Increasing the Stability of CNNs using a Denoising Layer Regularized by Local Lipschitz Constant

Anonymous Author(s) Affiliation Address email

Abstract

1	Recent studies have revealed that ConvNets are sensitive to small perturbations in
2	the input. This can cause fatal consequences in smart cars because of instability of
3	the road understanding module. In this paper, we propose an objective function
4	regularized by the local Lipschitz constant and train a layer for restoring images.
5	Our experiments on the GTSRB and the Caltech-Pedestrian datasets show that our
6	modular approach not only increases the accuracy of the classification ConvNets on
7	the clean datasets but it also increases the stability of the ConvNets against noise.
8	Comparing our method with similar approaches shows that it produces more stable
9	ConvNets while it is computationally similar or more efficient than these methods.

10 **1** Introduction

Understanding road is crucial for autonomous cars. Lane segmentation, pedestrian detection, traffic
sign recognition and car detection are some of the well known problems in this field. Convolutional
Neural Networks (ConvNets) have been successfully applied on these problems. Cireşan et al. (2012)
and Sermanet and Lecun (2011) proposed ConvNets that beat a human driver in classification of
traffic signs. Aghdam et al. (2015) also proposed a more accurate ConvNet with fewer parameters.
Similarly, Angelova et al. (2015) detected pedestrians using a cascade of ConvNets. Besides, Levi
et al. (2015) and Bittel et al. (2015) have proposed ConvNets for segmenting lanes in an image.

In real world applications, road understanding faces some practical challenges. For examples, if the weather is rainy or foggy, the camera mounted on the car may not acquire clean images. This may cause some artifacts on the image. In addition, based on the shutter speed, the image might be degraded by a motion if the car is being driven on an uneven route.

Despite the impressive results obtained by ConvNets, Szegedy et al. (2014) showed that small 22 perturbation of input images can alter their classification score. They study the reason by computing 23 24 the upper bound of the Lipschitz constant for each layer. The results suggest that instability of ConvNets might be due to the fact that they are highly non-linear functions. Hence, a small change 25 in the input may considerably change the output. Aghdam et al. (2016) empirically studied various 26 ConvNets trained on different datasets. In this work, they generated 1200 noisy images for each 27 sample in the test sets. The results showed that all the ConvNets in their experiments were unstable 28 to image degradation even when the samples were degraded using the Gaussian noise with $\sigma = 1$. 29 Similarly, Papernot et al. (2015) produced adversarial samples which were incorrectly classified by 30 the ConvNet. They produced these samples by modifying 4.02% of the input features. Goodfellow 31 et al. (2015) argued that the instability of ConvNets to adversarial examples is due to linear nature of 32 ConvNets. Based on this idea, they proposed a method for quickly generating adversarial examples. 33 They used these examples to reduce the test error. 34

³⁵ Gu and Rigazio (2014) stacked a denoising autoencoder (DAE) to their ConvNet and preprocessed ³⁶ the adversarial examples using the DAE before feeding them to the ConvNet. However, the resulting

Submitted to 29th Conference on Neural Information Processing Systems (NIPS 2016). Do not distribute.

network can be still attacked by new adversarial examples. Inspired by contractive autoencoders, they
 added a smoothness penalty to the objective function and trained a more stable network.

Contribution: In this paper, we train channel-wise filters for restoring images. Our objective function tries to locally reduce the nonlinearity of the restoration module. To be more specific, we train a convolution layer with 3 filters to restore an image as accurate as possible but also it generates nearly identical outputs for all perturbations in small neighborhood of an image. Our experiments on pedestrian detection and traffic sign classification datasets show that this lightweight restoration layer is able to effectively tackle with noisy images in real-time compared with other methods.

45 2 Proposed Method

⁴⁶ Denoting the *softmax* layer of a ConvNet (*i.e.* the last layer in a classification ConvNet) by $\mathcal{L}_{\theta}(x)$, ⁴⁷ the general idea is to find a parameter vector θ such that:

$$\forall_{\|\nu\| \le \epsilon} \mathcal{L}_{\theta}(x+\nu) = \mathcal{L}_{\theta}(x) \tag{1}$$

where ν is a noise vector whose magnitude is less than T. Solving the instability of ConvNets against 48 noise using the above formulation may require to add new terms to the loss function or generate 49 thousands of noisy samples. Instead, we propose a modular approach consisting of two ConvNets. 50 The first ConvNet is a denoising layer that we are going to mention in this section. The second 51 ConvNet is the one that is originally trained on training samples. In our approach, we connect the 52 denoising network to the classification network and feed the images to the denoising network. Our 53 aim is to train a denoising ConvNet that is able to restore the original image as accurate as possible 54 and it produces identical results for all the samples located within radius r from the current sample. 55 Mathematically, we are looking for two sets of parameters θ_1 and θ_2 such that: 56

$$\forall_{\|\nu\| \le T} \mathcal{L}_{\theta_1}(\mathcal{F}_{\theta_2}(x+\nu)) = \mathcal{L}_{\theta_1}(\mathcal{F}_{\theta_2}(x)).$$
(2)

where θ_1 indicates the parameters of the classification ConvNet and θ_2 denotes the parameters of the denoising ConvNet. The parameters θ_1 is already available by training the classification ConvNet on the training samples. Then, our goal is to find a function $\mathcal{F} : \mathbb{R}^n \to \mathbb{R}^n$ that is able to map all points around $x \in \mathbb{R}^n$ to the same point. If we can find such a function, the sample x and all its adversarial examples will be mapped to the same point. Then, the classification ConvNet will be able to produce the same output for all adversarial examples.

In contrast to Jain and Seung (2009), we do not restrict \mathcal{F} to the Gaussian noise. Furthermore, 63 contrary to Burger et al. (2012) and Hradi (2015) that model \mathcal{F} using 16M and 4.5M parameters, 64 our approaches requires determining only 75 parameters. From one perspective, \mathcal{F} can be seen as 65 an associative memory that is able to memorize patterns $X = \{x_1 \dots x_i \dots x_j\}, x_i \in \mathbb{R}^N$ in our 66 dataset and map every sample $\{x_i + \nu | \|\nu\| \le \epsilon\}$ to x_i . Here, x_i is an image patch and X is the set of 67 all possible image patches collected from all classes of objects in our dataset. Figure 1 illustrates 68 our approach. Our approach can be considered as a layer which is later connected to the input of a 69 classification layer and its aim is to reduce the effect of noise. 70

The two layers shown in this figure have identical architectures and they share all their parameters. Furthermore, we need the two layers during the training phase and we will only use one of them in the test phase. The layer consists of 3 convolution filters of size 5×5 which are separately applied on the red, green and blue channels of the noisy image. Also, the result of convolutions are passed through a *ReLU* activation function and they are concatenated in order to create the final image.

It should be noted that the noise generation module in Figure 1 is only used during the training phase.
In the test phase, the noise generation module is omitted. In this paper, we have only concentrated on
additive noise. The noise generation module creates noisy patterns with various probability density
functions. Five examples of the probability density functions have been shown in Figure 1.

⁸⁰ Given a set of clean image patches $\mathcal{X}_{clean} = \{x_{clean}^1, \dots, x_{clean}^N\}$ and their noisy versions $\mathcal{X}_{noisy} = \{x_{noisy}^1, \dots, x_{noisy}^N\}$, restoration ConvNets are usually trained by minimizing the Euclidean loss ⁸² function(Dong et al., 2014; Svoboda et al., 2016; Hradi, 2015; Burger et al., 2012):

$$E = \frac{1}{N} \sum_{i=1}^{N} \|x_{clean}^{i} - \mathcal{F}_{\theta}(x_{noisy}^{i})\|^{2} + \lambda \|\theta\|^{2}$$
(3)

where θ is the set of network weights and biases and λ is the regularization coefficient. The objective

⁸⁴ of this function is to train a restoration ConvNet which is able to restore clean images from noisy



Figure 1: The proposed layer for modelling \mathcal{F} in (2). C(s, k) shows a ReLU layer containing s filters of size $k \times k$. Probability density functions used for generating noise are shown in the left.

85 inputs as accurate as possible. However, we argue that training a ConvNet using the above loss

86 function could be accurate if χ_{clean} is clean in practice. But, this is not usually the case in datasets collected for road understanding problems. 87

This is illustrated in Figure 2 on the samples 88

from GTSRB (Stallkamp et al., 2012) dataset. 89

The green rectangles show contradictory patches 90

in each column. For example, the green patches 91

related to the speed limit sign are pointing to 92

the same pattern. However, one of these patches 93 are degraded due to camera motion. In the sec-94

ond and third columns, shadow and excessive 95

ambient light on the patches has caused the con-96

tradiction. In the last two columns, there are 97



Figure 2: Unclean training samples with contradictory patches (best viewed in color).

some irregularities due to camera noise. X_{noisy} 98

is usually generated from \mathcal{X}_{clean} . However, because of the above reasons, it might not be practical to 99 train the ConvNet using (3) due to contradictions in the database. 100

To tackle with this problem, we propose to add a new term to the objective function encouraging the 101 layer to learn a mapping in which $\|\mathcal{F}(x_{clean}) - \mathcal{F}(x_{noisy})\|$ is less than $\|x_{clean} - x_{noisy}\|$. This 102 is analogous to *locally* reducing the Lipschitz constant of the layer. Our final objective function is 103 defined as follows: 104

$$E = \frac{1}{N} \sum_{i=1}^{N} \left[w_1 \| x_{clean}^i - \mathcal{F}_{\theta}(x_{noisy}^i) \|^2 + w_2 \frac{\| \mathcal{F}(x_{clean}^i) - \mathcal{F}(x_{noisy}^i) \|}{\| x_{clean} - x_{noisy} \|} \right]$$
(4)

where N is the total number of the images. Also, $\mathcal{F}(x_{clean}^i)$ and $\mathcal{F}(x_{noisy}^i)$ are computed at the 105 same time using the top and bottom layers in Figure 1, respectively. Moreover, the noisy patterns are 106 generated on the fly. That said, the degradation module accepts a mini-batch of clean images and 107 outputs their degraded version along with identity mapping data. This helps the network not only 108 learn to restore noisy patches but also apply identity mapping on clean patches. 109

One the one hand, our layer learns to restore images where intensity values is in interval [0, 1]. On 110 the other, we initialize our filters close to averaging filters. Therefore, the output never becomes a 111 negative number. Since our approach is only one layer consisting of convolution operators and ReLU 112 functions, it is a linear operator which is applied on the input image. Formally, conditions f(kx) =113 kf(x) and f(x + y) = f(x) + f(y) hold in our approach. Taking into account one convolution kernel, $\mathcal{F}(x_{clean}^i) - \mathcal{F}(x_{noisy}^i)$ can be simplified as $Wx_{clean}^i - Wx_{noisy}^i = W(x_{clean}^i - x_{noisy}^i)$. Consequently, the second term in (4) is minimized by reducing ||W||. This is similar to regularizing 114 115 116 the objective function with an adaptive weight analogous to the difference between clean and noisy 117 samples. For this reason, we do not add other regularization terms to our objective function. In 118 terms of Lipschitz constant, the second term reduces the slope of the hyperplane represented by each 119 convolution kernel. Moreover, it helps to reduce the effect of contradictory patches. 120

121 **3 Experiments**

We carry out our experiments on German Traffic Sign Recognition Benchmark (GTSRB) (Stallkamp et al., 2012) and Caltech-Pedestrian(Dollár et al., 2009) datasets. These datasets have some important characteristics. First, they have been collected considering real scenarios (*e.g.* shadow, lightening, occlusion, camera motion) and they contain many degraded images. Second, the imaging device are noisy and they produce artifacts on the acquired images. Third, the resolution of images are low. Therefore, a slight change in the image may affect the classification score.

Note that we *do not* apply zero-padding in the training phase to avoid the impact of border effect on the loss function. Besides, the input is normalized to [0, 1]. All the weights in our layer are initialized using the normal distribution with mean value set to 1 and standard deviation set to 0.2. Taking into account the fact that each activation of the layer must be in interval [0, 1], the initial weights are divided by 25 in order to make the results of convolution kernels close to this interval. After we have the layer trained, zero-padding is applied on the input (the size of padding is 2 for each side).

Exploratory analysis: To evaluate the restoration accuracy of our method, we generate 150 noisy 134 images for each sample in the test set. Generating a noisy pattern is done in several steps. First, 135 136 we randomly select the *uniform* or the *normal* distribution with probability 0.5. Then, a noisy pattern is generated with $\mu = 0$ and $\sigma = U(0.5, 15)$ if the normal random number generator 137 is selected. In the case of uniform random number generator, the noisy pattern is generated in 138 interval [-3U(0.5, 15), 3U(0.5, 15)]. Next, the noisy pattern is sparsified with probability 0.25. The 139 sparsification is done by generating a binary mask using the *binomial* distribution with n = 1 and 140 p = U(0.5, 1.0). It is worth mentioning that we set the seed of random number generators to an 141 identical value for all methods. The noisy samples are fed into the layer and the *peak signal to noise* 142 *ratio* of the restored image is computed using the following equation: 143

$$psnr = 10 \log_{10} \left(\frac{255^2}{\frac{1}{HW} \sum_{m=i}^{H} \sum_{n=j}^{W} (x - x')^2} \right).$$
(5)

In the above equation, x is the clean image and x' is the noisy/restored image (after re-scaling to [0, 255]). In addition, we also study how the Lipschitz constant of our layer changes locally. This is done by computing $\|\mathcal{F}(x_{clean}^i) - \mathcal{F}(x_{noisy}^i)\|$ and $\nu = \|x_{clean} - x_{noisy}\|$. To compare our results with other similar methods, we also restored the images using the *bilateral*(d=5 and $\sigma_1 = \sigma_2 = 9$), the *median*(5 × 5) and *Gaussian*(5 × 5) filtering approaches. Figure 3 illustrates the scatter plot of the PSNR study (left) and the Lipschitz study (right) superimposed with a polynomial fitted on the data.

According to the results, both Gaussian and median filtering approaches are not able to restore image 150 accurately. This is due to the fact that objects in the GTSRB and the Caltech-Pedestrian datasets are 151 152 represented using low resolution images. On the one hand, details of objects are mainly determined using high frequency pixels. On the other hand, these pixels are close in the case of these low 153 resolution images. As the result, Gaussian and median filtering approaches oversmooth the images 154 which degrades the edges of objects. For this reason, PSNR of the filtered image is much lower than 155 the PSNR of the noisy image. In contrast, bilateral filtering preserves the edges and this is the main 156 reason that it has a higher PSNR compared with these two methods. Moreover, bilateral filtering 157 restores image with higher PSNR when the PSNR of the noisy image is less than 35. 158

The Lipschitz study shows that Gaussian and median filtering approaches produce close results regardless of the magnitude of noise. In contrary, images restored by bilateral filtering are scattered at a distance which is approximately similar to the distance of the noisy image from the clean image. We are looking for a filtering approach which is able to accurately restore images with smaller Lipschitz constant. Consequently, none of the three approaches are appropriate for our purpose.

However, the filter learned by our approach has a trade off between accuracy and the Lipschitz constant. Looking at the PSNR values, we observe that, on average, it is more accurate than these three methods. Besides, its Lipschitz constant is approximately linear. More importantly, the Lipschitz constant of our filter is less than 1 which means that restored images become closer after being filtered by our layer. Finally, we observe that the Lipschitz constant is very stable with very low variation in our approach. This means that, the filter learned by our objective function is not sensitive to the variations of input image.

We further analyze our filters in the frequency domain using the Fourier transform. Figure 4 illustrates the frequency response of our filters along with the Gaussian filter. First, our filters have higher



Figure 3: PSNR (left) and Lipschitz analysis (right) of the the Gaussian, median, bilateral and our approaches (best viewed in color and electronically).

- response to low frequencies than the Gaussian filter. For this reason, it passes some of the details in
- the image more than Gaussian filter. Second, they also have higher responses in very high frequencies. This helps our filters to preserve edges more than the Gaussian filter.



Figure 4: Comparing our filter with Gaussian filter in the frequency domain.

175 176

177

Quantitative analysis: We pick the ConvNet in Angelova et al. (2015) for detecting pedestrians and the ConvNet in Aghdam et al. (2015) for classification of traffic signs. We connect our learned

restoration filters to these ConvNets and fine-tune them on the original dataset (we do not augment
the dataset with noisy images). Then, the ConvNets are tested using noisy test sets. We repeat this
procedure (fine tuning the ConvNets) on Gaussian, median and bilateral filtering as well.

The noisy test sets are created by generating 1050 Guassian noise patterns with $\sigma \in \{0.3, 1, 2, 3, 4, 8, 10\}$ for each sample (150 images per each value of σ). Then, these noisy samples are fed to the above ConvNets (after connecting our layer to these ConvNets). To generate the same noisy samples for all methods in our experiment, we always seed the random number generator with a fixed value. Table 1 and Table 2 show the results on the GTSRB and the Caltech-Pedestrian datasets.

We observe that adding a Guassian or median3x3 layer to the GTSRB ConvNet increases the
classification accuracy of on clean images. This is due to the fact that some of the test samples might
be noisy for the reasons we discussed in Section 2. The Gaussian layer reduces the effect of noise

	accuracy (%) for different values of σ								
network	clean	0.3	1	2	3	4	8	10	overall
original	99.06	98.56	98.56	98.55	98.52	98.48	98.20	97.93	98.48
gaussian3x3	99.22	98.72	98.72	98.71	98.68	98.64	98.36	98.13	98.65
gaussian5x5	99.22	98.71	98.71	98.70	98.68	98.65	98.42	98.23	98.66
median3x3	99.15	98.66	98.66	98.65	98.63	98.60	98.38	98.19	98.62
median5x5	98.94	98.42	98.42	98.39	98.36	98.32	98.04	97.83	98.34
bilateral1	98.99	98.49	98.49	98.48	98.46	98.45	98.27	98.13	98.47
bilateral2	96.94	96.49	96.48	96.48	96.44	96.42	96.28	96.17	96.46
our filter	99.31	99.31	99.31	99.30	99.28	99.26	99.03	98.84	99.21

Table 1: Accuracy of the GTSRB ConvNet obtained by degrading the *test images* in the original dataset using a Gaussian noise with various values of σ .

	accuracy (%) for different values of σ								
network	clean	0.3	1	2	3	4	8	10	overall
original	92.39	91.97	91.97	91.97	91.95	91.92	91.67	91.49	91.92
gaussian3x3	92.36	91.89	91.89	91.87	91.85	91.80	91.61	91.48	91.84
gaussian5x5	92.01	91.52	91.52	91.49	91.46	91.43	91.23	91.12	91.47
median3x3	92.61	92.16	92.16	92.16	92.17	92.16	92.07	92.02	92.19
median5x5	92.18	91.67	91.67	91.66	91.64	91.61	91.46	91.34	91.65
bilateral1	92.74	92.27	92.26	92.25	92.25	92.22	92.13	92.03	92.27
bilateral2	92.27	91.82	91.82	91.83	91.82	91.81	91.83	91.84	91.88
our filter	92.86	92.84	92.83	92.81	92.78	92.76	92.56	92.46	92.74

Table 2: Accuracy of the Caltech-Pedestrian ConvNet obtained by degrading the *test images* in the original dataset using a Gaussian noise with various values of σ .

and it increases the accuracy of the ConvNet. Similarly, median3x3 and bilateral1 filtering increases 189 the accuracy of the Caltech-Pedestrian ConvNet on clean samples. However, while Gaussian filtering 190 works well on the GTSRB dataset it does not increase the accuracy on the Caltech-Pedestrian dataset. 191 Likewise, bilateral filtering improves the accuracy on the Caltech-Pedestrian dataset but they do not 192 increase the accuracy on the GTSRB dataset. Notwithstanding, the filters learned by our method 193 produce the most accurate results on both datasets. In addition, our layer produces a ConvNet with 194 highest stability against noise compared with approaches with similar computational complexity. 195 In fact, the computational complexity of our layer is identical to the Gaussian 5x5 and its less than 196 bilateral and median filtering approaches. 197

Analyzing results: Figure 5 illustrates some of the samples that are classified incorrectly by the original ConvNet but they are classified correctly after being smoothed by our method. The original image inside the green rectangle is degraded by shadow. Our layer filters the edges of the object and reduces the effect of shadow on the edges. The background of the image inside the red rectangle is smoothed by the layer. Edges in the original image inside the yellow rectangle has Bayer like pattern because of excessive lightening in the background. This effect is reduced by our filter. Finally, a general filtering is applied on the image inside the blue rectangle and makes it smoother. In sum, our method increases the accuracy by improving degraded edges and smoothing background noise.



Figure 5: Images that are correctly classified after being filtered by our layer. Left to right: Original image, difference with restored, restored image and normalized difference.

205

206 4 Conclusion

In this paper, we proposed a lightweight approach for increasing stability of ConvNets. Our method 207 trains a ReLU layer containing 3 channel-wise filters. We proposed a new objective function 208 consisting of the sum of square error penalized by the local Lipschitz constant of the filters. We 209 showed that the Lipschitiz constant in this particular configuration act as an adaptive L_2 regularizer. 210 Our experiments on the GTSRB and the Caltech-Pedestrian datasets showed that this approach 211 increases the accuracy of the original ConvNets on the clean test sets. In addition, the stability of 212 ConvNets increased against noise. Besides, since it is a modular approach, we do not need to train 213 a large ConvNet using thousands of noisy samples to increase the stability. Rather, we train the 214 classification ConvNet on the clean dataset. Then, we train our restoration layer on the noisy training 215 set. Finally, the classification ConvNet is fine-tune for one epoch using the clean training set. 216

217 **References**

- Aghdam, H. H., Heravi, E. J., and Puig, D. (2015). Recognizing Traffic Signs using a Practical Deep
 Neural Network. In *Second Iberian Robotics Conference*, Lisbon. Springer.
- Aghdam, H. H., Heravi, E. J., and Puig, D. (2016). Analyzing the Stability of Convolutional Neural
 Networks Against Image Degradation. In *Proceedings of the 11th International Conference on*
- 222 Computer Vision Theory and Applications.
- Alhussein Fawzi, Omar Fawzi, and Pascal Frossard (2015). Analysis of classifiers' robustness to
 adversarial perturbations. (2014):1–14.
- Angelova, A., Krizhevsky, A., View, M., View, M., Vanhoucke, V., Ogale, A., and Ferguson, D.
 (2015). Real-Time Pedestrian Detection With Deep Network Cascades. *Bmvc2015*, pages 1–12.
- Bittel, S., Kaiser, V., Teichmann, M., and Thoma, M. (2015). Pixel-wise Segmentation of Street with
 Neural Networks. pages 1–7.
- Burger, H. C., Schuler, C. J., and Harmeling, S. (2012). Image denoising Can plain neural networks
 compete with BM3D .
- Cireşan, D., Meier, U., Masci, J., and Schmidhuber, J. (2012). Multi-column deep neural network for
 traffic sign classification. *Neural Networks*, 32:333–338.
- Dollár, P., Wojek, C., Schiele, B., and Perona, P. (2009). Pedestrian detection: A benchmark. 2009
 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops,
 CVPR Workshops 2009, pages 304–311.
- Dong, C., Loy, C. C., and He, K. (2014). Image Super-Resolution Using Deep Convolutional
 Networks. *arXiv preprint*, 8828(c):1–14.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. (2015). Explaining and Harnessing Adversarial
 Examples. *Iclr 2015*, pages 1–11.
- Gu, S. and Rigazio, L. (2014). Towards Deep Neural Network Architectures Robust to Adversarial
 Examples. *arXiv*:1412.5068 [cs], (2013):1–9.
- 242 Hradi, M. (2015). Convolutional Neural Networks for Direct Text Deblurring. *Bmvc*, (1):1–13.
- Jain, V. and Seung, S. (2009). Natural Image Denoising with Convolutional Networks. pages 769–776.
- Levi, D., Garnett, N., and Fetaya, E. (2015). StixelNet: a deep convolutional network for obstacle detection and road segmentation. *Bmvc*, pages 1–12.
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. (2015). The
 Limitations of Deep Learning in Adversarial Settings.
- Sabour, S., Cao, Y., Faghri, F., and Fleet, D. J. (2015). Adversarial Manipulation of Deep Representations. *arXiv preprint arXiv:1511.05122*, (2015):1–10.
- Sermanet, P. and Lecun, Y. (2011). Traffic sign recognition with multi-scale convolutional networks.
 Proceedings of the International Joint Conference on Neural Networks, pages 2809–2813.
- Stallkamp, J., Schlipsing, M., Salmen, J., and Igel, C. (2012). Man vs. computer: Benchmarking
 machine learning algorithms for traffic sign recognition. *Neural Networks*, 32:323–332.
- Svoboda, P., Hradis, M., Marsik, L., and Zemcik, P. (2016). CNN for License Plate Motion
 Deblurring.
- 257 Szegedy, C., Zaremba, W., and Sutskever, I. (2014). Intriguing properties of neural networks.