March 2020

LEARNING AN IMAGE ENHANCEMENT PIPLINE

HP INC

Learning an Image Enhancement Pipeline

Abstract: An agent infers an optimal set of parameters for a set of selected enhancement transformations for each specific image and allows an image pipeline to be learns without having any prior knowledge of the target or ground truth images.

This disclosure relates to the field of image processing.

A technique is disclosed that trains an agent to infer an optimal set of parameters for a set of selected enhancement transformations for each specific image.

Lot of efforts has been expended to find ways to improve documents quality after being scanned or shot by a camera, and to provide the best printed quality through human grading of print quality, such as by computer vision engineers. They typically work on a relatively small set of images and the pipeline ultimately depends on their level of expertise to come up with a hand-coded solution.

Some applications, such as some scanner features, try to infer the best pipeline enhancement for a few different documents types or based on some properties of documents. It is a challenging task, as these usually focus on one single set of parameters for all documents, resulting in suboptimal enhancement pipelines that tends to optimally learn only a single task well (e.g. deblurring, denoising, JPEG artifacts reduction, etc.). It's very hard to come up with a generic enhancement pipeline that will work for most of the images one could encounter. Most of the literature improves only one aspect of the image with a recognition that improving for many factors would lead to too complex NN architecture and would require an excessive amount of resources.

Today, some mobile applications embed photo editor into their applications to let user performing fine-grained transformations on their documents/photos.

According to the present disclosure, a hybrid model based on Reinforcement Learning learns the optimal set of parameters of a few selected enhancement transformations for each specific image, and learns the pipeline without having any prior knowledge of the target or ground truth images.

A few functions are pre-selected before learning the image enhancement pipeline. A selection of N transformation function that will be used to enhance the images are pre-selected. Each of these is usually a well-known enhancement operations used in image enhancement pipelines such as lightning enhancement, color enhancement, de-blur, de-noise, contrast, etc. Each of these may take one or more parameters. These parameters are the elements that the agent will learn to makes its optimal inferences.

A selection of Image Quality Assessment (IQA) scores may be relevant to the specific image enhancement that one would look for. For each IQA function, one can associate a weight wi.

From the Reinforcement Learning framework perspective, there is an Environment and an Agent. The Environment has available a set of images that will be used for the many epochs executed by the learning agent. For each epoch, a subset of these images (say X) is selected, and the goal is to optimize the overall Reward signal during that epoch.

Environment:

The Environment will do the following:

1) For all images, always maintain a matrix of all images latent vector thanks to a Convolutional Network (such as for example VGG, ResNet, MobileNet, or another) and parameters computed by the agent (0 if none). That constitutes the Overall State (a map of images in some sort) of the Environment.

2) For all images, compute the scores for all selected IQA functions.

3) Compute the Reward as represented by weighted average of all IQA scores for the last enhanced image. If that reward signal is greater than the last one for that image, replace it in the Overall State by the new latent vector and set of parameters.

Agent:

The Agent infers (and learns) the optimal parameters for each of the selected transformation functions that are part of the enhancement pipeline, based on the latent vector of a specific image. A parameter is defined that may be used to enable/disable the related transformation function. These parameters define the actions that the agent will take, by applying the transformation functions with the inferred parameters to the image, resulting in an enhanced image. In addition, a memory of the agent parameters is built, and latent vectors are inferred for a specific image (matrix).

With a frequency F (i.e. a few rounds), the Agent initiates a learning phase, performing backpropagation over all its local NN used to infer the function parameters. The cost function may be based on several different strategies, such as MSE. While cost functions usually rely on a set of known target / ground truth images, that can't be done in this case since these are unknown.

As a result, the Overall State of the environment will always represent the best set of images known at the current stage of learning. In that way, the target / ground truth set of images is built while the Reinforcement Learning loop is executed. Thus the Agent will compute the MSE or other cost function over that set of best (to date) enhanced images.

The disclosed technique advantageously learns likely optimal parameters for each specific image to be enhanced, in contrast to the prior technique of a computer vision engineer configuring a fixed set of enhancement transformation with some hand-coded heuristic to decide what needs to be done to enhance each photo. The technique enables the creation of a learned enhancement pipeline specific to a printer or a press.

Disclosed by Sebastien Tandel, HP Inc.