

Technical Disclosure Commons

Defensive Publications Series

May 31, 2019

PRINT DEFECTS GENERATOR USING VARIATIONAL AUTOENCODER FOR INSPECTION SYSTEMS

HP INC

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

INC, HP, "PRINT DEFECTS GENERATOR USING VARIATIONAL AUTOENCODER FOR INSPECTION SYSTEMS",
Technical Disclosure Commons, (May 31, 2019)
https://www.tdcommons.org/dpubs_series/2244



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Print Defects Generator Using Variational Autoencoder for Inspection Systems

Abstract

Deep Learning (DL) applications usually need a large amount of labeled data to perform successfully. Data augmentation is a well-known method to avoid overfitting during a training process by increasing the number of training images with simple transformations that do not modify the true nature of the sample. However, these modifications are usually handcrafted and problem-dependent as it is difficult to determine which one is better. For these reasons, there is an increasing interest in studying automatic mechanisms to generate new instances using the available data. Variational Autoencoders (VAE) are powerful generative models. They're a modern take on autoencoders—a type of network that aims to encode an input to a low-dimensional latent space and then decode it back—that mixes ideas from deep learning with Bayesian inference. VAEs are usually easier to train compared to Generative Adversarial Nets (GANs), although the samples are sometimes blurrier. In this work, we present VAE based print defect generator that requires minimum (few dozen, e.g. 20-30) samples per defect type. First, we design a small compact Fully Convolutional Network (FCN) inspired by "Small U-Net" which is encoder-decoder type network used for detecting vehicles in a video stream [https://github.com/vxy10/p5_VehicleDetection/blob/master/README.md]. This is mainly achieved by using Separable convolutions which separates the learning of spatial features and channel-wise features and using global max pooling instead of fully-connected layers. In order to improve further the generalization capabilities of our network (which is called CompactSmallU-Net) we use simple data augmentation (flipping, color channels swap, etc.), drop-out and batch normalization while training the VAE. Next, we create a database (DB) of images with real looking defects by using seamless cloning [<https://www.learnopencv.com/tag/seamless-cloning/>] to overlay the generated defects on top of (defects free) images. Finally, Siamese based Single Shot Detector (SSD) can be used to train a press inspection system (AAA- Automatic Alert Agent) for detecting and classifying print defects (e.g., drips & dents) given pair of RGB images (digital and scanned) from the DB.

Problem statement

The main challenge in any inspection system is to maintain low false alarm and miss detect rates. In our case the problem is even more complicated as we are using low cost scanners (integrated into our presses) which create artifacts that erroneously classified as true defects. By using DL object detection algorithm called Siamese SSD (Single shot detector) the false alarms rate was reduced significantly from about ~0.5% to 0.005%.

In order to reach our goal of smart press that will enable the operator to receive an accurate information on the relevant part (e.g. consumable) to replace, defect classification module should be added on top of our detection module. Fig. 1 presents some defect types (drips, dents, wrinkles, paper cuts, etc.) with the corresponding troubleshooting. The main challenge is collecting large amounts of samples per defect type, which can be expensive and time consuming, but still needed in order to build models that can generalize better. Data augmentation is a common method that artificially creates new data based on modifications of the existing data (rotation, flipping, etc.). However, the heuristics underlying these modifications are very dependent on the specific problem at hand.

Our solution

The key idea is to automatically generate new samples that look similar to the training data, by using Variational Autoencoders (VAE). VAE is a neural network that learns to reproduce input by mapping it into a smaller low-dimensional latent space which “forces” the network to learn a more compact representation of

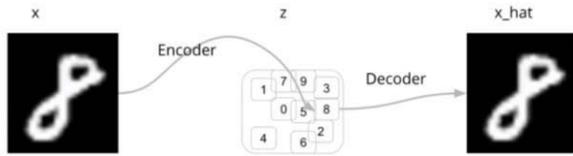


Figure 2 – An autoencoder maps input into compressed representation and then decoding it back

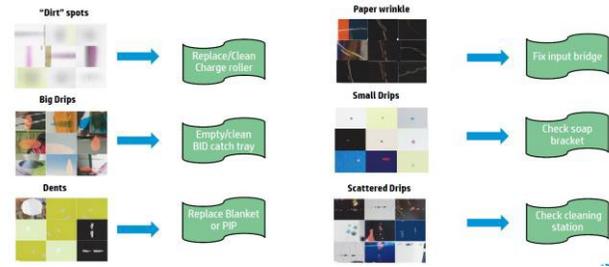


Figure 1- Example of several defect types and corresponding troubleshooting

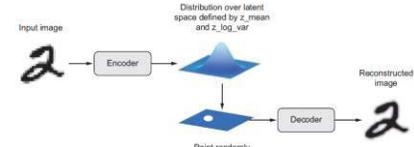


Figure 3 – A VAE maps an input image into probability distribution, followed by sampling from it to generate new samples [2]

the data (Fig. 2). A VAE, instead of compressing its input image into a fixed code in the latent space (like in “regular” auto-encoders) maps it into parameters of probability distribution (e.g., mean and variance of Gaussian). New samples can be generated by sampling from the latent space back to the original input (Fig. 3). The latent space is continuously meaningful as the decoder learns that not only is a single point in latent space referring to a sample of that class, but all nearby points refer to the same as well. The parameters of a VAE are trained with two loss functions: a reconstruction loss (usually mean squared error) forces the decoded samples to match the initial inputs, and a regularization loss to helps learn encodings, *all* of which are as close as possible to each other while still being distinct, allowing smooth interpolation, and enabling the construction of *new* samples. In order to reduce overfitting when having relatively small amount of data (e.g., few dozen) we have designed a small FCN network (*CompactSmallU-Net*). The network (Fig. 4) has the same architecture as small-Unet but without the skip connections and with some modifications. First, convolution layers were replaced with separable convolutions. Next, two new convolutional + global max pooling layers (VAR and AVG) with 256 filters were added to the end of the encoder (instead of fully connected layers which are common in VAE networks). Standard data augmentation (flipping, color channels swap, etc.), drop-out (10%) and batch-normalization are used to reduce further overfitting. After training the network with several dozen samples of specific print defect, never-before-seen defects are generated by sampling and decoding points from the latent space.

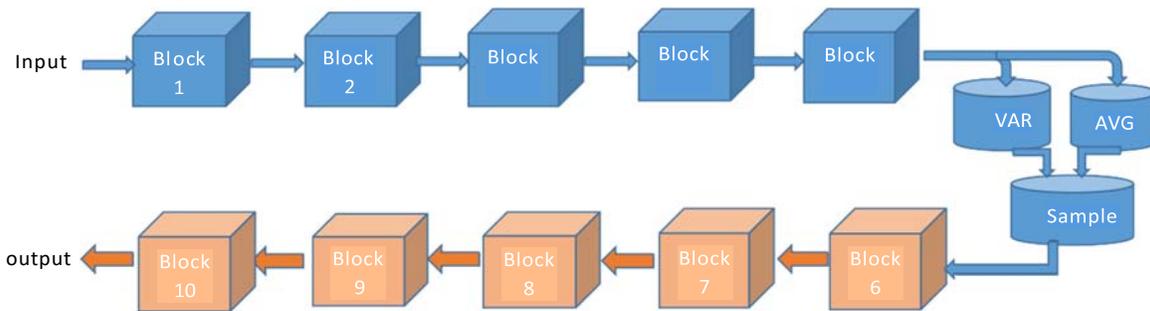


Figure 4 – CompactSmallU-Net. The network has about 180,000 parameters. All convolutions are 3X3.
 Blocks 1-5 are composed from (Sep_Conv-Batch norm) x 2 - Max Pool-Drop Out. # filters per block: 8,16,32,64,128
 Blocks 6-10 are composed from (Sep_Conv-Batch norm) x 2 - Up_Sampling. # filters per block: 64,32,16,8,3

The generated defects are used for creating large DB of images (e.g. few thousands) with real-looking print defects. This is done by using seamless cloning algorithm that overlays the generated defects on top of defect-free images while making a composition that looks seamless and natural. Finally, Siamese based SSD which uses a single network for both detection and classification can be trained with images from the generated DB.

Evidence the solution works

We have assessed the performance of the proposed method on two defects types: small drips and dents. Drips are random scattered drops on the printed page, while dents are mechanical damage to the press. We have collected small dataset of 22 Dents and 32 Drips as shown in Fig. 5 and 6, respectively. The defects in each category have relatively smooth background as we have found that it converges much faster and does not pose a limitation since every defect (resized to 64X64 pixels) can be overlaid on any texture using seamless cloning.



Figure 5 – Dents DB for training the VAE



Figure 6 – Dripping DB for training the VAE

500K training iterations were used with initial learning rate of 0.001 and Adam optimizer to train separately two VAE networks- one for generating dents and the other for generating drips. Fig. 7 demonstrates 100 VAE generated dents, while Fig. 8 is the same but for drips. It can be noted that the generated defects look realistic.

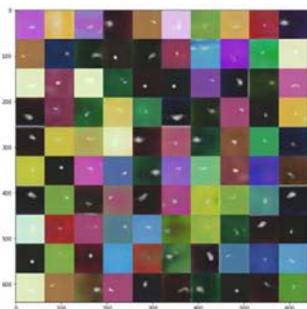


Figure 7 – 100 VAE generated dents

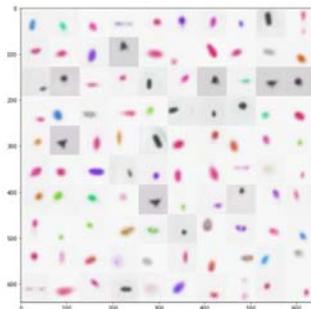


Figure 8 – 100 VAE generated drips

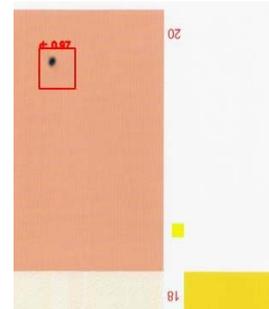


Figure 9 – Example of overlaying drip (in red rectangle) on defect free image using seamless-cloning

Fig. 9 is an example of overlaying one of the generated drips on top of one image from the defect free DB. The composition looks seamless and natural and therefore can be used for training Siamese based SSD (detection example is the red bounding box around the defect in Fig. 9). For better visibility on textured areas contrast stretching can be applied on the luminance channel of the generated defects before seamless cloning.

Competitive approaches

GANs are an alternative to VAEs for generating data, by forcing the generated images to be statistically almost indistinguishable from real ones. However, it is much harder to train compared to VAE as it contains two networks that try to beat each other. Although VAEs may create blurrier samples it is not limitation in our case as we simulate quite simple defects. To the best of our knowledge this work is the first to generate print defects using VAE and in particular drips and dents. Also, we have designed novel compact network that requires minimum (few dozen) samples per defect type and thus when combined with seamless cloning result in defects that look realistic. The generated defects can be used to train deep learning-based defect detection and classification system.

Current status

We have successfully trained VAEs that generate drips and dents. It will be integrated into our inspection system by using it to train the Siamese based SSD detection and classification module.

Next steps

We plan to train one VAE for each high-priority defect type class, thus creating the ability to pinpoint the defect's root cause.

Disclosed by Oren Haik, Oded Perry and Eli Chen, HP Inc.