

Analyzing and Mitigating JPEG Compression Defects in Deep Learning

Max Ehrlich¹

Larry Davis¹

Ser-Nam Lim²

Abhinav Shrivastava¹

¹University of Maryland

²Facebook AI

{maxehr, lsd}@umiacs.umd.edu

sernamlim@fb.com

abhinav@cs.umd.edu

Abstract

With the proliferation of deep learning methods, many computer vision problems which were considered academic are now viable in the consumer setting. One drawback of consumer applications is lossy compression, which is necessary from an engineering standpoint to efficiently and cheaply store and transmit user images. Despite this, there has been little study of the effect of compression on deep neural networks and benchmark datasets are often losslessly compressed or compressed at high quality. Here we present a unified study of the effects of JPEG compression on a range of common tasks and datasets. We show that there is a significant penalty on common performance metrics for high compression. We test several methods for mitigating this penalty, including a novel method based on artifact correction which requires no labels to train.

1. Introduction

The JPEG compression algorithm [1] has remained the most popular compression algorithm since the early 90s. Despite rapid advances in video compression and application of those technologies to create superior still image compression algorithms, JPEG is still ubiquitous. JPEG is considered a simple solution to storage and transmission of user data. It is well supported and compresses data quite well despite glaring quality loss at low bitrates. Meanwhile, computer vision, driven by the deep learning revolution of 2012 [2], reaches new milestones regularly with respect to performance, and is continually adopted in the mainstream industry and consumer-facing applications.

Despite this continuing application of deep learning to consumer methods, and the immense popularity of JPEG compression in consumer application, the effect of JPEG compression on deep learning models has been poorly studied. Many consumer applications rely on pretrained models either in whole or in part since labeled datasets are computationally expensive to create. These pretrained models are often designed for academic scenarios and use standard datasets. While many of these datasets, like ImageNet [3],




Figure 1: Study results at a glance. Each point shows the performance loss, after applying mitigation, for a model evaluated on quality 10 JPEGs (lower is better), comparing Supervised Fine-Tuning to Task-Targeted Artifact Correction. ▲ Orange triangles are classification models evaluated with Top-1 accuracy on ImageNet, ■ Green squares are detection models evaluated using mAP on MS-COCO, and ◆ blue diamonds are semantic segmentation models evaluated using mIoU on ADE20k. Point size shows the inference throughput (larger is better).

COCO [4], and others, are JPEG compressed, it is often at reasonable quality levels. This is, in general, not reflective of the quality levels seen in consumer applications where the compression level is not always under the control of the deployed system. In light of this, we believe that a comprehensive study which characterizes performance loss due to JPEG compression is long overdue.

In this study, we quantify the effect of JPEG compression over standard performance metrics on a range of datasets and computer vision tasks. We do this by evaluating pretrained models in a systematic and consistent way. In addition to this, we also explore several strategies for mitigating any performance penalty, including a novel method which we call **Task-Targeted Artifact Correction**. This method uses the error of the downstream task logits on JPEG im-

ages with respect to uncompressed versions of the same images, and as such is completely self-supervised: it requires no training labels making it ideal for consumer applications where labels are often expensive to obtain and unreliable. We show that for severe to moderate JPEG compression, there is a steep performance penalty and that this penalty can be mitigated effectively by using our proposed strategies. The study results are summarized in Figure 1, which can be used to quickly choose a model and mitigation technique for a particular application.

Our study helps future scientists make informed decisions when faced with the following two important questions: (1) Is JPEG compression affecting the model? (2) If so, can anything be done about it? Our overall findings are as follows:

1. Heavy to moderate JPEG compression incurs a significant performance penalty on standard metrics, which we show on a set of standard computer vision tasks and datasets. We consider this study to be the primary contribution of our paper and have taken great care to ensure fair comparisons over a plethora of commonly used models. Furthermore, our benchmarking code is fully pluggable and will be released so that models can continue to be evaluated as they are developed. This study is **significantly** more comprehensive than any prior work studying any form of compression (see Section 2 for an overview of prior studies.)
2. If the target application has enough labeled data, and preservation of the model result on uncompressed or losslessly compressed images is not important, then fine tuning the task model using JPEG as a data augmentation strategy effectively mitigates this performance loss.
3. If the target application lacks labels, supervised training is impractical, or performance on uncompressed or losslessly compressed images must be preserved, then JPEG artifact correction can be used as a pre-processing step. We show that off-the-shelf artifact correction improves the performance of downstream tasks greatly, and our self-supervised technique makes a further improvement, approaching that of fine tuning the downstream task directly and surpassing it on some tasks all without requiring ground truth labels.

2. Prior Work

We begin by reviewing several recent works which attempt to characterize computer vision task performance under various image corruptions. We then review deep learning methods which make direct use of JPEG data and conclude with a review of artifact correction techniques which consider downstream tasks.

Analysis of Compression Defects. Recently, several works emerged that consider network performance in the presence of JPEG compression. In dos Santos *et al.* [5],

the authors revisit Gueguen *et al.* [6] and compare the performance of the previous model under different compression settings. They find that the model still faces a performance penalty for compressed images even though it is trained on DCT coefficients. Mandelli *et al.* [7] use two models from the EfficientNet [8] family and compare multimedia-forensics tasks (camera model identification and generated image detection) with computer vision tasks (ImageNet [3] and LSUN [9] classification). The authors test on several JPEG quality factors as well as examine cropping defects (which misalign the JPEG grid). They find that while all models have a loss of performance on low-quality images, computer vision tasks tended to be more robust than multimedia-forensics tasks. Benz *et al.* [10] show a method for improving batch normalization, which increases robustness to several image corruptions. While they do not treat compression specifically, the ImageNet-C [11] dataset includes compressed images as a type of corruption. Hendrycs and Diettrich, in addition to introducing the ImageNet-C [11] dataset, perform a limited benchmarking of several classification models over several JPEG compression settings, including other corruptions they introduce. An older work examining this problem by Zheng *et al.* [12] considers classification and ranking using a technique called stability training, which tries to match network output on an uncorrupted image with that of the corrupted image. They consider moderate JPEG compression.

Deep Learning with JPEG Data. There have been several works in recent years which attempt to merge deep learning with low-level JPEG primitives. Ghosh and Chellappa [13] include a DCT as part of their initial layer and show a performance improvement on classification. Gueguen *et al.* [6] read JPEG DCT coefficients directly into their network and show that the DCT representation requires fewer parameters to learn comparable results to pixels, yielding a speed improvement. Ehrlich and Davis [14] formulate a fully JPEG domain residual network and again show a speed improvement. Lo and Hang [15] show a method for semantic segmentation on DCT coefficients and show both a performance and speed improvement over using pixels. Deguerre *et al.* [16] show a method for object detection on DCT coefficients and again show performance and speed improvement over pixels. Ehrlich *et al.* [17] formulate a JPEG artifact correction network on DCT coefficients and attain state-of-the-art results for color images. Choi and Han use a network to learn task-guided quantization matrices that maximize task performance after JPEG compression [18].

Artifact Correction for Task Improvement. Several recent works have studied using artifact correction in the presence of downstream tasks [19]–[21]. Galteri *et al.* train a GAN [22] for JPEG artifact correction with the primary goal of making an image for human consumption. In [19],

they use an ensemble of networks for each JPEG quality level and manually pick the network at inference time; in [20] they use an auxiliary network to classify a JPEG to its quality level and automatically pick the artifact correction network at inference time. In both works, they show an improvement on object detection tasks. Katakol *et al.* [21] tests semantic segmentation with several different compression algorithms using adversarial restoration.

3. Methodology

Our goal is to simulate a system receiving a JPEG compressed image compressed at some unknown quality level q ¹. We assume that q is chosen uniformly at random from the range [10, 90], below quality 10 there is little information preserved and above quality 90, the image is nearly identical to the uncompressed version. In all cases, the images are compressed before being put through any transformations that the target model prescribes². For evaluation, the images are compressed at each value in the given range in steps of 10, and each model and each mitigation method are evaluated on these images. We use the Independent JPEG Group’s libjpeg [23] software for compression.

For any mitigation methods that require fine tuning, training images are randomly compressed from the same range as a form of data augmentation [24]. The images are always compressed at some quality factor, there are no uncompressed images used for fine tuning. Since the simulated system is assuming JPEG inputs, there is no need for non-JPEG images for fine tuning. The models are trained with a learning rate starting from 1×10^{-3} and ending at 1×10^{-6} for 200 epochs using a cosine annealing [25] learning rate scheduler. We use stochastic gradient descent with momentum for the optimizer with the momentum set to 0.9 and the weight decay set to 5×10^{-4} . As a rule, we only use validation sets for reporting final numbers as most datasets do not provide labeled test data. When training different mitigation techniques for the same model, the same batch size is used³.

For methods requiring artifact correction, we use QGAC [17]. It is important to understand that correction-based mitigation has only recently become practical as the state-of-the-art has shifted from “quality-aware” models to “quality blinded” models. Quality aware models train a different model for each JPEG quality setting, *i.e.*, there would be a unique model for quality 10 JPEGs, quality 20 JPEGs, *etc.*

This is a critical limitation because the JPEG quality is not stored in the JFIF file format, making quality-aware models impossible to use in a real setting. Theoretically, these models could be examined in the contrived scenerio of the study, however, this would require a combinatorially large number of models to be trained - one per quality level per AC model per task model. QGAC is quality-blinded meaning that a single model handles all JPEG quality settings. Additionally, it supports color images, achieves state-of-the-art performance on common benchmarks, and has public code and weights making it a prime candidate for use by other researchers. For completeness, we conducted a limited study using two common quality-aware models in Section 4.4 by limiting the compression quality settings and downstream tasks to make the training tractable.

We evaluate the models with no mitigation as a baseline to characterize the effect of JPEG compression when applied directly to machine learning models. We additionally evaluate two common mitigation methods. Finally, we propose a new mitigation method which requires no supervision to train, something which is crucial outside of the academic setting. We now briefly describe these methods.

Baseline. The baseline method simply passes the JPEG compressed evaluation images to the model unchanged. This serves as an important quantifier for how the JPEG quality affects pre-trained models. This method requires no fine tuning.

Supervised Fine-Tuning. In the Supervised Fine-Tuning mitigation, the task network is fine tuned using the protocol given above with JPEG images as input. In other words, we JPEG compress all images from the training set of the given task and fine tune the network from pretrained weights on the given task to improve performance on JPEG inputs. There are two major drawbacks to this method. The first is that it requires a training set of labeled data. This is not always easy to obtain in consumer applications. The second is that it sacrifices performance on uncompressed images in exchange for performance on compressed images. As we show in Section 4, this method provides good results and a fast runtime, especially at training time (See Appendix E for throughput results).

Artifact Correction. For the artifact correction method, we JPEG compress the evaluation images and then perform artifact correction on them before passing them to the task model. This method requires no fine tuning, and we use pretrained weights for the artifact correction model and for the task model. This method does not sacrifice performance on clean images, since these can be detected (by file extension or mime type) and artifact correction can be skipped in this case; however, this method is not trainable and we will show in the results section that it gives worse performance than fine tuning methods.

¹Most JPEG compression software specifies a scalar quality value in [0, 100] in lieu of a target bitrate, this quality value is not part of the JPEG standard.

²For example, cropping to 224×224 for ImageNet based models. Note that while this is more difficult to implement, the reverse process would incur an unrealistically greater loss of quality.

³*e.g.*, fine tuning a ResNet 18 uses the same batch size as fine tuning artifact correction using resnet18 as the task supervision, but may not use the same batch size as COCO object detection experiments.

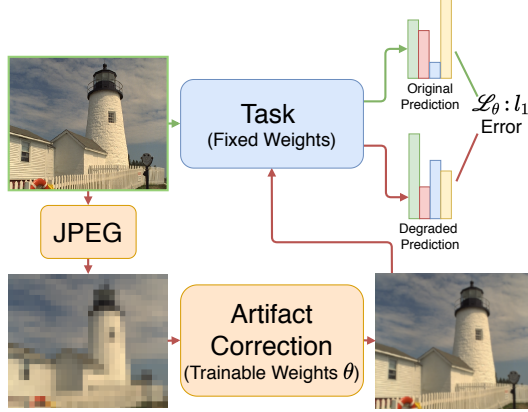


Figure 2: Task-Targeted Artifact Correction. The original image (green) is passed through the task network to obtain the unnormalized prediction logits. Then the image is JPEG compressed and artifact corrected using an artifact correction network with trainable weights. The corrected image is then passed through the task network. The l_1 error between the logits on the original image and the logits on the corrected image is used as loss to tune the artifact correction network weights.

Task-Targeted Artifact Correction. In this novel mitigation technique, we fine tune the artifact correction networking using error on the task network logits. To do this, we minimize the l_1 distance between the task network logits computed on an uncompressed and compressed version of the same image. Formally, given a mini-batch of images B , we minimize

$$\mathcal{L}_\theta = \|\text{Task}(B) - \text{Task}(\text{AC}(\text{JPEG}_q(B); \theta))\|_1, \quad (1)$$

where Task is any task network, AC is the artifact correction network with parameters θ , and q is the JPEG quality level. This is shown schematically in Figure 2. In this way, the artifact correction network learns to correct in a way that maximizes the performance of the downstream network. Note that this method is entirely self-supervised, no ground truth labels are used (the only loss comes from the difference in behavior on the compressed vs. uncompressed versions of the images). Not only is this easier to deploy in a consumer setting where labeled data is hard to obtain, but it also alleviates any concern that the artifact correction network may be learning to perform the task for the task network, artificially increasing the number of parameters and making for an unfair comparison. By using l_1 error on the network output, the method is applicable to many different tasks with little or no modification. This method again does not sacrifice performance on uncompressed images since the task network weights are unchanged and an uncompressed input can be detected and the artifact correction step skipped. We show in Section 4 that this method approaches fine tuning the task network directly, even exceeding it in some tasks, despite having no access to ground truth labels.

In addition to this, Task-Targeted Artifact correction supports flexible training scenarios. We show in Section 4.3 that the networks weights are **transferrable** *e.g.*, an artifact correction network which was trained for one network can be used for other networks. This allows a lightweight training setup where a “fast-to-train” network is used to create the artifact correction network weights and then reused for models which would be cumbersome to train. Similarly, Task-Targeted Artifact Correction supports **multi-head** training, where multiple downstream tasks are used at the same time during training. While this method bears a superficial similarity to stability training [12], both transfer and multihead are impossible with stability training and the exact loss formulation we use (Equation 1) is universal: it can be easily applied to many tasks.

4. Results

Since our study is focused on both **analyzing** (Sections 4.1 and 4.2) and **mitigating** (Sections 4.3, and 4.4), our results are organized into one of these two categories. First, we show an abridged version of the study on JPEG compression, the full study results are available in Appendix D. After this, we present a preliminary study of two recent forensics models which were tested using the same methodology as the main study. We then examine the transferability and multihead scenarios mentioned in Section 3 in detail, a unique property of Task-Targeted Artifact Correction that makes it a flexible mitigation option.

4.1. Analyzing: Abridged Study Results

We show a subset of our results highlighting interesting, unusual, or edge cases as well as discussing overall findings for each major task. We aim to provide a good covering of common tasks and datasets and picked several models and a single representative dataset per task. While our work is the most comprehensive study of its kind to date, it is not exhaustive: there exist models and datasets which we did not test on. In most cases, we started from pretrained weights available in commodity deep learning libraries.

The results presented in the body of the paper are restricted to qualities [10, 50] (heavy to moderate compression) for brevity and because these are the most interesting results. Results on the full range of qualities we considered ([10, 90]) along with tables of results are available in our appendices. In each case, the results are shown as plots giving loss in performance on a task appropriate metric vs. JPEG quality level. We studied Classification, Object Detection, Instance and Semantic Segmentation as a set of computer vision tasks.

The abbreviated results are shown in Figure 3, which shows the results of all models with no mitigation separated by task, and Figure 4, which shows the behavior of individual models with mitigations applied compared to no




Figure 3: Performance loss due to JPEG compression by task. The plots show all models from a single task with no mitigation applied. For segmentation tasks, the format of the model name is `Encoder Model + Decoder Model` and “ds” indicates that the model was trained with deep supervision. Note that methods which use a Pyramid Pooling Module (PPM) decoder always use deep supervision.



Figure 4: Performance loss of tested models with mitigation applied. **• Circle:** No Mitigation, **+ Cross:** Off-the-Shelf Artifact Correction, **◆ Diamond:** Task-Targeted Artifact Correction, **■ Square:** Supervised Fine-Tuning. The models in this figure correspond to those shown in Figure 3. These plots are best viewed digitally and are intended to show the trend of models for each task. For plots with individual models, please see Appendix D.

mitigation in condensed form. Please see Appendix D for full plots and tables for individual models. All performance measures show the absolute drop in performance, *e.g.*, a 14% drop indicates that the model performs 14% worse on JPEG images at the given quality than on uncompressed images. For the reference numbers that we used for this comparison, see Appendix D.4. As a general rule, we observed that more complex tasks incurred a larger performance penalty on JPEG compressed inputs. Additionally, the more complex the task, the more likely it was to be aided by Task-Targeted Artifact Correction, and the less likely it was to be aided by Supervised Fine-Tuning. In Figure 5, we show qualitative results for detection, instance and semantic segmentation. Please see Appendix C to view these results in more detail. We now briefly discuss the details of the study for each task.

Classification. We tested classification models using the ImageNet [3] dataset. We tested the following models: MobileNetV2 [26], ResNet 18, 50, and 101 [27], ResNeXt 50 and 101 [28], VGG 19 [29], InceptionV3 [30], and EfficientNet B3 [8]. The pretrained weights for the task networks in this section come from the torchvision library [31]. The evaluation metric used was Top-1 Accuracy. Models

in this task generally responded better to Supervised Fine-Tuning than to Task-Targeted Artifact Correction with the notable exception of MobileNetV2 and EfficientNet which responded better to Task-Targeted Artifact Correction. We ran GradCAM [32] to examine Class Activation Maps for JPEG inputs as well as all mitigations. We found that JPEG degrades the gradient quality as well as induces localization errors. Please see Appendix A for this analysis.

Detection and Instance Segmentation. Next, we show results on object detection and instance segmentation. These models were tested using the MS-COCO dataset [4]. We tested three detection models: Fast R-CNN [33], Faster R-CNN [34], and RetinaNet [35], and we used Mask R-CNN [36] for instance segmentation. The pretrained weights come from the Detectron2 library [37]. In all cases, we use a model with a ResNet 50 [27] backbone, and for the R-CNNs we use a Feature Pyramid Network [38] for the detector. For Task-Targeted Artifact Correction training, we use loss on only the backbone features rather than the detection logits. The evaluation metric used was mean average precision (mAP). Models in this section responded well to all mitigation techniques, with Task-Targeted Artifact Correction helping the most for low-quality settings.





Figure 5: Qualitative results. Top: FasterRCNN Detection, Middle: MaskRCNN Instance Segmentation, Bottom: HRNetV2 + C1 Semantic Segmentation. All inputs were compressed at quality 10. Note the poor quality of the results with no mitigation. In particular, the detection prediction of “boat”, while incorrect, is reasonable based on the deformed shape and lack of texture of the compressed image. Despite the compression, the network is displaying some understanding that there is an ocean scene. Similarly, for MaskRCNN, skateboard is predicted seemingly on the shape of the region alone. As mitigations are applied, the confidence of the correct detections for both detection and instance segmentation increases. For semantic segmentation, the result is mostly incorrect with seemingly random classification regions which mitigations are able to clean up significantly.

The moderate quality settings, however, responded better to Supervised Fine-Tuning. For a further analysis of detection, we used TIDE [39] and determined that missed detections make up the bulk of errors caused by JPEG compression, please see Appendix B for this analysis.

Semantic Segmentation. For semantic segmentation, we show results on ADE20K using the accompanying code [40], [41]. This code is organized into pluggable encoders and decoders. We test the following encoders: MobileNetV2 [26], ResNet 18, 50, and 101 [27], and HRNet [42]. We test the following decoders: C1 with deep supervision [43], PSPNet (the Pyramid Pooling Module) with and without deep supervision [43], and UPerNet [44]. Similar to the object detection experiments, we train Task-Targeted Artifact Correction using loss from the encoder features only. The metric used is mean of per-class intersection-over-union of classified pixels (mIoU). In general, segmentation models were greatly affected by compressed inputs and Supervised Fine-Tuning was not an effective mitigation, with some models performing worse after fine tuning than with no mitigation at all. Conversely, Task-Targeted Artifact Correction was able to effectively mitigate performance loss for low and moderate qualities.

4.2. Analyzing: Limited Study on Forensics

In addition to the above computer vision tasks, we conducted a limited study on two recent forensics models: Wang *et al.* [24] and Chai *et al.* [45]. The goal of both of these models is to detect an image which, in whole or in part, was generated by a GAN [22], and can be expressed as a two-class classification problem with classes “real” and “fake”. Wang *et al.* uses a ResNet 50 architecture [27] to classify the images. Chai *et al.* uses a slightly modified Xception [46] model to classify patches, forming an image label by majority vote among the patches. Both methods use their own bespoke datasets and provide pretrained models which we obtained and used for our experiments. The evaluation metric we used was intended to match the one presented in each paper. For Wang *et al.* that is accuracy of the real/fake prediction, while for Chai *et al.* it is accuracy of the predicted patches. The results are shown in Figure 6. Despite the similar formulation, both networks have completely different behavior on compressed images than traditional ImageNet [3] classification and furthermore have completely different behavior to each other. We note that both models have nearly 100% accuracy on clean images they were trained on. Under JPEG compression, even up to the quality 90 (the maximum that we tested) performance of Chai *et al.* stays around 50% indicating the equivalent to a



(a) Forensic task performance with no mitigation. (b) Forensics task performance with mitigation

Figure 6: Forensic model results. • **Circle**: No Mitigation, + **Cross**: Off-the-Shelf Artifact Correction, ♦ **Diamond**: Task-Targeted Artifact Correction, ■ **Square**: Supervised Fine-Tuning.

random guess while Wang *et al.* is able to match its performance by quality 50. While both models incorporated slight JPEG-as-data-augmentation into their pretrained weights, they are greatly aided by additional Fine-Tuning although artifact correction of any kind provided little to no benefit. Since these methods are intended to spot potentially malicious behavior, we find it troubling that a simple perturbation like JPEG is capable of fooling the networks to such a degree and we hope that this study will provide an impetus for further investigation.

4.3. Mitigating: Transferability and Multihead

In this section, we show results which are intended to give a better understanding of our proposed Task-Targeted Artifact Correction. In particular, we examine the transferability of the targeted models, *e.g.*, if a model which was trained for one network and task, like a lightweight MobileNetV2 for ImageNet classification, for example, can be used to mitigate performance loss of a more complex and harder to train network like a ResNet-101, or even used in a wholly different task like detection. We also examine if this performance mitigation holds when a correction network is targeted to multiple models simultaneously, in other words, when multiple task networks are providing supervisory signal. These scenarios improve the practical application of the method by allowing for more flexible training and model reuse in addition to the stated advantage of not requiring labels. A table of results for all plots in this section is given in Appendix F.

Intra-Task Transferability. For intra-task transferability, we tested two scenarios. We first trained Task-Targeted Artifact Correction models on ResNet 18 [27] and MobileNet V2 [26], two of our smallest models which are fast to train. We then tested these models on ImageNet [3] classification using ResNet 101 [27] as the downstream network, one of our largest and slowest to train models. The result shown in Figure 7a is comparable with the task-targeted model trained on the ResNet 101 itself, indicating good transfer.

Inter-Task Transferability. For inter-task transferability, we used the same artifact correction models trained in the

Table 1: Comparison with Common AC Baselines for Task-Targeted Artifact Correction. ARCNN and IDCN are both “quality aware” models. Metrics are classification: top-1 accuracy, detection: mAP, segmentation: mIoU. QGAC outperforms both.

Task (Network)	AC Network	Quality 10	Quality 20
MobileNetV2 (Classification)	ARCNN	61.19	66.89
	IDCN	62.62	65.72
	QGAC	64.64	68.63
FasterRCNN (Detection)	ARCNN	24.99	27.02
	IDCN	25.99	28.23
	QGAC	31.43	33.85
HRNetV2 + C1 (Semantic Segmentation)	ARCNN	29.35	35.90
	IDCN	32.62	37.00
	QGAC	34.14	37.61

previous section but now used them to test using COCO [4] object detection on Faster R-CNN [47] and ADE20K [40], [41] semantic segmentation using the HRNetV2 [42] encoder with C1 decoder. These results are shown in Figures 7b and 7c, and again show good transfer with comparable performance to task-targeted networks trained for each of the networks.

Multihead. We tested two setups for multihead training: a two task setup and a three task setup. For the two task setup, we train the artifact correction network using a ResNet 18 [27] as well as a Faster R-CNN [47] providing downstream loss. For the three task setup we add in semantic segmentation using HRNetV2 [42] + C1. To train this, we alternate batches, taking one batch from ImageNet, performing a full forward pass using the Resnet 18, then taking a batch from COCO and performing another full forward pass using the Faster R-CNN backbone and in three task case following the same procedure with ADE20K. The l_1 loss of the features for each separate network is taken and the backpropagated loss is summed. The result, shown in Figure 8 is on-par with artifact correction models trained with a single model, and in the classification case, both multi-head models perform better than the single task corrector indicating better generalization.

4.4. Mitigating: Limited Comparison with Other Artifact Correction Models

As discussed in Section 3, artifact-correction based mitigation has only recently become viable using quality-blind methods. Using quality-aware methods, which train a unique model for each quality setting, is impossible in real scenarios because the quality setting is not stored in the JPEG file, and in the context of the study, leads to an intractable training protocol which would require models to be trained per-quality level and per-model. Nevertheless, since quality aware models dominated the literature for many years, there is some interest in their behavior in this context. To make this tractable we restrict this part of the study to two models: ARCNN [48], a standard baseline




Figure 7: Transfer Results. In all plots, we add an evaluation using artifact correction weights that were trained on ResNet-18 and MobileNetV2, our lightest weight models. Note that “Fine-Tuned” and “Task-Targeted Artifact Correction” methods are both trained using their respective task network directly *e.g.* in (a) they use a ResNet 101. -- dashed lines indicate results shown in Section 4.1.




Figure 8: Multihead Results. In all plots, we add an evaluation using artifact correction weights that were trained using multiple task networks. For the two task setup, we used ResNet-50 and FasterRCNN. For the three task setup, we used ResNet-50, FasterRCNN, and HRNetV2 + C1. Note that HRNetV2 + C1 has no two-task multihead model. -- dashed lines indicate results shown in Section 4.1.

in artifact correction, and IDCN [49] a recent model. AR-CNN was modified to handle color images and IDCN handles them natively. Both models were ported to our framework and retrained to within 0.5dB PSNR of the published numbers. We additionally restrict the quality settings to 10 and 20 only (the quality settings reported by IDCN). Table 1 shows that both models perform worse than QGAC, further motivating our use of it for the main study.

5. Conclusion

In this paper we conducted a large scale study of JPEG compression on common computer vision tasks and datasets. Our study shows that JPEG compression has a steep penalty across the board for heavy to moderate compression settings. We also tested several strategies for mitigating this performance penalty including a novel method which requires no labels. Our proposed mitigation strategy achieves better results than other unlabeled mitigations and since it requires no labels it is ideal for consumer facing applications where labels are often hard to obtain. For complex tasks, our method outperforms the supervised method despite having no access to ground truth labels. Our method promotes model reuse by allowing transfer of weights between tasks and multihead training. For further results, including throughput for each tested model (Appendix E) we

invite readers to our appendices.

We hope to extend this work by considering more compression methods. Despite their relative disuse, there are other still image compression techniques, such as JPEG 2000 [50], HEIF [51], WebP [52], and BPG [53] that should be considered. Consideration of learned-image-compression algorithms is also a useful addition. Finally, video processing models are becoming commonplace, and video compression is almost exclusively lossy. An extended study should consider these models over different video compression settings.

It is our hope that this will be a living study that it will be updated as new models, compression algorithms, and mitigation techniques are developed. Our benchmarking code is pluggable to allow for experimentation with different models, and will be made freely available. Further, since our Task-Targeted Artifact Correction method shows good transferability, we plan to make the weights we trained for this study available for general use. Many applications will benefit from using our TTAC weights with no modification.

Acknowledgement This project was partially supported by independent grants from Facebook AI, DARPA SemaFor (HR001119S0085) and DARPA SAIL-ON (W911NF2020009) programs.

References

- [1] G. K. Wallace, “The jpeg still picture compression standard,” *IEEE transactions on consumer electronics*, vol. 38, no. 1, pp. xviii–xxxiv, 1992.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [4] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft coco: Common objects in context,” in *European conference on computer vision*, Springer, 2014, pp. 740–755.
- [5] S. F. dos Santos, N. Sebe, and J. Almeida, “The good, the bad, and the ugly: Neural networks straight from jpeg,” in *2020 IEEE International Conference on Image Processing (ICIP)*, IEEE, 2020, pp. 1896–1900.
- [6] L. Gueguen, A. Sergeev, B. Kadlec, R. Liu, and J. Yosinski, “Faster neural networks straight from jpeg,” in *Advances in Neural Information Processing Systems*, 2018, pp. 3933–3944.
- [7] S. Mandelli, N. Bonettini, P. Bestagini, and S. Tubaro, “Training cnns in presence of jpeg compression: Multimedia forensics vs computer vision,” *arXiv preprint arXiv:2009.12088*, 2020.
- [8] M. Tan and Q. V. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” *arXiv preprint arXiv:1905.11946*, 2019.
- [9] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao, “Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop,” *arXiv preprint arXiv:1506.03365*, 2015.
- [10] P. Benz, C. Zhang, A. Karjauv, and I. S. Kweon, “Revisiting batch normalization for improving corruption robustness,” *arXiv preprint arXiv:2010.03630*, 2020.
- [11] D. Hendrycks and T. Dietterich, “Benchmarking neural network robustness to common corruptions and perturbations,” *arXiv preprint arXiv:1903.12261*, 2019.
- [12] S. Zheng, Y. Song, T. Leung, and I. Goodfellow, “Improving the robustness of deep neural networks via stability training,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4480–4488.
- [13] A. Ghosh and R. Chellappa, “Deep feature extraction in the dct domain,” in *Pattern Recognition (ICPR), 2016 23rd International Conference on*, IEEE, 2016, pp. 3536–3541.
- [14] M. Ehrlich and L. S. Davis, “Deep residual learning in the jpeg transform domain,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 3484–3493.
- [15] S.-Y. Lo and H.-M. Hang, “Exploring semantic segmentation on the dct representation,” in *Proceedings of the ACM Multimedia Asia on ZZZ*, 2019, pp. 1–6.
- [16] B. Deguerre, C. Chatelain, and G. Gasso, “Fast object detection in compressed jpeg images,” *arXiv preprint arXiv:1904.08408*, 2019.
- [17] M. Ehrlich, L. Davis, S.-N. Lim, and A. Shrivastava, “Quantization guided jpeg artifact correction,” *Proceedings of the European Conference on Computer Vision*, 2020.
- [18] J. Choi and B. Han, “Task-aware quantization network for jpeg image compression,” in *European Conference on Computer Vision*, Springer, 2020, pp. 309–324.
- [19] L. Galteri, L. Seidenari, M. Bertini, and A. Del Bimbo, “Deep generative adversarial compression artifact removal,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 4826–4835.
- [20] —, “Deep universal generative adversarial compression artifact removal,” *IEEE Transactions on Multimedia*, 2019.
- [21] S. Katakol, B. Elbarashy, L. Herranz, J. van de Weijer, and A. M. Lopez, “Distributed learning and inference with compressed images,” *arXiv preprint arXiv:2004.10497*, 2020.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [23] I. J. Group. (). “Libjpeg,” [Online]. Available: <http://libjpeg.sourceforge.net>.
- [24] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, “Cnn-generated images are surprisingly easy to spot...for now,” in *CVPR*, 2020.
- [25] I. Loshchilov and F. Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” *arXiv preprint arXiv:1608.03983*, 2016.

- [26] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, “Mobilenetv2: Inverted residuals and linear bottlenecks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated residual transformations for deep neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1492–1500.
- [29] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [30] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [31] S. Marcel and Y. Rodriguez, “Torchvision the machine-vision package of torch,” in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1485–1488.
- [32] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [33] R. Girshick, “Fast r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- [34] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1137–1149, 2016.
- [35] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [36] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [37] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, *Detectron2*, <https://github.com/facebookresearch/detectron2>, 2019.
- [38] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature pyramid networks for object detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [39] D. Bolya, S. Foley, J. Hays, and J. Hoffman, “Tide: A general toolbox for identifying object detection errors,” in *ECCV*, 2020.
- [40] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, “Semantic understanding of scenes through the ade20k dataset,” *arXiv preprint arXiv:1608.05442*, 2016.
- [41] —, “Scene parsing through ade20k dataset,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [42] K. Sun, Y. Zhao, B. Jiang, T. Cheng, B. Xiao, D. Liu, Y. Mu, X. Wang, W. Liu, and J. Wang, “High-resolution representations for labeling pixels and regions,” *arXiv preprint arXiv:1904.04514*, 2019.
- [43] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [44] T. Xiao, Y. Liu, B. Zhou, Y. Jiang, and J. Sun, “Unified perceptual parsing for scene understanding,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 418–434.
- [45] L. Chai, D. Bau, S.-N. Lim, and P. Isola, “What makes fake images detectable? understanding properties that generalize,” *arXiv preprint arXiv:2008.10588*, 2020.
- [46] F. Chollet, “Xception: Deep learning with depthwise separable convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
- [47] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [48] C. Dong, Y. Deng, C. C. Loy, and X. Tang, “Compression artifacts reduction by a deep convolutional network,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 576–584.
- [49] B. Zheng, Y. Chen, X. Tian, F. Zhou, and X. Liu, “Implicit dual-domain convolutional network for robust color image compression artifact reduction,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 11, pp. 3982–3994, 2019.

- [50] M. W. Marcellin, A. Bilgin, M. J. Gormish, and M. P. Boliek, “An overview of jpeg-2000,” in *Proceedings of the Conference on Data Compression*, ser. DCC '00, USA: IEEE Computer Society, 2000, p. 523, ISBN: 0769505929.
- [51] M. Hannuksela, E. Aksu, V. M. VadaKital, and J. Lainema, “Overview of the high efficiency image file format,” in *JCTVC-V0072*, 2015.
- [52] A. WebP, *New image format for the web*, 2020.
- [53] F. Bellard, *Bpg image format*. [Online]. Available: <https://bellard.org/bpg/>.

A. GradCam Visualizations

In order to provide more insight into the potential cause of misclassifications on JPEG images, besides the obvious answer of “quality degradation”, we provide visualizations using GradCam [32]. The visualizations are shown in Figures 9 - 11, see Figure 14 for enlarged inputs as well as predicted and ground-truth classes (the fine tuned model also predicts the correct class on this image). The visualizations compare the gradient and class-activation-maps (CAMs) for each mitigation technique to passing the JPEG directly (“no mitigation”) and to passing the original image directly. What we see is quite telling. The gradient with no mitigation is degraded significantly with respect to the original input, however the CAM indicates that it is still focusing in the correct location in the image. Fine-tuning the model greatly improves the quality of the gradient, however the CAM localization is now off. Off-the-shelf artifact correction improves the gradient quality, however the localization is now less constrained and the network appears unsure of where in the image to focus. Finally, Task-Targeted artifact correction seems to make an improvement in the gradient while preserving the CAM localization.

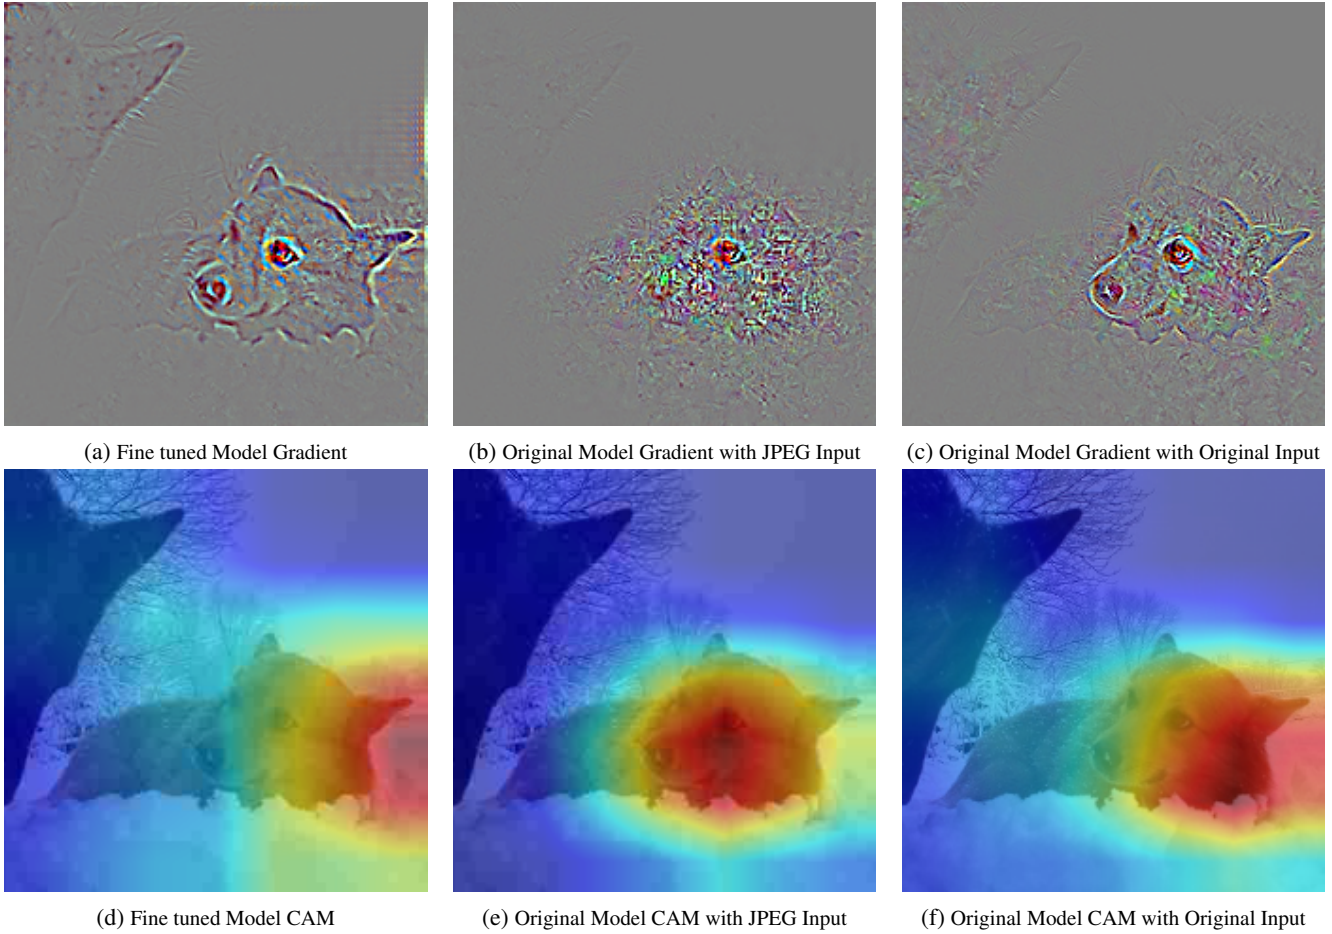





Figure 9: Fine Tuned Model Comparison




(a) Off-the-Shelf AC Gradient




(b) Original Model Gradient with JPEG Input




(c) Original Model Gradient with Original Input



(d) Off-the-Shelf AC CAM




(e) Original Model CAM with JPEG Input




(f) Original Model CAM with Original Input


Figure 10: Off-the-Shelf Artifact Correction Comparison




(a) Task-Targeted AC Gradient




(b) Original Model Gradient with JPEG Input




(c) Original Model Gradient with Original Input



(d) Task-Targeted AC CAM



(e) Original Model CAM with JPEG Input



(f) Original Model CAM with Original Input


Figure 11: Task-Targeted Artifact Correction Comparison

B. Detection Errors

Here we look deeper at the detection errors produced by JPEG compressed inputs using TIDE [39]. TIDE computes a breakdown of exactly which errors contributed to mAP loss during evaluation of detection and instance segmentation and shows the breakdown graphically in a condensed yet informative format. We ran TIDE evaluation on FasterRCNN for box detection and MaskRCNN for instance segmentation with no mitigations applied to understand how JPEG effects specific detection errors.


The results show similar behavior for both methods. On low quality JPEGs, the bulk of the errors are missed detections. This can be seen in the pie chart showing the relative proportions of missed detections, which is roughly 50% for quality 10, and in the high number of false negatives in the bar chart on the lower right. As the quality increases, the proportion of missed detections gradually decreases and at high quality, localization errors make up a larger proportion of the errors. It should be noted that although the proportion of error attributed to localization increases, the detections overall are much more reliable on high quality JPEGs as expected. This can be seen in the significantly lower false negative rate as well as the scale of the x-axis of the bar chart in the bottom left of the images.

coco_instances_results




(a) Q=10

coco_instances_results



(b) Q=50


coco_instances_results



(c) Uncompressed


Figure 12: FasterRCNN TIDE Plots

coco_instances_results




(a) Q=10

coco_instances_results



(b) Q=50

coco_instances_results



(c) Uncompressed

Figure 13: MaskRCNN TIDE Plots

C. Qualitative Results

Since the proposed Task-Targeted Artifact Correction is at its core an image-to-image regression technique, we provide some qualitative results here that show images with their downstream task network behavior. All of the images in this section were compressed at quality 10 before being corrected. Where appropriate we also visualize the result of the Supervised Fine-Tuning method for comparison.



(a) JPEG Q=10, Prediction: “Norwich terrier”, Fine-Tuned Prediction: “Pembroke, Pembroke Welsh corgi”



(b) Off-the-Shelf Artifact Correction, Prediction: “basenji”



(c) Task-Targeted Artifact Correction, Prediction: “Pembroke, Pembroke Welsh corgi”




(d) Original, Prediction: “Pembroke, Pembroke Welsh corgi”


Figure 14: MobileNetV2, Ground Truth: “Pembroke, Pembroke Welsh corgi”



(a) JPEG Q=10




(b) Off-the-Shelf Artifact Correction




(c) Task-Targeted Artifact Correction



(d) Supervised Fine-Tuning




(e) Original

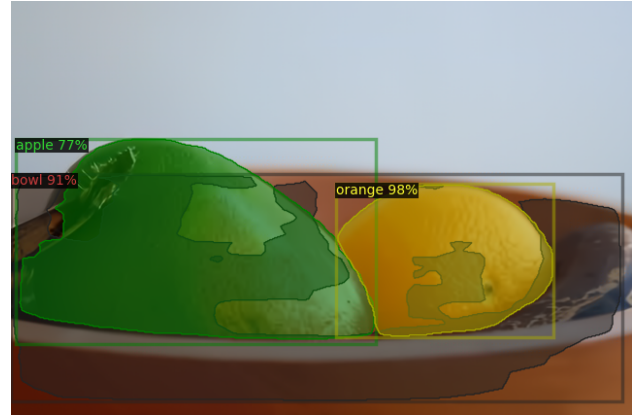


(f) Ground Truth


Figure 15: FasterRCNN




(a) JPEG Q=10




(b) Off-the-Shelf Artifact Correction




(c) Task-Targeted Artifact Correction



(d) Supervised Fine-Tuning




(e) Original



(f) Ground Truth


Figure 17: MaskRCNN



(a) JPEG Q=10




(b) Degraded Prediction




(c) Ground Truth




(d) Off-the-Shelf Artifact Correction



(e) Off-the-Shelf Artifact Correction Prediction




(f) Ground Truth




(g) Task-Targeted Artifact Correction



(h) Task-Targeted Artifact Correction Prediction




(i) Ground Truth



(j) JPEG Q=10



(k) Supervised Fine-Tuning Prediction



(l) Ground Truth

Figure 18: HRNetV2 + C1

D. Full Study Results

Here we give the full results of the study including plots and tables of results for JPEG quality levels [10, 90]. The results are shown visually in plots similar to those given in the body of the paper and the raw numbers are provided in tables.

D.1. Plots of Results







Figure 19: Classification




(a) MobileNetV2




(b) VGG-19




(c) InceptionV3




(d) ResNeXt 50




(e) ResNeXt 101




(f) ResNet 18



(g) ResNet 50



(h) ResNet 101



(i) EfficientNet B3







Figure 21: Detection and Instance Segmentation




(a) FastRCNN



(b) FasterRCNN



(c) RetinaNet



(d) MaskRCNN







Figure 23: Semantic segmentation




(a) HRNetV2 + C1




(b) MobileNetV2 + C1




(c) ResNet 18 + PPM




(d) ResNet 50 + UPerNet



(e) ResNet 50 + PPM



(f) ResNet 101 + UPerNet



(g) ResNet 101 + PPM






Figure 25: Forensics



(a) Chai *et al.*



(b) Wang *et al.*

D.2. Tables of Results

Model	Metric	Reference	Mitigation	Q=10	Q=20	Q=30	Q=40	Q=50	Q=60	Q=70	Q=80	Q=90
EfficientNet B3	Top-1 Accuracy	83.98	Supervised Fine-Tuning	79.78	81.84	82.47	82.68	82.78	82.75	82.83	82.85	82.83
			None	77.24	81.11	81.95	82.52	82.67	82.91	83.10	83.37	83.75
			Off-the-Shelf Artifact Correction	75.92	80.02	81.47	82.12	82.44	82.71	82.94	83.23	83.70
			Task-Targeted Artifact Correction	81.03	82.71	83.21	83.53	83.64	83.71	83.73	83.80	83.76
InceptionV3	Top-1 Accuracy	77.33	Supervised Fine-Tuning	75.11	77.25	77.77	77.89	78.13	78.13	78.24	78.26	78.32
			None	69.38	74.15	75.44	75.98	76.38	76.69	76.95	77.14	77.30
			Off-the-Shelf Artifact Correction	71.21	75.04	76.09	76.42	76.68	76.79	76.97	77.06	77.13
			Task-Targeted Artifact Correction	73.65	75.89	76.53	76.82	76.93	76.99	77.09	77.15	77.10
MobileNetV2	Top-1 Accuracy	70.72	Supervised Fine-Tuning	65.65	69.21	69.92	70.20	70.37	70.53	70.50	70.55	70.54
			None	57.23	65.55	67.87	68.95	69.47	69.98	70.24	70.60	70.86
			Off-the-Shelf Artifact Correction	57.33	65.25	67.76	68.93	69.60	70.07	70.40	70.71	70.58
			Task-Targeted Artifact Correction	64.64	68.63	69.71	70.18	70.32	70.44	70.50	70.52	70.34
ResNet-101	Top-1 Accuracy	76.91	Supervised Fine-Tuning	74.63	76.50	77.07	77.20	77.27	77.29	77.43	77.44	77.53
			None	66.12	73.00	74.65	75.39	75.83	76.29	76.51	76.79	76.96
			Off-the-Shelf Artifact Correction	67.91	73.64	75.09	75.84	76.23	76.52	76.56	76.80	76.74
			Task-Targeted Artifact Correction	72.99	75.53	76.30	76.60	76.59	76.72	76.70	76.72	76.59
ResNet-18	Top-1 Accuracy	68.84	Supervised Fine-Tuning	65.49	68.46	69.07	69.16	69.36	69.33	69.38	69.53	69.49
			None	57.62	65.26	67.07	67.68	68.08	68.30	68.61	68.84	68.92
			Off-the-Shelf Artifact Correction	61.19	66.39	67.87	68.39	68.61	68.77	68.97	68.99	68.90
			Task-Targeted Artifact Correction	63.83	67.06	68.04	68.24	68.35	68.48	68.52	68.60	68.50
ResNet-50	Top-1 Accuracy	75.31	Supervised Fine-Tuning	73.18	75.46	76.02	76.24	76.36	76.42	76.52	76.52	76.55
			None	63.43	71.20	73.23	74.10	74.43	74.63	75.01	75.09	75.34
			Off-the-Shelf Artifact Correction	66.90	72.45	73.95	74.60	74.93	75.18	75.26	75.42	75.30
			Task-Targeted Artifact Correction	70.48	73.56	74.39	74.81	74.94	75.00	74.98	74.98	74.89
ResNeXt-101	Top-1 Accuracy	78.81	Supervised Fine-Tuning	75.60	78.00	78.50	78.71	78.86	78.97	79.01	78.98	79.06
			None	68.83	74.84	76.39	77.05	77.60	78.00	78.16	78.56	78.75
			Off-the-Shelf Artifact Correction	71.19	75.88	77.14	77.80	78.15	78.30	78.57	78.66	78.61
			Task-Targeted Artifact Correction	74.73	77.33	78.08	78.29	78.55	78.62	78.68	78.73	78.68
ResNeXt-50	Top-1 Accuracy	76.99	Supervised Fine-Tuning	74.21	76.23	76.79	77.01	77.08	77.18	77.16	77.30	77.17
			None	66.96	73.21	74.85	75.62	76.07	76.37	76.63	76.88	77.06
			Off-the-Shelf Artifact Correction	68.05	73.56	75.11	75.95	76.38	76.59	76.71	76.99	76.90
			Task-Targeted Artifact Correction	72.22	75.45	76.09	76.62	76.86	76.83	76.85	76.99	76.81
VGG-19	Top-1 Accuracy	73.44	Supervised Fine-Tuning	69.50	72.66	73.29	73.74	73.83	73.85	73.95	74.14	74.11
			None	59.27	68.08	70.49	71.53	71.99	72.42	72.80	73.24	73.46
			Off-the-Shelf Artifact Correction	61.93	68.79	70.82	71.83	72.50	72.94	73.13	73.40	73.44
			Task-Targeted Artifact Correction	67.50	71.32	72.33	72.76	73.03	73.16	73.50	73.48	73.44

Table 2: Results for classification models.

Model	Metric	Reference	Mitigation	Q=10	Q=20	Q=30	Q=40	Q=50	Q=60	Q=70	Q=80	Q=90
FasterRCNN	mAP	35.37	Supervised Fine-Tuning	29.09	33.34	34.72	35.08	35.49	35.82	35.96	36.06	36.17
			None	20.35	30.03	32.59	33.43	34.04	34.31	34.73	34.93	35.25
			Off-the-Shelf Artifact Correction	28.45	31.86	33.10	33.85	34.05	34.47	34.70	34.77	34.71
			Task-Targeted Artifact Correction	31.43	33.85	34.29	34.81	34.81	34.97	35.01	34.88	34.81
FastRCNN	mAP	34.02	Supervised Fine-Tuning	28.01	31.94	33.08	33.56	33.88	34.17	34.42	34.44	34.66
			None	19.99	29.04	31.22	32.19	32.65	33.00	33.34	33.40	33.80
			Off-the-Shelf Artifact Correction	27.62	30.91	32.04	32.56	32.78	33.18	33.28	33.48	33.44
			Task-Targeted Artifact Correction	30.11	32.31	33.07	33.31	33.39	33.53	33.69	33.68	33.59
MaskRCNN	mAP	32.84	Supervised Fine-Tuning	26.32	30.48	31.79	32.21	32.55	32.83	33.11	33.20	33.32
			None	18.35	27.58	29.83	30.80	31.32	31.62	32.02	32.29	32.62
			Off-the-Shelf Artifact Correction	25.82	29.35	30.67	31.32	31.59	31.85	32.03	32.24	32.16
			Task-Targeted Artifact Correction	28.48	30.85	31.71	32.00	32.19	32.24	32.35	32.43	32.26
RetinaNet	mAP	33.57	Supervised Fine-Tuning	27.64	31.97	33.03	33.50	33.80	34.12	34.30	34.33	34.40
			None	18.76	28.23	30.63	31.59	32.27	32.57	32.88	33.02	33.42
			Off-the-Shelf Artifact Correction	26.74	29.90	31.24	31.87	32.19	32.60	32.86	33.02	32.93
			Task-Targeted Artifact Correction	29.66	31.86	32.73	32.97	32.98	33.13	33.24	33.23	33.09

Table 3: Results for detection models.

Model	Metric	Reference	Mitigation	Q=10	Q=20	Q=30	Q=40	Q=50	Q=60	Q=70	Q=80	Q=90
HRNetV2 + C1	mIoU	40.59	Supervised Fine-Tuning	34.76	37.35	38.74	38.78	39.27	39.75	39.98	39.86	39.96
			None	24.95	35.16	38.03	38.52	39.02	40.09	40.50	40.41	40.54
			Off-the-Shelf Artifact Correction	32.30	36.54	38.40	38.52	40.08	40.44	40.46	40.22	40.60
			Task-Targeted Artifact Correction	34.14	37.61	39.23	39.24	39.92	40.53	40.62	40.39	40.55
MobileNetV2 (dilated) + C1 (ds)	mIoU	29.52	Supervised Fine-Tuning	19.07	22.37	23.43	23.62	23.60	24.15	24.44	24.37	24.46
			None	13.92	24.03	27.13	27.75	27.73	28.86	29.37	29.35	29.43
			Off-the-Shelf Artifact Correction	21.17	25.27	27.31	27.16	29.14	29.32	29.26	29.06	29.54
			Task-Targeted Artifact Correction	24.74	27.37	28.44	28.33	29.19	29.56	29.54	29.38	29.52
ResNet101 + UPerNet	mIoU	41.08	Supervised Fine-Tuning	35.32	37.41	38.27	38.28	38.55	38.59	38.72	38.58	38.70
			None	26.14	36.70	39.45	39.81	39.55	40.47	40.98	40.97	41.07
			Off-the-Shelf Artifact Correction	33.90	37.39	39.12	39.38	40.32	40.58	40.78	40.79	41.04
			Task-Targeted Artifact Correction	35.82	38.67	39.96	39.98	40.22	40.79	40.97	40.91	41.00
ResNet101 (dilated) + PPM	mIoU	40.26	Supervised Fine-Tuning	31.86	35.45	36.73	36.94	36.91	37.33	37.67	37.55	37.65
			None	25.68	35.19	37.76	38.43	38.24	39.27	40.03	40.17	40.21
			Off-the-Shelf Artifact Correction	31.44	35.86	38.01	38.26	39.54	39.73	39.94	40.06	40.22
			Task-Targeted Artifact Correction	33.99	37.63	39.04	39.11	39.38	39.73	40.07	40.11	40.10
ResNet18 (dilated) + PPM	mIoU	36.65	Supervised Fine-Tuning	29.84	32.33	33.08	33.01	33.38	33.61	33.50	33.29	33.33
			None	21.16	31.99	34.72	35.36	35.41	36.16	36.56	36.60	36.59
			Off-the-Shelf Artifact Correction	28.64	32.59	34.56	34.53	35.96	36.21	36.29	36.25	36.64
			Task-Targeted Artifact Correction	31.69	34.55	35.80	35.80	36.12	36.50	36.66	36.54	36.60
ResNet50 + UPerNet	mIoU	39.21	Supervised Fine-Tuning	32.88	35.11	35.94	35.90	36.41	36.58	36.63	36.49	36.55
			None	24.29	34.78	37.34	37.71	37.70	38.57	39.12	39.13	39.16
			Off-the-Shelf Artifact Correction	31.83	35.52	37.20	37.26	38.44	38.67	38.87	38.86	39.12
			Task-Targeted Artifact Correction	34.36	36.94	38.17	38.07	38.55	38.93	39.14	39.06	39.09
ResNet50 (dilated) + PPM	mIoU	38.91	Supervised Fine-Tuning	32.26	35.33	36.04	36.04	36.53	36.75	36.93	36.71	36.92
			None	23.05	33.95	36.66	37.07	37.40	38.58	38.93	38.70	38.86
			Off-the-Shelf Artifact Correction	28.36	32.69	35.24	35.31	37.74	38.04	38.18	38.13	38.73
			Task-Targeted Artifact Correction	31.92	35.43	37.04	36.92	38.05	38.69	38.79	38.52	38.74

Table 4: Results for segmentation models.

Model	Metric	Reference	Mitigation	Q=10	Q=20	Q=30	Q=40	Q=50	Q=60	Q=70	Q=80	Q=90
Chai et. al	Patch Accuracy	93.83	Supervised Fine-Tuning	53.04	54.81	56.12	56.86	57.34	57.66	58.07	58.58	59.41
			None	49.87	49.46	49.48	49.86	50.33	50.74	51.12	51.83	53.26
			Off-the-Shelf Artifact Correction	50.56	50.20	50.34	50.63	51.14	51.63	52.26	53.22	54.12
			Task-Targeted Artifact Correction	51.99	52.46	52.70	52.98	53.21	53.33	53.48	53.58	53.28
Wang et. al	Accuracy	99.96	Supervised Fine-Tuning	89.70	96.69	98.41	99.04	99.44	99.54	99.65	99.88	99.91
			None	54.25	91.08	97.56	98.79	99.28	99.38	99.60	99.81	99.84
			Off-the-Shelf Artifact Correction	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00
			Task-Targeted Artifact Correction	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00

Table 5: Results for forensics models.

D.3. Reference Results

The following table gives the reference numbers of the pretrained weights as evaluated by our system on uncompressed images.

Model	Value
ImageNet Classification, Metric: Top-1 Accuracy	
ResNet 18	68.84
ResNet 50	75.31
ResNet 101	76.91
ResNeXt 50	76.99
ResNeXt 101	78.81
VGG 19	73.44
MobileNetV2	70.72
InceptionV3	77.33
EfficientNet B3	83.98
COCO Object Detection and Instance Segmentation, Metric: mAP	
FastRCNN	34.02
FasterRCNN	35.38
RetinaNet	33.57
MaskRCNN	32.84
ADE20k Semantic Segmentation, Metric: mIoU	
HRNetV2 + C1	40.59
MobileNetV2 (dilated) + C1	29.52
ResNet 18 (dilated) + PPM	36.65
ResNet 50 (dilated) + PPM	38.91
ResNet 101	41.08
ResNet 101 (dilated) + PPM	40.26
Forensics (dataset varies), Metric: Accuracy (exact formulation varies)	
Chai <i>et al.</i>	93.84
Wang <i>et al.</i>	99.96




Figure 27: Throughput results for all tested models.

E. Throughput

Although artifact correction is mentioned in prior works and presented here as a viable compression mitigation technique, we would be remiss if we did not note the slower throughput of these methods. In Figure 27 we show the training and inference throughput for batches of size 1 of both the artifact correction mitigation as well as the supervised fine tuning mitigation. These results are critical when considering which mitigation method is most viable for a particular application: although Task-Targeted Artifact Correction is more flexible, it comes with a cost in throughput.

Table 6: Transfer and multihead results. Reference indicates the performance of the pretrained weights on uncompressed images. **Best** result in bold, second best underlined.

Mitigation	Q=10	Q=20	Q=30	Q=40	Q=50
HRNetV2 + C1 , Reference: 40.59 mIoU (Semantic Segmentation)					
None	24.95	35.16	38.03	38.52	39.02
Off-the-Shelf Artifact Correction	32.30	36.54	38.40	38.52	40.08
Supervised Fine-Tuning	34.76	37.35	38.74	38.78	39.27
Task-Targeted Artifact Correction	34.14	<u>37.61</u>	<u>39.23</u>	<u>39.24</u>	39.92
MobileNetV2 Transfer	33.20	37.05	38.93	38.95	39.33
ResNet18 Transfer	33.77	37.44	39.22	39.21	39.25
Multihead (Three Model)	<u>34.38</u>	37.68	39.39	39.39	<u>39.72</u>
Faster RCNN , Reference: 35.37 mAP (Object Detection)					
None	20.35	30.03	32.59	33.43	34.04
Off-the-Shelf Artifact Correction	28.45	31.86	33.10	33.85	34.05
Supervised Fine-Tuning	29.09	33.34	34.72	35.08	35.49
Task-Targeted Artifact Correction	31.43	33.85	<u>34.29</u>	<u>34.81</u>	<u>34.81</u>
MobileNetV2 Transfer	30.05	33.04	33.86	34.35	34.48
ResNet18 Transfer	30.72	33.30	34.20	34.57	34.66
Multihead (Two Model)	<u>31.09</u>	33.39	34.19	34.67	34.68
Multihead (Three Model)	30.96	<u>33.41</u>	<u>34.29</u>	34.68	34.70
ResNet-101 , Reference: 76.91, Top-1 Accuracy (Image Classification)					
None	66.12	73.00	74.65	75.39	75.83
Off-the-Shelf Artifact Correction	67.91	73.64	75.09	75.84	76.23
Supervised Fine-Tuning	74.63	76.50	77.07	77.20	77.27
Task-Targeted Artifact Correction	<u>72.99</u>	<u>75.53</u>	<u>76.30</u>	<u>76.60</u>	<u>76.59</u>
MobileNetV2 Transfer	72.18	75.35	76.15	76.49	76.58
ResNet18 Transfer	71.80	75.05	76.00	76.40	76.49
ResNet-50 , Reference: 75.31, Top-1 Accuracy (Image Classification)					
None	63.43	71.20	73.23	74.10	74.43
Off-the-Shelf Artifact Correction	66.90	72.45	73.95	74.60	74.93
Supervised Fine-Tuning	73.18	75.46	76.02	76.24	76.36
Task-Targeted Artifact Correction	70.48	73.56	74.39	74.81	74.94
Multihead (Two Model)	<u>71.66</u>	74.14	74.90	<u>75.05</u>	75.10
Multihead (Three Model)	71.49	<u>74.23</u>	<u>74.96</u>	75.05	<u>75.15</u>

F. Multihead Results

Table 6 shows the raw numbers for our transfer and multi-head experiments.