

A Generative Adversarial Approach with Residual Learning for Dust and Scratches Artifacts Removal

Ionuț Mironică
mironica@adobe.com
Adobe Research Romania
Bucharest, Romania

ABSTRACT

Retouching can significantly elevate the visual appeal of photos, but many casual photographers lack the expertise to operate in a professional manner. One particularly challenging task for old photo retouching remains the removal of dust and scratches artifacts. Traditionally, this task has been completed manually with special image enhancement software and represents a tedious task that requires special know-how of photo editing applications.

However, recent research utilizing Generative Adversarial Networks (GANs) has been proven to obtain good results in various automated image enhancement tasks compared to traditional methods. This motivated us to explore the use of GANs in the context of film photo editing. In this paper, we present a GAN based method that is able to remove dust and scratches errors from film scans. Specifically, residual learning is utilized to speed up the training process, as well as boost the denoising performance.

An extensive evaluation of our model on a community provided dataset shows that it generalizes remarkably well, not being dependent on any particular type of image. Finally, we significantly outperform the state-of-the-art methods and software applications, providing superior results.

CCS CONCEPTS

• **Computing methodologies** → **Image processing; Computational photography; Reconstruction.**

KEYWORDS

Generative Adversarial Networks, Dust and Scratches artifacts removal

ACM Reference Format:

Ionuț Mironică. 2020. A Generative Adversarial Approach with Residual Learning for Dust and Scratches Artifacts Removal. In *Proceedings of The 2nd workshop on Structuring and Understanding of Multimedia heritAge Contents (ACM Multimedia 2020)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Photographic materials are rather unstable compared to other cultural objects and their degradation is much faster than those of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM Multimedia 2020, October 12–16, 2020, Online

© 2020 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

paintings, sculptures or architecture. Dust and scratches on original negatives distinctly appear as light or dark colored artifacts on a scan. These unsightly artifacts have become a major consumer concern during the last years. Dust and scratches [22] are two completely different phenomena, but they are often categorized together. "Dust" may refer to any small foreign body that is accidentally lying on the film surface that obscures the underlying emulsion image, or conversely, it can have a certain degree of transparency. On the other hand, in the case of a "scratch", the film material is missing. They can be superficial or they may penetrate the emulsion layer, completely removing parts of the scanned image. In this moment, the only technology that is able to reduce the effect of dust and scratches are scanners with dedicated hardware that detects the error areas [23]. This hardware is either prohibitively expensive or very slow to use. Without hardware assistance, dust and scratch removal algorithms generally resort to blurring, thereby losing image details. Different from existing software solutions as Photoshop or Lightroom [10] that alter the entire image appearance, we present a new alternative for dust and scratch detection that effectively differentiates between defects and image details and it is able to fill the gaps with realistic patches.

Deep neural networks (DNNs) and GANs [7, 18] have shown very promising results for various image restoration tasks [6, 9, 16, 25]. However, the design of network architectures remains a major challenge for achieving further improvements. In this paper, we first propose an algorithm that is specifically designed to remove dust and scratches errors from old film images. A novel adversarial learning-based model that can exploit the multi-scale redundancies of natural dust and scratches artifacts is proposed. Experimental results show that the proposed method obtains better results than any application and algorithm available for this particular problem.

The technical contributions of this work are summarized as follows:

- Development of a realistic scenario to test the proposed system. In this respect, we create a new benchmark dataset that represents a realistic, natural and challenging scenario in terms of diversity of film photos. The dataset contains pairs of photos with or without artifacts provided by the community.
- Identify and integrate appropriate architectural features and training procedures which lead to a boosted GAN performance for dust and scratches removal. The proposed steps of improvement include:
 - We introduce a new deep learning representation for the addressed problem of dust and scratches removal model; We find that residual learning can greatly benefit the CNN

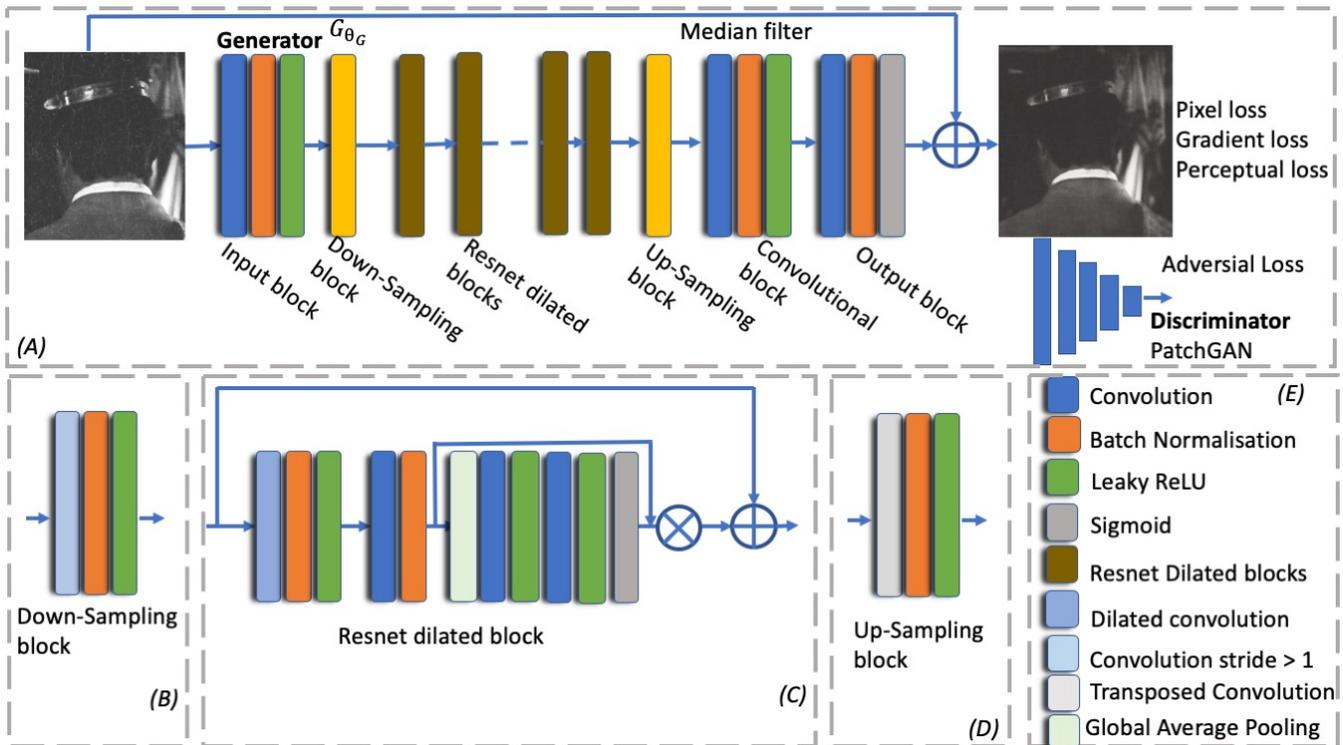


Figure 1: Overview of our proposed approach. Figure A presents the overview for the Generator and the Discriminator architectures, while the Figures B, C, D and E expose the technical details of the Generator architecture: the down-sampling and up-sampling blocks, Resnet dilated blocks with the attention mechanism and the agenda with all the used layers on different colors.

learning as they can not only speed up the training but also boost the denoising performance;

- Using a GAN structure, we enable the learning of photo retouching. To our knowledge, this is the first GAN applied to this problem. Furthermore, our approach scales with image resolution and generates no distortion artifacts in the image;
- We combine several novel loss functions that are able to preserve the visual quality of restored photos.
- We demonstrate its generality in terms of applications by applying it to a wide category of images;
- We achieve better performance than current state-of-the-art algorithms and applications.

The paper has the following structure: Section II reviews related work in the literature, identifying the main drawbacks and possible improvements, Section III details the proposed architecture, Section IV provides implementation details and a quantitative evaluation of the results, while Section V provides conclusions and identifies future work.

1 RELATED WORK

During the last years, most of the old photo restoration research targeted automatic colorization techniques [8, 17]. Even if the removal of artifacts on films has long been a challenge, not much research has been conducted in this particular field. Most of the approaches involved the breaking down of the problem into two parts: the automatic detection of regions with errors and the in-painting process. In the detection stage the scanned image pixels are categorized using a binary classification into two classes: “normal” and “flawed”, thus generating a detection mask [5]. The main problem of this approach is that a low performance of the detection may generate many displeasing digital artifacts. A second category of algorithms is inspired by the recent advances of deep learning image segmentation. They are trying to automatically correct the errors without any intermediary process of pixel categorization.

The following in-painting operation is assisted by the detection mask, which indicates where the restoration is necessary and where the image has to be left unedited. In-painting uses the information located near the boundaries of the degraded regions to restore continuous edges by connecting the truncated fragments together. In [19] an in-painting method based on manual selection of imperfections is introduced. They repair the imperfections by



Figure 2: Examples of images from the dataset. First row shows training images and the second row provides images with synthetic errors used in the training process and the third row contains images from the test dataset.

proposing an isotropic diffusion model extended with the notion of user-provided diffusion barriers. A fully automatic solution to detect and remove dust and scratches artifacts is presented in [2]. They use a local detection step that takes advantage of local contrast difference which enables the algorithm to separate defects from large image edges. Then, they generate a credibility map and apply a weighted bilateral filter for local repairs. An algorithm focused on correction and detection of vertical scratches is presented in [3]. These types of errors are commonly observed on photographic films due to a mechanical defect in the cameras, but solves only a small part of the scratch errors that may occur. In [27], the authors introduced an approach that exploits the image information inside the damaged region. By generating the artifact formation model, they make use of contextual information in the image and creates a color consistency constraint on dust to remove these artifacts.

Recently, there is an emerging interest in applying deep convolutional neural networks for solving various image-to-image translations problems [4, 8, 26]. A similar approach is proposed in [21]. They adapted a SegNet [1] encoder-decoder architecture trained with cross entropy loss that is able to automatically remove the dust and scratches artifacts. They also proposed a new dataset, but it contains only a limited amount of black and white photo patches.

However, the vast majority of proposed dust and scratches removal solutions are not able to completely restore the content of the photographs corrupted by noise. Following a careful analysis of

the existing methods and what they are missing in order to create a viable solution, we propose a series of innovative elements. Thus, we introduce an appropriate deep learning architecture and training procedures which lead to a boosted GAN performance for dust and scratches removal. We also concern about the size of the proposed model because our aim is to develop a solution for practical use. Thus, we carefully designed the model so that it has impressive performance using a limited number of learning parameters. The last, but not the least, we develop a realistic scenario by creating a new diverse and challenging benchmark dataset.

2 PROPOSED APPROACH

Inspired by the recent advances in semantic segmentation and GANs, the architecture of the proposed network is illustrated in Figure 1. The approach contains two parts: a generator that represents a feed-forward reconstruction CNN network and a discriminator which is similar to the PatchGAN architecture [12, 20]. Given an input image I deteriorated by dust and scratches, our goal is to estimate a noise-free image \hat{T} . Given training image pairs $((I_n, T_n), n = 1, \dots, N$, the following equations should be solved:

$$\hat{\Theta} = \underset{\Theta_G}{\operatorname{argmin}} = \frac{1}{N} \sum_{c=1}^N l(G_{\Theta_G}(I_n), T_n) \quad (1)$$

We first present the details of the generator network G_{Θ_G} followed by the loss functions.

2.1 Image Reconstruction Network

Our generator network contains several blocks: input, output, down-sampling, up-sampling and dilated residual blocks. The input block receives a $W \times H \times 3$ color image, while the output is bound by a sigmoid function that normalizes the network output values to $[0, 1]$ interval. After the input layer, there is a down-sampling layer that reduces the size of the input. It is followed by seven dilated residual blocks. The architecture of the dilated residual blocks is presented in Figure 1C. Each block contains a channel attention mechanism to the feature maps from convolutional layers such that different features are weighted differently according to global statistics of the activations. In order to take information from different scales, the blocks contain dilated convolutions which helps the network to reach a global contextual consistency in the spatial domain.

Each block uses Batch Normalization (BN) [11], which has become a common practice in deep learning to accelerate the training of deep neural networks. In the case of Deep Convolutional GANs, it was proven that applying BN to discriminator architectures can be very beneficial to stabilize the GAN learning and to prevent a mode collapse due to poor initialization [18]. Internally, BN preserves content-related information by reducing the covariance shift within a mini-batch during training. It uses the internal mean and variance of the batch to normalize each feature channel. After the generator, the discriminator is used in a form of 70×70 PatchGAN [12] fully convolutional architecture.

2.2 Residual Learning

The discriminator and the generator operate on the residual between the original and a median image. This allows both the generator and discriminator to concentrate only on the important sources of noise and be more specific to the dust and scratches artifacts. Since these regions are challenging to reconstruct and in-painting well, they correspond to the largest perceptual errors. This can also be viewed as subtracting a data dependent baseline which helps to reduce variance.

2.3 Training loss

In our algorithm we want to reduce the dissimilarity between the approximation \hat{T} image and the correct T image using the Pixel loss, the Gradient Loss and the Perceptual loss. We also want to improve the perceptual quality of the corrected image, for which we have included the Adversarial loss. The perceptual quality of an image \hat{T} represents the degree to which it looks like a natural image and has nothing to do with its similarity to any reference image.

Pixel Loss: We define the Pixel loss between T and \hat{T} images as:

$$L_{Pixel} = \|\hat{T} - T\| \quad (2)$$

where $\|\cdot\|$ represents the Mean Square Error (MSE) distance.

Gradient Loss: We penalize the difference between edges as the difference of the gradient operator along x and y axes.

$$L_{Gradient} = \|\nabla_x \hat{T} - \nabla_x T\| + \|\nabla_y \hat{T} - \nabla_y T\| \quad (3)$$

where ∇_x and ∇_y represent the gradients over x and y axis.

Perceptual Loss: Metrics such as MSE and SSIM only focus on low-level information in the image, while it is also necessary to

use some kind of higher-level information to improve the visual quality. The basic idea is to employ a content extractor. Then, if the enhanced image and the ground truth are similar, their corresponding outputs from the content extractor should also be similar. We use the VGG-19 network [20] as a content extractor. In particular, we define the context loss based on the output of the ReLU activation layers, by the absolute difference of the representations of the corrected image and the ground truth [13]:

$$L_{Perceptual/i/j} = \frac{1}{W_{i,j} H_{i,j} C_{i,j}} \sum_{x=1}^{W_{i,j}} \sum_{y=1}^{H_{i,j}} \sum_{z=1}^{C_{i,j}} \|\Phi_{i,j}(\hat{T}) - \Phi_{i,j}(T)\| \quad (4)$$

where $W_{i,j}, H_{i,j}, C_{i,j}$ describe the dimensions of the respective feature maps within the VGG-19 network and $\Phi_{i,j}$ represents the feature maps obtained by the j -th convolution layer in the i -th block. The final $L_{Perceptual}$ become:

$$L_{Perceptual} = \sum_i \|L_{Perceptual/i}(\hat{T}) - L_{Perceptual/i}(T)\| \quad (5)$$

Adversarial Loss: We add an adversarial loss to improve the realism of the corrected images. In this way, the generator network G tries to generate film images, while discriminator network D aims to distinguish between fake generated images and real-world images. We define an opponent discriminator D as a VGG [20] architecture while our proposed network acts the role of the generator. We used the adversarial loss function proposed in [25]:

$$L_{Adv} = \mathbb{E}_{x \sim p_x} [\log(D(x))] - \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))] \quad (6)$$

To summarize, our loss is defined as:

$$L = w_1 L_{Pixel} + w_2 L_{Gradient} + w_3 L_{Context} + w_4 L_{Adv} \quad (7)$$

where we empirically set the weights as $w_1 = 1, w_2 = 2, w_3 = 1$ and $w_4 = 0.01$ respectively throughout our experiments.

3 EXPERIMENTS

3.1 Implementation Details

Training data. Real corrected images with ground truth are difficult to obtain. To generate enough training data, we adopt a fusion of synthetic noise and real data as our train dataset. Using a set of 2,500 film photos without dust and scratches errors, we apply a diverse category of 100 synthetic artifacts. However, because we want to validate our algorithm on real-world scenarios, for test data we use a set of 500 photos with real noise that were provided and manually corrected by several professional photographers. Figure 2 contains several examples from the training dataset, synthetic images and the test dataset provided by the community.

Training details. Our implementation is based on PyTorch. Because we have a reduced number of training photos and their size are higher than $10MP$, we have extracted 128×128 patches from the training images. We train the model for 200 epochs using the Adam optimizer [14]. The base learning rate is set to 5×10^{-4} and we decrease the learning rate on each epoch using a standard weight decay that is set to l/N , where l represents the learning rate and N is the number of epochs. The weights are initialized as in [15]. Because BN usually requires a large batch size value [26], we set the batch size to a minimum of 32.

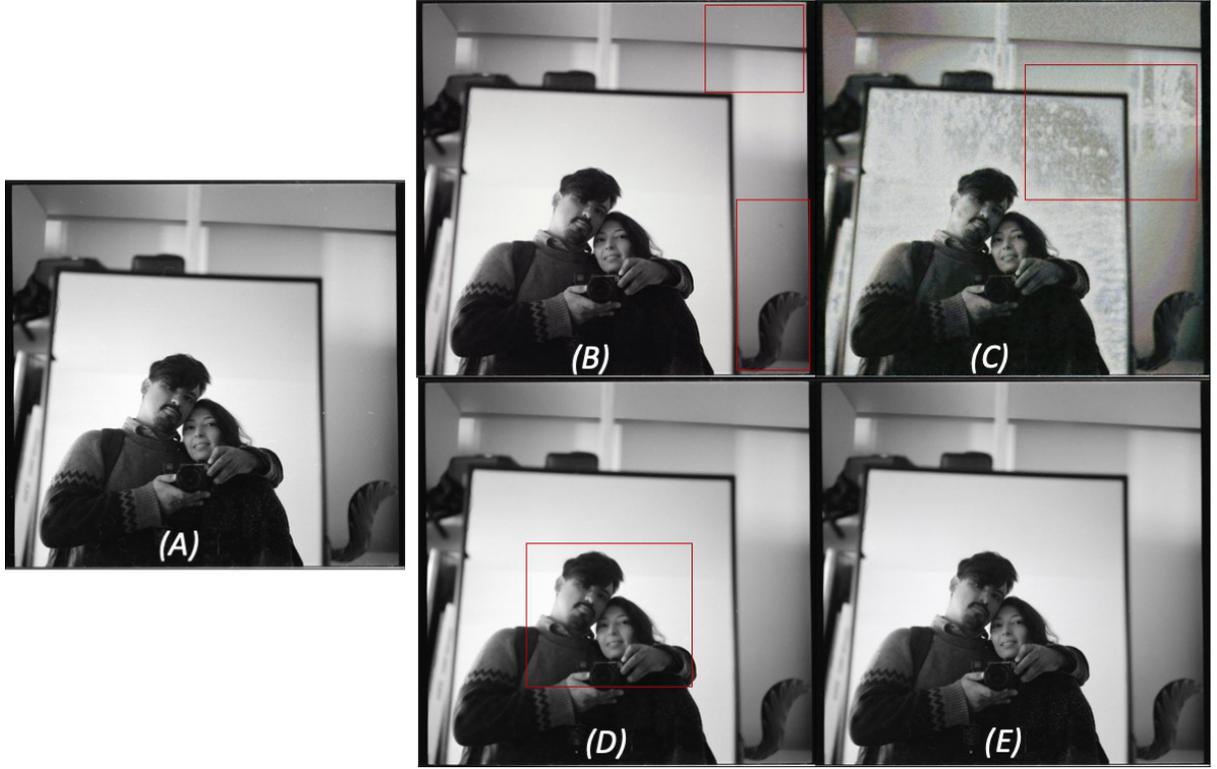


Figure 3: Comparison between the proposed approach and other state-of-the-art algorithms. Image (A) represents the input image that contains the dust&scratches artifacts. Images (B), (C) and (D) provide the output for the [12], [21] and [10]. Figure (E) shows the output provided by our approach.

To prevent overfitting, we perform image augmentation, i.e., we apply random transformations to input samples during network training on the fly. We apply random rotation and resizing transformations on both input images and synthetic errors. Also, we apply random intensities of the synthetic images.

3.2 Experimental Settings

Through extensive experiments, we qualitatively and quantitatively validate our model and learning framework. We provide full-reference quality assessment using two evaluation parameters: Peak Signal to Noise Ratio (PSNR) and the Structure SIMilarity index (SSIM). PSNR is defined by equation:

$$\text{PSNR} = 10 \log_{10} \frac{(2^d - 1)^2 WH}{\sum_{i=1}^W \sum_{j=1}^H (p[i, j] - p'[i, j])^2} \quad (8)$$

where d is the bit depth of pixel, W the image width, H the image height, and $p[i, j]$, $p'[i, j]$ is the i th-row j th-column pixel in the ground truth and corrected image respectively.

The SSIM [24] is a method for predicting quality of digital images. is calculated on various windows of an image. The measure between two windows x and y of the same size $N \times N$ is:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (9)$$

Table 1: Comparison with State-of-the-Art

Algorithm	PSNR	SSIM
Pix2Pix [12]	31.91	0.8
Strubel et al. [21]	25.07	0.65
Photoshop [10]	22.17	0.62
Proposed approach	35.12	0.92

where μ_x , μ_y , σ_x^2 and σ_y^2 represents the average and variances of x and y , σ_{xy} the correlation coefficient of x and y and c_1 and c_2 two variables that stabilize the division with weak denominator. Larger values of PSNR and SSIM indicate better performance.

3.3 Comparison with State-of-the-Art

We compare our method with other deep learning architectures, like GAN-based method Pix2pix-GAN [12] and the method proposed in [21]. For fair comparison, we fine-tune these models on our training dataset. Also, we compare our method with the state-of-the-art algorithm from Adobe Photoshop application [10]. Table 1 summarizes the results of all competing methods on our test dataset. The quality metrics include PSNR and SSIM.

As can be seen, the proposed model obtains better results than the other methods in terms of PSNR and SSIM index, and the differences in output are evident, as presented in Figure 3. The differences between the proposed approach and other algorithm are considerable: Pix2Pix is not able to remove most of the artifacts, the approach proposed in [21] removes the artifacts but generates color artifacts because it is not suitable for color images. The algorithm embedded in Photoshop is also able to remove most of artifacts but the generated image become blurry and many relevant details are removed from the generated image. Figure 3E contains the generated image using the proposed approach that is able to preserve the image details, while removing almost all the dust and scratches artifacts.

To conclude, using the proposed model significantly improves the PSNR and the SSIM, yielding much better results than the state-of-the-art algorithms.

4 CONCLUSIONS

The work presented in this paper improved the state-of-the-art for automatic dust and scratches removal on film photos using conditional adversarial networks. The proposed GAN architecture integrates techniques from the literature to ensure good training stability and to increase the contribution of the adversarial loss during training, which prevents the GAN from collapsing into blurry images.

The main contributions of this paper are: i) the design of a method for the restoration of dust and scratches artifacts ii) introduction of a benchmark dataset that create a realistic scenario to test the proposed system and iii) a thorough comparison with the state-of-the-art deep learning architectures and software applications.

Preliminary experiments with the proposed architecture showed good results for the restoration of other artifacts (i.e. noise removal, digital dust). As future work, we plan to investigate its extension to other single and multiple distortions. Also, we want to extend our research to old film videos, which can provide temporal information regarding dust and scratches.

REFERENCES

- [1] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39, 12 (2017), 2481–2495.
- [2] Ruth Bergman, Ron Maurer, Hila Nachlieli, Gitit Ruckenstein, Patrick Chase, and Darryl Greig. 2008. Comprehensive solutions for automatic removal of dust and scratches from images. *Journal of Electronic Imaging* 17, 1 (2008), 013010.
- [3] Bernard Besserer and Cedric Thiré. 2004. Detection and tracking scheme for line scratch removal in an image sequence. In *European Conference on Computer Vision*. Springer, 264–275.
- [4] Jingwen Chen, Jiawei Chen, H. Chao, and Ming Yang. 2018. Image Blind Denoising with Generative Adversarial Network Based Noise Modeling. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), 3155–3164.
- [5] Mohamed A Elgharib, François Pittié, and Anil Kokaram. 2013. Blotch and scratch removal in archived film using a semi-transparent corruption model and a ground-truth generation technique. *EURASIP Journal on Image and Video Processing* 2013, 1 (2013), 33.
- [6] Leonardo Galteri, Lorenzo Seidenari, Marco Bertini, and Alberto Del Bimbo. 2017. Deep generative adversarial compression artifact removal. In *In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4826–4835.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [8] Marc Górriz, Marta Mrak, Alan F Smeaton, and Noel E O'Connor. 2019. End-to-End Conditional GAN-based Architectures for Image Colourisation. *International Workshop on AI for Smart TV Content Production, Access and Delivery* (2019).
- [9] Yuanming Hu, Hao He, Chenxi Xu, Baoyuan Wang, and Stephen Lin. 2018. Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)* 37, 2 (2018), 1–17.
- [10] Adobe Inc. 2020. Photoshop. (2020). "http://www.adobe.com/Photoshop"
- [11] S. Ioffe and C. Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [13] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*.
- [14] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [15] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 136–144.
- [16] Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. 2018. MBLLEN: Low-light Image/Video Enhancement Using CNNs. *British Machine Vision Conference* (2018).
- [17] Kamyar Nazeri, Eric Ng, and Mehran Ebrahimi. 2018. Image colorization using generative adversarial networks. In *International conference on articulated motion and deformable objects*. Springer, 85–94.
- [18] A. Radford, L. Metz, and S. Chintala. 2016. Unsupervised representation learning with deep convolutional generative adversarial networks.
- [19] Manuel M. Oliveira Brian Bowen Richard and McKenna Yu-Sung Chang. 2001. Fast digital image inpainting. In *Proceedings of the International Conference on Visualization, Imaging and Image Processing (VIIP)*.
- [20] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations* (2015).
- [21] David Strubel, Marc Blanchon, and Fofi David. 2019. Deep learning approach for artefacts correction on photographic films. In *Fourteenth International Conference on Quality Control by Artificial Vision*, Vol. 11172. International Society for Optics and Photonics, 111720M.
- [22] Giorgio Trumpy and Rudolf Gschwind. 2015. Optical detection of dust and scratches on photographic film. *Journal on Computing and Cultural Heritage (JOCCH)* 8, 2 (2015), 1–19.
- [23] Chi-Pin Tsai and Hung-chou Chiu. 2001. Image scanner with image correction function. US Patent 6,327,056.
- [24] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [25] Kaixuan Wei, Jiaolong Yang, Ying Fu, David Wipf, and Hua Huang. 2019. Single image reflection removal exploiting misaligned training data and network enhancements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8178–8187.
- [26] Kaixuan Wei, Jiaolong Yang, Ying Fu, David Wipf, and Hua Huang. 2019. Single Image Reflection Removal Exploiting Misaligned Training Data and Network Enhancements.
- [27] Changyin Zhou and Stephen Lin. 2007. Removal of image artifacts due to sensor dust. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1–8.



Figure 4: Qualitative results of our method. First row shows input images and the second row images provides the output of the proposed algorithm

