations: Wang et al. proposed several data augmentations which help their classifier generalize to unseen architectures:

- *Gaussian blur (Blur):* Before cropping, with 50% probability, the images are blurred with $\sigma \sim \text{Uniform}\{0, 3\}$
- *JPEG-compressiong (JPEG):* with 50% probability images are JPEF-ed by two popular image libraries, OpenCV and PIL, with quality $\sim \text{Uniform}\{30, \ldots, 100\}$.
- *Blur + JPEG (Blur & JPEG (0.5)):* with 50% probability images are possibly blurred and JPEG-ed, each with 50% probability.
- *Blur + JPEG (Blur & JPEG (0.1)):* similar to the previous, but the probability is dropped to 10%.

All augmentations are applied before cropping and flipping. Additionally, we also evaluate no augmentations (*No Aug.*) as a baseline. Note that the JPEG standard only specifies quality setting in the range $[30, 95]$. However, the authors sample in the range $[30, 100]$. We follow their sampling, but clip all values to the allowed range.

Training details: At the start of the training, we randomly sample 10% of the training set for validation. We train all model using Adam [Kingma and Ba, 2015] with $\beta_1 = 0.9$, $\beta_2 = 0.999$, a batch size of 64, and an initial learning rate of 1^{-4}. For training the spectral classifier, we change the initial learning rate to 1^{-3} since we train the underlying model from scratch. We drop the learning rate by $10\times$ when the validation accuracy stagnates for 5 epochs. When the learning rate reaches 1^{-7}, we terminate training and select the best classifier based on the validation accuracy.

Table 1: **Average Precision (AP) of training a ResNet50 classifier on different subsets of the training data.** We report the average precision over 11 different CNN generators for both, our results (Reproduced) and the results from the original publication (Wang et al.). Chance is 50% and the best possible results is 100%. Note that the classifier is trained on ProGAN and we thus display the results in gray. The mean Average Precision (mAP) is obtained by taking the mean of the individual APs. Note that we apply blurring and JPEG compression with 50% probability during training.

Setting	Result	ProGAN	StyleGAN	BigGAN	CycleGAN	StarGAN	GauGAN	CRN	IMLE	SITD	SAN	DeepFake	**mAP**
2-class	Reproduced	99.3	81.0	71.8	88.4	83.5	89.7	98.8	99.5	82.0	72.3	61.5	84.3
	Wang et al.	98.8	78.3	66.4	88.7	87.3	87.4	94.0	97.3	85.2	52.9	58.1	81.3
4-class	Reproduced	99.8	90.6	79.7	92.5	91.2	93.5	98.1	98.7	95.0	74.6	75.7	90.0
	Wang et al.	99.8	87.0	74.0	93.2	92.3	94.1	95.8	97.5	87.8	58.5	59.6	85.4
8-class	Reproduced	100.0	96.7	81.8	92.7	92.6	95.0	98.2	99.3	85.7	69.8	77.9	90.0
	Wang et al.	99.9	94.2	78.9	94.3	91.9	95.4	98.9	99.4	91.2	58.6	63.8	87.9
16-class	Reproduced	100.0	98.5	87.0	94.5	95.0	96.8	99.2	99.3	84.5	77.9	76.8	91.8
	Wang et al.	100.0	98.2	87.7	96.4	95.5	98.1	99.0	99.7	95.3	63.1	71.9	91.4
20-class	Reproduced	100.0	99.4	88.8	94.1	96.7	96.2	98.5	99.1	92.9	72.3	93.7	93.8
	Wang et al.	100.0	98.5	88.2	96.8	95.4	98.1	98.9	99.5	92.7	63.9	66.3	90.8

We abstained from validating additional combinations of (Adam-) hyperparameters, learning rate, and batch size. The selection of data augmentation and data set variety is already computational challenging and can be viewed as a hyperparameter in its own right. However, we note that this might be an interesting direction for future experiments.

Evaluation criteria: We follow the authors and use Average Precision (AP) as our evaluation metric. It is a ranking-based metric which is not sensitive to the "base rate" of the fraction of fake images and thus commonly used throughout the literature [Zhou et al., 2018, Huh et al., 2018, Wang et al., 2019b]. Additionally, we compute the mean Average Precision (mAP) over all data sets, to compute an overall tendency. During testing, all images are center cropped without resizing to match the training dimension.

Computational requirements: All our experiments are run on two desktop machines. Each machine is running Ubuntu 18.04, with 64GB RAM, a AMD Ryzen 7 3700X 8-Core Processor, and two GeForce RTX 2080Ti. Training a single ResNet50 model on two GPUs on the entire data set (20 classes) takes around 42 hours. Most of the experiments are conducted on smaller data set (2–8 classes), which significantly reduces the training time to roughly 6 hours. Note that the time does not scale linearly, we assume the difference to stem from the fact that the smaller data set still fits into RAM and thus avoids expensive reading from disk. Overall, we estimate the entire experiments to take around 380 GPU hours. This number does not include initial testing or debugging runs.

3 Results

Overall we were able to reproduce the original results and verify the two claims of the publication. In Section 3.1, we examine the first claim, evaluating if more diverse data help generalization. After repeating the related experiment from the original publication, we rerun the experiment with different settings to investigate further. In Section 3.2, we examine the second claim, studying in detail the different data augmentations proposed by the authors. Finally, in Section 3.3, we depart from the original experiments and investigate further to gain additional insights into the proposed methods.

3.1 Claim: *More diverse data improves generalization*

First, we investigate how diversity of the data set affects the results. We aim at reproducing the first half of Table 2 in the original publication. Similar to the original paper, we train five different classifiers on subsets of the original training data ($\{2, 4, 8, 16, 20\}$-classes) and use blurring and JPEG compression with a probability of 50% during training. The results are displayed in Table 1.

Results: Overall, the results are similar to the originally reported numbers. However, we always seem to slightly outperform the results by Wang et al.. We verified our implementation against the original published code and noticed two substantial differences: First, we implemented Gaussian filtering with the build-in methods provided by PyTorch [Paszke et al., 2019], while Wang et al. build their own method using SciPy [Virtanen et al., 2020]. Second, we apply blurring in conjunction with JPEG compression with a probability of 50%. In contrast, Wang et al. first apply

Table 2: **Average Precision of training a ResNet50 classifier on the entire data set with different augmentations.** We report the average precision for both our results (Reproduced) and the the original publication (Wang et al.). Chance is 50%, the best possible results is 100%, and we highlight the ProGAN results in gray since the classifier is trained on similar data. We train on the entire data set (20 classes).

Name	Result	ProGAN	StyleGAN	BigGAN	CycleGAN	StarGAN	GauGAN	CRN	IMLE	SITD	SAN	DeepFake	**mAP**
No Aug.	Reproduced	100.0	96.8	73.5	81.9	100.0	68.2	95.1	88.8	97.1	87.2	98.4	89.7
	Wang et al.	100.0	96.3	72.2	84.0	100.0	67.0	93.5	90.3	96.2	93.6	98.2	90.1
Blur	Reproduced	100.0	94.0	71.7	81.9	100.0	71.0	63.5	63.1	93.8	83.9	98.6	83.8
	Wang et al.	100.0	99.0	82.5	90.1	100.0	74.7	66.6	66.7	99.6	53.7	95.1	84.4
JPEG	Reproduced	100.0	99.4	89.4	93.3	94.5	96.4	99.1	99.3	91.1	69.7	92.5	93.2
	Wang et al.	100.0	99.0	87.8	93.2	91.8	97.5	99.0	99.5	88.7	78.1	88.1	93.0
Blur + JPEG (0.5)	Reproduced	100.0	99.4	88.8	94.1	96.7	96.2	98.5	99.1	92.9	72.3	93.7	93.8
	Wang et al.	100.0	98.5	88.2	96.8	95.4	98.1	98.9	99.5	92.7	63.9	66.3	90.8
Blur + JPEG (0.1)	Reproduced	100.0	99.5	85.7	93.3	97.9	93.6	99.2	99.4	95.5	77.7	95.0	94.2
	Wang et al.	100.0	99.6	84.5	93.5	98.2	89.5	98.2	98.4	97.2	70.5	89.0	92.6

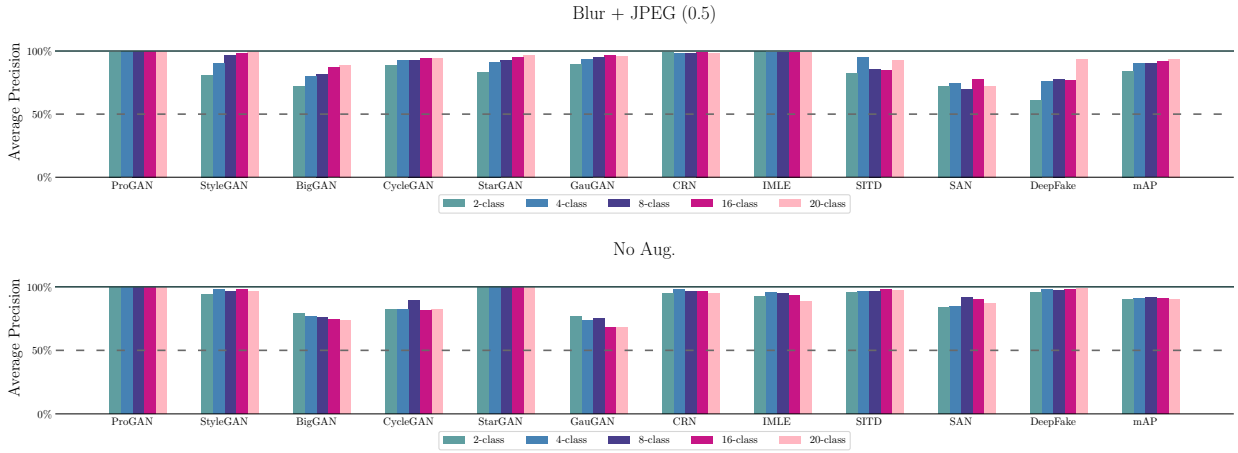

Figure 1: **Comparison of different data set sizes.** When training with augmentations we observe a steady increase in performance while increasing the diversity of the data set. We get mixed results for training without augmentations. Apparently, the higher data diversity can only be utilized with the aid of data augmentations. Best viewed in color.

blurring with a probability of 50% and *then* apply JPEG with a probability of 50%. Note that the paper describes this correctly and we will discuss this further in Section 4.2. We attribute our better performance to these differences.

The original publication observed that data diversity helped with generalization up to a specific point. The authors hypothesized that at this point (16 to 20 classes) the training set may be diverse enough for practical generalization. We agree that data set diversity helps with generalization. We cannot fully agree that there exists a point of diminishing returns. While this is true for a majority of the data sets, utilizing more classes allows our classifiers to achieve significantly better results on the DeepFake ($76.8 \rightarrow 93.7$ AP) and SITD ($84.5 \rightarrow 92.9$ AP) data sets.

We suspect that the diminishing returns may be connected to the aggressive data augmentations used in the original experiments. Thus, we perform a control experiment where we again train multiple classifiers on subsets of the training data with all augmentations disabled. Note the original publication only assessed subsets with augmentations enabled.

The results are presented in Figure 1, a full table can be found in the supplementary material. Contrary to our hypothesis, the data augmentation help the model generalize in the presence of more diverse data. We assume that it acts as a regularizer, preventing the model from overfitting to the training set. However, in Section 3.3.1 we find that using data augmentations in conjunction with a data set generated by StyleGAN2 hurts generalization. Thus, we overall conclude that more diverse data improves generalization, albeit only in the combination with data augmentations dependent on the data set.

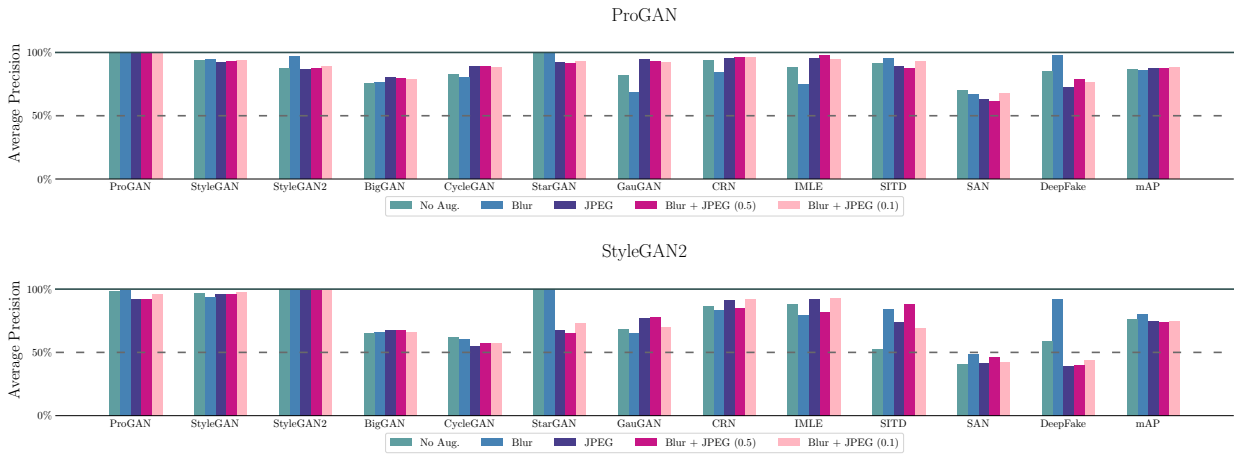

Figure 2: **Comparison between two training sets generated by ProGAN and StyleGAN2.** The classifier trained on data sets generated by ProGAN generalizes significantly better to other architectures. Blurring seems to help the classifier trained on StyleGAN2 to generalize better, the other augmentations hurt performance. Best viewed in color.

3.2 Claim: *Data augmentations improve generalization*

Second, we investigate the claim that data augmentations improve generalization and reproduce the rest of Table 2. Specifically, we train five classifier on the entire training set, with each classifier trained on one data augmentation introduced in Section 2.2. The results are presented in Table 2.

Results: Again, our results are close to the original experiments. We also observe the phenomena that we seem to generally outperform Wang et al.. A notable difference is the blurring results for SAN. We hypothesize that the hand-crafted implementation of Wang et al. differs from the PyTorch implementation we use.

The authors state that data augmentations improve performance and note the SAN and DeepFake data sets as exceptions. Since SAN is a super-resolution model, the only difference between real and generated data is high-frequency components. These components get destroyed by JPEG compression and blurring, which makes a performance drop seem likely. Our model recovers better from using a combination of blurring and JPEG compression for the SAN and DeepFake data sets. We hypothesize that our implementation may apply more aggressive preprocessing (we always apply blurring and JPEG compression), but on fewer training examples (our implementation is all or nothing; Wang et al. can apply blurring, JPEG compression, or both). This implies that other data augmentation strategies might yield even better results and provide an interesting direction for future research in that area.

3.3 Results beyond the original paper

We perform further experiments to get more insights into the original paper. First, we investigate if the results are dependent on the generator used to create the data set. Thus, we created a second data set, consisting of images sampled from StyleGAN2 [Karras et al., 2020] instances and compare them against data sampled from ProGAN. Second, we investigate if the results are dependent on the classifier used. We compare the ResNet50 against a different image classifier (VGG-11) and a classifier which leverages spectral analysis [Frank et al., 2020] to asses if the data augmentations can benefit other classifiers.

3.3.1 Research Question: *Do the results depend on the generator which creates the data set?*

We want to investigate if the generalization capabilities depend on the generator which creates the training set. As described in Section 2.1, we utilized pre-trained StyleGAN2 models to create a novel data set.

We then trained five ResNet50 classifier using the augmentation proposed by Wang et al.. For comparison, we also trained five classifier on LSUN *cat and horse* data generated by ProGAN. In contrast to our previous experiments, we also include a StyleGAN2 test set which was provided by Wang et al. after the original publication. The results are displayed in Figure 2 and the corresponding Table can be found in the supplementary material.

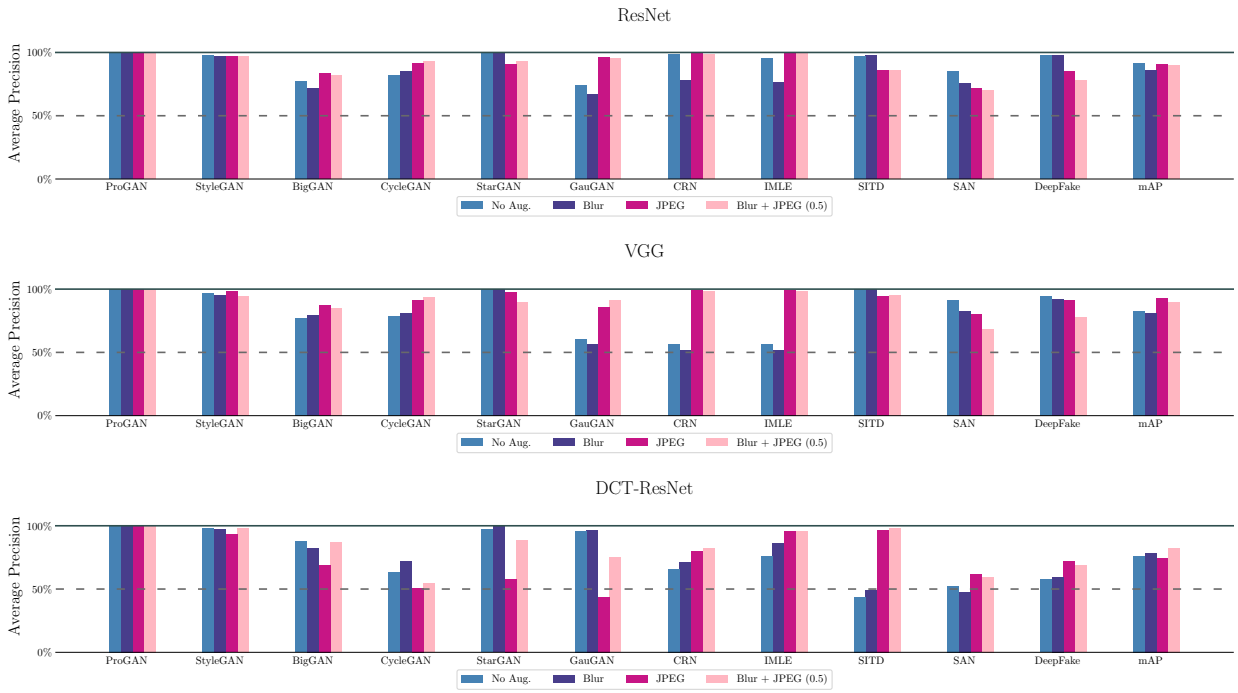

Figure 3: **Different classifier trained with data augmentations.** The data augmentations seem to transfer to other classifiers and, in some instances, are the key to generalization. These experiments were conducted on eight classes from the training data set. Best viewed in color.

We can observe multiple interesting phenomena:

1) Data augmentations help ProGAN generalise to StyleGAN2: When utilizing the proposed data set augmentation, the classifiers trained on ProGAN generalize to detect images from StyleGAN2 ($86.3 \rightarrow 97.1$ AP). However, the same does not apply for StyleGAN2. Data generated by StyleGAN2 already achieves quite a high performance on ProGAN, thus adding data augmentations does only slightly improve or even hurts performance significantly.

2) Data augmentations are specific to the data set: We can observe that despite blurring, the other proposed augmentations worsen the results of StyleGAN2 classifiers. This process is reversed for ProGAN, where blurring hurts mAP, but all other augmentations improve it. Thus data augmentations seem to be specific to the training data set used and should be evaluated carefully.

3) ProGAN serves as a better prior for generalization: If we compare the individual performance and the overall mAP, the classifiers trained on ProGAN seem to better generalize to other architectures. We hypothesize that this is due to the fact that StyleGAN2 is a very recent proposal (released in 2020), while all other generators were proposed earlier (2016-2019). ProGAN is itself an older generator (2018), thus the images generated by it might have more in common with older architectures.

We conclude that the results indeed depend on the generator used for creating the data set. The data augmentations do not transfer and our results imply that going forward, practitioners might have to continuously update their training data sets to generalize to new architectures. Note that we only performed experiments with two classes, adding more classes might improve the performance of StyleGAN2 as a training set.

3.3.2 Research Question: *Do data augmentations transfer to other classifier?*

Finally, we want to investigate the question if the results transfer to other classifiers. We train two other classifiers to compare our results: a VGG-11 model and a classifier which utilizes spectral information (DCT-ResNet). The results are depicted in Figure 3, a table can be found in the supplementary material. The overall results transfer across model architectures and even to the classifier which uses spectral information. This is especially interesting since Frank et al. found that their classifier relied on the high-frequency information, which blurring and JPEG should ruin. Wang et al.

speculated that the data augmentations forced their models to concentrate on differences in the low to mid frequencies, which might help generalization. Overall, we conclude that data augmentations transfer to other classifiers. However, we also note that there exists no silver bullet which helps generalization in general.

4 Discussion

Since we focused our efforts on reproducing the paper from scratch, we can only partially comment on the source code. However, we later used it to compare our implementation against. In the following, we want to give an overview of our observations.

4.1 What was easy?

The overall description in the paper is very detailed. The authors discuss in detail how they collected the data for the test set, provide a good description of their training procedure, and summarize their results clearly with the aid of tables and figures. Additionally, the authors include a long appendix with further details. We did find one description in the paper confusing, which we discuss in more detail in Section 4.2.

As discussed above, we only had a limited exposure to using the authors' code. The parts we looked at were of good code quality, however, sometimes could use more comments. The README file of the repository is very detailed, it includes clear instruction on how to set up the code, download the corresponding data, train models from scratch, and evaluate the results. Additionally, the authors include links to download pre-trained models. We used these models to validate our results and found them quite intuitive to use.

4.2 What was difficult?

We came across very few hurdles when reproducing the code. However, we did end up slightly miss-implementing one of the data augmentations. In the paper, the authors describe their introduced data augmentations in Section 4.1. The description for the *Blur+JPEG* augmentation states: "the image is possibly blurred and JPEG-ed, each with 50% probability". However, the shorthand explicitly uses a plus sign and in the remainder of the paper only the shorthand is used. Thus, based on the frequent use of the plus sign, we assumed the augmentation *always* applies the blurring and JPEG compression together. This resulted in us implementing a slightly different version of the data augmentation, where we apply both in conjunction with 50% probability. We want to stress that the paper states this correctly, we simply missed the last part of the description.

4.3 Communication with original authors

Due to the quality of the paper and the availability of the code repository, we could resolve every question. We did message the authors regarding our different results for the Blur and JPEG augmentation. They agreed that it is difficult to conclude an exact reason why and how augmentations help generalization. For further research direction, they suggested further research into which features the CNN learns.

5 Conclusion

In summary, we believe the paper to be reproducible. We successfully recreated the original experiments from scratch. While our results are slightly different, the overall trend remains the same. However, our additional experiments revealed several limitations to these claims.

More diverse data sets help the classifier generalize to unseen architectures, but: the choice of generator is crucial and data set-specific augmentations have to be used. When using no augmentations, more data diversity hurts the performance of the classifier trained on ProGAN. When we instead use StyleGAN2 as the generator for the data set, we achieve significantly lower performance for all tested configurations. This suggest the possibility that future practitioners might have to constantly update their training sets with new architectures.

Also, the data augmentations seem to be specific for the training data set used. When using ProGAN generated data as the training set, the augmentations transfer across different classifiers. However, only blurring helps StyleGAN2 generalize, while the other proposed augmentations hurt performance. Thus, we advice practitioners to carefully investigate which augmentations help their current situation.

References

Nasir Ahmed, T_ Natarajan, and Kamisetty R Rao. Discrete cosine transform. *IEEE Transactions on Computers*, 1974.

Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations (ICLR)*, 2019.

Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to See in the Dark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Qifeng Chen and Vladlen Koltun. Photographic Image Synthesis with Cascaded Refinement Networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.

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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging Frequency Analysis for Deep Fake Image Recognition. In *International Conference on Machine Learning (ICML)*, 2020.

Karen Hao. The Biggest Threat of Deepfakes isn't the Deepfakes Themselves. *MIT Technology Review*, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Minyoung Huh, Andrew Liu, Andrew Owens, and Alexei A Efros. Fighting Fake News: Image Splice Detection via Learned Self-Consistency. In *European Conference on Computer Vision (ECCV)*, 2018.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations (ICLR)*, 2018.

Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-Realistic Single Image Super-Resolution using a Generative Adversarial Network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Ke Li, Tianhao Zhang, and Jitendra Malik. Diverse Image Synthesis from Semantic Layouts via conditional IMLE. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.

Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017.

Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised Image-to-Image Translation Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic Image Synthesis with Spatially-Adaptive Normalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. FaceForensics++: Learning to Detect Manipulated Facial Images. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.

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. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 2015.

Tom Simonite. Artificial Intelligence Is Coming for Our Faces. *Wired*, 2019.

Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations (ICLR)*, 2015.

Nicholas Thompson and Issie Lapowsky. How Russian Trolls Used Meme Warfare to Divide America. *Wired*, 2017.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 2020.

Run Wang, Lei Ma, Felix Juefei-Xu, Xiaofei Xie, Jian Wang, and Yang Liu. Fakespotter: A simple baseline for spotting ai-synthesized fake faces. *arXiv preprint arXiv:1909.06122*, 2019a.

Sheng-Yu Wang, Oliver Wang, Andrew Owens, Richard Zhang, and Alexei A Efros. Detecting Photoshopped Faces by Scripting Photoshop. In *IEEE International Conference on Computer Vision (ICCV)*, 2019b.

Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A. Efros. CNN-generated images are surprisingly easy to spot...for now. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365*, 2015.

Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2019.

Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. Learning Rich Features for Image Manipulation Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.