



Enhancement of Retinal Fundus Images via Pixel Color Amplification

Alex Gaudio¹✉, Asim Smailagic¹, and Aurélio Campilho^{2,3}

¹ Carnegie Mellon University, Pittsburgh, PA 15213, USA
agaudio@andrew.cmu.edu, asim@cs.cmu.edu

² INESC TEC, Porto, Portugal

³ Faculty of Engineering, University of Porto, Porto, Portugal
campilho@fe.up.pt

Abstract. We propose a pixel color amplification theory and family of enhancement methods to facilitate segmentation tasks on retinal images. Our novel re-interpretation of the image distortion model underlying dehazing theory shows how three existing priors commonly used by the dehazing community and a novel fourth prior are related. We utilize the theory to develop a family of enhancement methods for retinal images, including novel methods for whole image brightening and darkening. We show a novel derivation of the Unsharp Masking algorithm. We evaluate the enhancement methods as a pre-processing step to a challenging multi-task segmentation problem and show large increases in performance on all tasks, with Dice score increases over a no-enhancement baseline by as much as 0.491. We provide evidence that our enhancement preprocessing is useful for unbalanced and difficult data. We show that the enhancements can perform class balancing by composing them together.

Keywords: Image enhancement · Medical image analysis · Dehazing · Segmentation · Multi-task learning

1 Introduction

Image enhancement is a process of removing noise from images in order to improve performance on a future image processing task. We consider image-to-image pre-processing methods intended to facilitate a downstream image processing task such as Diabetic Retinopathy lesion segmentation, where the goal is to identify which pixels in an image of a human retina are pathological. In this setting, image enhancement does not in itself perform segmentation, but rather it elucidates relevant features. Figure 1 shows an example enhancement with our method, which transforms the color of individual pixels and enhances fine detail.

Our main contributions are to re-interpret the distortion model underlying dehazing theory as a theory of pixel color amplification. Building on the widely known Dark Channel Prior method [5], we show a novel relationship between three previously known priors and a fourth novel prior. We then use these four




Fig. 1. Comparing unmodified image (left) to our enhancement of it (right).

priors to develop a family of brightening and darkening methods. Next, we show how the theory can derive the Unsharp Masking method for image sharpening. Finally, show that the pre-processing enhancement methods improve performance of a deep network on five retinal fundus segmentation tasks. We also open source our code for complete reproducibility [4].

2 Related Work

Natural images are distorted by refraction of light as it travels through the transmission medium (such as air), causing modified pixel intensities in the color channels of the image. A widely used physical theory for this distortion has traditionally been used for single image dehazing [2, 5, 8, 14]:

$$\mathbf{I}(\mathbf{x}) = \mathbf{J}(\mathbf{x})t(\mathbf{x}) + \mathbf{A}(1 - t(\mathbf{x})), \quad (1)$$

where each pixel location, \mathbf{x} , in the distorted RGB image, \mathbf{I} , can be constructed as a function of the distortion-free radiance image \mathbf{J} , a grayscale transmission map image \mathbf{t} quantifying the relative portion of the light ray coming from the observed surface in $\mathbf{I}(\mathbf{x})$ that was not scattered (and where $t(\mathbf{x}) \in [0, 1] \forall \mathbf{x}$), and an atmosphere term, \mathbf{A} , which is typically a RGB vector that approximates the color of the uniform scattering of light. Distortion is simply a non-negative airlight term $\mathbf{A}(1 - t(\mathbf{x}))$. We refer to [2] for a deeper treatment of the physics behind the theory in a dehazing context. Obtaining a distortion free image \mathbf{J} via this theory is typically a three step process. Given \mathbf{I} , define an atmosphere term \mathbf{A} , solve for the transmission map \mathbf{t} , and then solve for \mathbf{J} . We develop new insights into this theory by demonstrating ways in which it can behave as a pixel amplifier when \mathbf{t} and \mathbf{A} are allowed to be three channel images.

The well known Dark Channel Prior (DCP) method [5, 7] addresses the dehazing task for Eq. (1) by imposing a prior assumption on RGB images. The assumption differentiates the noisy (hazy) image, \mathbf{I} , from its noise free (dehazed) image, \mathbf{J} . That is, in any haze-free multi-channel region of a RGB image, at least one pixel has zero intensity in at least one channel ($\{(0, g, b), (r, 0, b), (r, g, 0)\}$), while

	Inverted DCP (Illumination Correction)
Dark Channel Prior (Dehazing)	
$\mathbf{A} = (r, g, b).$ (2)	$\mathbf{A} = (1, 1, 1).$ (7)
$\tilde{t}(\mathbf{x}) = 1 - \min_c \min_{\mathbf{y} \in \Omega_I(\mathbf{x})} \frac{I^c(\mathbf{y})}{A^c}$ (3)	$\mathbf{J} = 1 - f_{\text{DCP}}(1 - \mathbf{I}, \mathbf{A})$ (8)
$t(\mathbf{x}) = \text{guidedFilter}(\mathbf{I}, \tilde{t}(\mathbf{x})).$ (4)	Bright Channel Prior (Exposure Correction)
$\mathbf{J}(\mathbf{x}) = \frac{\mathbf{I}(\mathbf{x}) - \mathbf{A}}{\max(t(\mathbf{x}), \epsilon)} + \mathbf{A}.$ (5)	$\tilde{t}(x) = 1 - \max_c \max_{\mathbf{y} \in \Omega_I(\mathbf{x})} \frac{I^c(\mathbf{y})}{A^c}$ (9)
$\implies \mathbf{J} = f_{\text{DCP}}(\mathbf{I}, \mathbf{A})$ (6)	$t(\mathbf{x}) = \text{guidedFilter}(\mathbf{I}, \tilde{t}(\mathbf{x})).$ (10)
	$\implies \mathbf{J} = f_{\text{BCP}}(\mathbf{I}, \mathbf{A})$ (11)

Fig. 2. Left: Dark Channel Prior (DCP) method for dehazing. Given an (inverted) image \mathbf{I} and atmosphere \mathbf{A} , obtain transmission map \mathbf{t} and then recover \mathbf{J} , the undistorted image. **Top and Bottom Right:** Two priors based on inversion of the Dark Channel Prior.

a hazy region will have no pixels with zero intensity ($r > 0, g > 0, b > 0$). The assumption is invalid if any channel of a distorted image is sparse or if all channels of the undistorted image are not sparse. To quantify distortion in an image, the assumption justifies creating a fourth channel, known as the dark channel, by applying a min operator convolutionally to each region of the images \mathbf{I} and \mathbf{J} . Specifically, $\tilde{I}^{\text{dark}}(\mathbf{x}) = \min_c \min_{\mathbf{y} \in \Omega_I(\mathbf{x})} \frac{I^c(\mathbf{y})}{A^c}$, where c denotes the color channel (red, green or blue) and $\Omega_I(\mathbf{x})$ is a set of pixels in \mathbf{I} neighboring pixel \mathbf{x} . The min operator causes \tilde{I}^{dark} to lose fine detail, but an edge-preserving filter known as the guided filter [6] restores detail $\mathbf{I}^{\text{dark}} = g(\tilde{\mathbf{I}}^{\text{dark}}, \mathbf{I})$. While $\mathbf{J}^{\text{dark}}(\mathbf{x})$ always equals zero and therefore cancels out of the equations, $\mathbf{I}^{\text{dark}}(\mathbf{x})$ is non-zero in hazy regions. By observing that the distortion free image \mathbf{J}^{dark} is entirely zero while \mathbf{I}^{dark} is not entirely zero, solving Eq. (1) for \mathbf{t} leads to Eq. (4) and then Eq. (5) in Fig. 2. In practice, the denominator of (5) is $\max(t(\mathbf{x}), \epsilon)$ to avoid numerical instability or division by zero; this amounts to preserving a small amount of distortion in heavily distorted pixels. Figure 2 summarizes the mathematics.

The DCP method permits various kinds of inversions. The bright channel prior [15] solves for a transmission map by swapping the min operator for a max operator in Eq. (4). This prior was shown useful for exposure correction. Figure 2 shows our variation of the bright channel prior based more directly on DCP mathematics and with an incorporated guided filter. Another simple modification of the DCP method is to invert the input image \mathbf{I} to perform illumination correction [3, 12, 13]. The central idea is to invert the image, apply the dehazing equations, and then invert the dehazed result. We demonstrate the mathematics of this inverted DCP method in Fig. 2. Color illumination literature requires the assumption that $\mathbf{A} = (1, 1, 1)$, meaning the image is white-balanced.

In the dehazing context, this assumption would mean the distorted pixels are too bright, but in the color illumination context, distorted pixels are too dark. In the Methods section, we expand on the concept of brightness and darkness as pixel color amplification, show the theory supports other values of \mathbf{A} , and we also expand on the concept of inversion of Eqs. (4) and (5) for a wider variety of image enhancements.

3 Methods

The distortion theory Eq. (1) is useful for image enhancement. In Sect. 3.1, we show how the theory is a pixel color amplifier. In Sect. 3.2, we show ways in which the theory is invertible. We apply these properties to derive a novel prior and present a unified view of amplification under four distinct priors. Sections 3.2 and 3.2 apply the amplification theory to three specific enhancement methods: whole image brightening, whole image darkening and sharpening.

3.1 The Distortion Theory Amplifies Pixel Intensities

We assume that \mathbf{A} , \mathbf{I} and \mathbf{J} share the same space of pixel intensities, so that in any given channel c and pixel location \mathbf{x} , the intensities A^c , $I^c(\mathbf{x})$ and $J^c(\mathbf{x})$ can all have the same maximum or minimum value. We can derive the simple equation $t(\mathbf{x}) = \frac{I^{(c)}(\mathbf{x}) - A^{(c)}}{J^{(c)}(\mathbf{x}) - A^{(c)}} \in [0, 1]$ from Eq. (1) by noting that the distortion theory presents a linear system containing three channels. The range of \mathbf{t} implies the numerator and denominator must have the same sign. For example, if $A^{(c)} \geq I^{(c)}(\mathbf{x})$, then the numerator and denominator are non-positive and $J^{(c)}(\mathbf{x}) \leq I^{(c)}(\mathbf{x}) \leq A^{(c)}$. Likewise, when $A^{(c)} \leq I^{(c)}(\mathbf{x})$, the order is reversed $J^{(c)}(\mathbf{x}) \geq I^{(c)}(\mathbf{x}) \geq A^{(c)}$. These two ordering properties show the distortion theory amplifies pixel intensities. The key insight is that the choice of \mathbf{A} determines how the color of each pixel in the recovered image \mathbf{J} changes. Models that recover \mathbf{J} using Eq. (1) will simply amplify color values for each pixel \mathbf{x} in the direction $\mathbf{I}(\mathbf{x}) - \mathbf{A}$.

Atmosphere Controls the Direction of Amplification in Color Space.

The atmosphere term \mathbf{A} is traditionally a single RGB color vector with three scalar values, $A = (r, g, b)$, but it can also be an RGB image matrix. As a RGB vector, \mathbf{A} does not provide precise pixel level control over the amplification direction. For instance, two pixels with the same intensity are guaranteed to change color in the same direction, even though it may be desirable for these pixels to change color in opposite directions. Fortunately, considering \mathbf{A} as a three channel RGB image enables precise pixel level control of the amplification direction. It is physically valid to consider \mathbf{A} as an image since the atmospheric light may shift color across the image, for instance due to a change in light source. As an image, \mathbf{A} can be chosen to define the direction of color amplification $I^c(\mathbf{x}) - A^c(\mathbf{x})$ for each pixel and each color channel independently.

Transmission Map and Atmosphere Both Control the Rate of Amplification. Both the transmission map \mathbf{t} and the magnitude of the atmosphere term \mathbf{A} determine the amount or rate of pixel color amplification. The effect on amplification is shown in the equation $\mathbf{J} = \frac{\mathbf{I}-\mathbf{A}}{\mathbf{t}} + \mathbf{A}$, where the difference $\mathbf{I} - \mathbf{A}$ controls the direction and magnitude of amplification and \mathbf{t} affects the amount of difference to amplify. The transmission map itself is typically a grayscale image matrix, but it can also be a scalar constant or a three channel color image. Each value $t(\mathbf{x}) \in [0, 1]$ is a mixing coefficient specifying what proportion of the signal is not distorted. When $t(\mathbf{x}) = 1$, there is no distortion; the distorted pixel $\mathbf{I}(\mathbf{x})$ and corresponding undistorted pixel $\mathbf{J}(\mathbf{x})$ are the same since $\mathbf{I}(\mathbf{x}) = \mathbf{J}(\mathbf{x}) + 0$. As $t(\mathbf{x})$ approaches zero, the distortion caused by the difference between the distorted image \mathbf{I} and the atmosphere increases.

3.2 Amplification Under Inversion

The distortion theory supports several kinds of inversion. The Eqs. (4) and (5) are invertible. The input image \mathbf{I} can also undergo invertible transformations. We prove these inversion properties and show why they are useful.

Inverting Eq. (4) Results in a Novel DCP-Based Prior. We discussed in Related Work three distinct priors that provide a transmission map: the traditional DCP approach with Eq. (4); the bright channel prior in Eq. (10); and color illumination via Eq. (8). Bright channel prior and color illumination respectively perform two types of inversion; the former changes the min operator to a max operator while the latter inverts the image $1 - \mathbf{I}$. Combining these two inversion techniques results in a novel fourth prior. In Table 1, we show the four transmission maps. We show that each prior has a solution using either the min or max operator, which is apparent by the following two identities:

$$\text{solve_t}(\mathbf{I}, \mathbf{A}) = 1 - \min_c \min_{y \in \Omega_I(\mathbf{x})} \frac{I^c(\mathbf{y})}{A^c} \equiv \max_c \max_{y \in \Omega_I(\mathbf{x})} \frac{1 - I^c(\mathbf{y})}{A^c} \quad (12)$$

$$\text{solve_t}(\mathbf{I}, \mathbf{A}) = 1 - \max_c \max_{y \in \Omega_I(\mathbf{x})} \frac{I^c(\mathbf{y})}{A^c} \equiv \min_c \min_{y \in \Omega_I(\mathbf{x})} \frac{1 - I^c(\mathbf{y})}{A^c} \quad (13)$$

The unified view of these four priors in Table 1 provides a novel insight into how they are related. In particular, the table provides proof that the Color Illumination Prior and Bright Channel Prior are inversely equivalent and utilize statistics of the maximum pixel values across channels. Similarly, DCP and our prior are also inversely equivalent and utilize statistics of the minimum pixel values across channels. This unified view distinguishes between weak and strong amplification, and amplification of bright and dark pixel neighborhoods.

In Fig. 3, we visualize these four transmission maps to demonstrate how they collectively perform strong or weak amplification of bright or dark regions of the input image. In this paper, we set $\mathbf{A} = \mathbf{1}$ when solving for \mathbf{t} . Any choice of $A^c \in (0, 1]$ is valid, and when all A^c are equal, smaller values of \mathbf{A} are guaranteed to amplify the differences between these properties further.

Table 1. Four transmission maps derived from variations of Eq. (4). For clear notation, we used the vectorized functions $\mathbf{t} = \text{solveMin.t}(\mathbf{I}, \mathbf{A}) = 1 - \min_c \min_{\mathbf{y} \in \Omega_{I(\mathbf{x})}} \frac{I^c(\mathbf{y})}{A^c}$ and $\mathbf{t} = \text{solveMax.t}(\mathbf{I}, \mathbf{A}) = 1 - \max_c \max_{\mathbf{y} \in \Omega_{I(\mathbf{x})}} \frac{I^c(\mathbf{y})}{A^c}$.

	Amplify dark areas	Amplify bright areas
Weak amplification	$\text{solveMin.t}(1 - \mathbf{I}, \mathbf{A} = 1)$ $1 - \text{solveMax.t}(\mathbf{I}, \mathbf{A} = 1)$ Color Illumination Prior	$\text{solveMin.t}(\mathbf{I}, \mathbf{A} = 1)$ $1 - \text{solveMax.t}(1 - \mathbf{I}, \mathbf{A} = 1)$ Standard Dark Channel Prior
Strong amplification	$1 - \text{solveMin.t}(\mathbf{I}, \mathbf{A} = 1)$ $\text{solveMax.t}(1 - \mathbf{I}, \mathbf{A} = 1)$ Our novel prior	$1 - \text{solveMin.t}(1 - \mathbf{I}, \mathbf{A} = 1)$ $\text{solveMax.t}(\mathbf{I}, \mathbf{A} = 1)$ Bright Channel Prior




Fig. 3. The transmission maps (right) obtained from source image (left) selectively amplify bright or dark regions. Dark pixels correspond to a larger amount of amplification. We randomly sample a retinal fundus image from the IDRiD dataset (see Sect. 4.1). We set the blue channel to all ones when computing the transmission map for the top right and bottom left maps because the min values of the blue channel in retinal fundus images are noisy. (Color figure online)

Inverting Eq. (5) Motivates Brightening and Darkening. Given an image \mathbf{I} , transmission map \mathbf{t} and an atmosphere \mathbf{A} , solving for the recovered image \mathbf{J} with Eq. (5) can be computed two equivalent ways, as we demonstrate by the following identity:

$$\mathbf{J} = \text{solve}_J(\mathbf{I}, \mathbf{t}, \mathbf{A}) \equiv 1 - \text{solve}_J(1 - \mathbf{I}, \mathbf{t}, 1 - \mathbf{A}) \tag{14}$$

The proof is by simplification of $\frac{\mathbf{I}-\mathbf{A}}{\mathbf{t}} + \mathbf{A} = 1 - \left(\frac{(1-\mathbf{I})-(1-\mathbf{A})}{\mathbf{t}} + (1 - \mathbf{A}) \right)$. It implies the space of possible atmospheric light values, which is bounded in $[0, 1]$, is symmetric under inversion.

We next prove that solving for \mathbf{J} via the color illumination method [3, 12, 13] is equivalent to direct attenuation $\mathbf{J} = \frac{\mathbf{I}}{\mathbf{t}}$, a fact that was not clear in prior work. As we noted in Eq. (8), color illumination solves $\mathbf{J} = 1 - \frac{(1-\mathbf{I})-\mathbf{A}}{\mathbf{t}} + \mathbf{A}$ under the required assumption that $\mathbf{A} = \mathbf{1}$. We can also write the atmosphere as $\mathbf{A} = 1 - \mathbf{0}$.

Then, the right hand side of (14) leads to $\mathbf{J} = 1 - \text{solve_J}(1 - \mathbf{I}, \mathbf{t}, \mathbf{A} = 1 - \mathbf{0}) = \frac{\mathbf{I} - \mathbf{0}}{\mathbf{t}} + \mathbf{0}$. Therefore, color illumination actually performs whole image brightening with the atmosphere $\mathbf{A} = (0, 0, 0)$ even though the transmission map uses a white-balanced image assumption that $\mathbf{A} = (1, 1, 1)$. Both this proof and the invertibility property Eq. (14) motivate Sect. 3.2 where we perform brightening and darkening with all priors in Table 1.




Fig. 4. Whole image brightening (left) and darkening (right) using the corresponding four transmission maps in Fig. 3. Note that $\mathbf{A} = \mathbf{1}$ when solving for \mathbf{t} , but $\mathbf{A} = \mathbf{0}$ or $\mathbf{A} = \mathbf{1}$, respectively, for brightening or darkening. (Color figure online)

Application to Whole Image Brightening and Darkening. Brightening versus darkening of colors is a matter of choosing an amplification direction. Extremal choices of the atmosphere term \mathbf{A} result in brightening or darkening of all pixels in the image. For instance, $\mathbf{A} = (1, 1, 1)$ guarantees for each pixel \mathbf{x} that the recovered color $\mathbf{J}(\mathbf{x})$ is darker than the distorted color $\mathbf{I}(\mathbf{x})$ since $\mathbf{J} \leq \mathbf{I} \leq \mathbf{A}$, while $\mathbf{A} = (0, 0, 0)$ guarantees image brightening $\mathbf{J} \geq \mathbf{I} \geq \mathbf{A}$. More generally, any \mathbf{A} satisfying $1 \geq A^c \geq \max_{\mathbf{x}} I^c(\mathbf{x})$ performs whole image brightening and any \mathbf{A} satisfying $0 \leq A^c \leq \min_{\mathbf{x}} I^c(\mathbf{x})$ performs whole image darkening. We utilize the four distinct transmission maps from Table 1 to perform brightening $\mathbf{A} = \mathbf{0}$ or darkening $\mathbf{A} = \mathbf{1}$, resulting in eight kinds of amplification. We visualize these maps and corresponding brightening and darkening techniques applied to retinal fundus images in Fig. 4. Our application of the Bright Channel Prior and Color Illumination Prior for whole image darkening is novel. Utilizing our prior for brightening and darkening is also novel.

Application to Image Sharpening. We show a novel connection between dehazing theory and *unsharp masking*, a deblurring method and standard image

sharpening technique that amplifies fine detail [9]. Consider \mathbf{A} as a three channel image obtained by applying a non-linear blur operator to \mathbf{I} , $\mathbf{A} = \text{blurry}(\mathbf{I})$. Solving Eq. (1) for \mathbf{J} gives $\mathbf{J} = \frac{1}{t}\mathbf{I} - \frac{(1-t)}{t}\mathbf{A}$. Since each scalar value $t(\mathbf{x})$ is in $[0, 1]$, we can represent the fraction $t(\mathbf{x}) = \frac{1}{u(x)}$. Substituting, we have the simplified matrix form $\mathbf{J} = \mathbf{u} \circ \mathbf{I} - (\mathbf{u} - 1) \circ \text{blurry}(\mathbf{I})$ where the \circ operator denotes element-wise multiplication with broadcasting across channels. This form is precisely *unsharp masking*, where \mathbf{u} is either a constant, or \mathbf{u} is a 1-channel image matrix determining how much to sharpen each pixel. The matrix form of \mathbf{u} is known as locally adaptive unsharp masking. Thus, we show the distortion theory in Eq. (1) is equivalent to image sharpening by choosing \mathbf{A} to be a blurred version of the original input image.

We present two sharpening algorithms, Algorithm 1 and 2, and show their respective outputs in Fig. 5. Sharpening amplifies differences between an image and a blurry version of itself. In unevenly illuminated images, the dark or bright regions may saturate to zero or one respectively. Therefore, the use of a scalar transmission map (Algorithm 1), where all pixels are amplified, implies that the input image should ideally have even illumination. The optional guided filter in the last step provides edge preserving smoothing and helps to minimize speckle noise, but can cause too much blurring on small images, hence the if condition.

Algorithm 2 selectively amplifies only the regions that have an edge. Edges are found by deriving a three channel transmission map from a Laplacian filter applied to a morphologically smoothed fundus image. We enhance edges by recursively sharpening the Laplace transmission map under the theory. Figure 6 shows the results of sharpening each image in Fig. 4 with Algorithm 2.

4 Experiments

Our primary hypothesis is that enhancement facilitates a model’s ability to learn retinal image segmentation tasks. We introduce a multi-task dataset and describe our deep network implementation.

4.1 Datasets

The **Indian Diabetic Retinopathy Dataset (IDRiD)** [11] contains 81 retinal fundus images for segmentation, with a train-test split of 54:27 images. Each image is 4288×2848 pixels. Each pixel has five binary labels for presence of: Microaneurysms (MA), Hemorrhages (HE), Hard Exudates (EX), Soft Exudates (SE) and Optic Disc (OD). Only 53:27 and 26:14 images present HE and SE, respectively. Table 2 shows the fraction of positive pixels per category is unbalanced both across categories (left columns) and within categories (right columns).




Fig. 5. Sharpening a retinal fundus image with Algorithm 1 (middle) and 2 (right). Image randomly sampled from IDRiD training dataset (described in Sec. 4.1). (Color figure online)

<p>Algorithm 1: Image Sharpening, simple</p> <hr/> <p>Input: I (input fundus image) Result: J (sharpened image) $A = \text{blur}(I, \text{blur_radius});$ $t = 0.15;$ if $\min(\text{img_width}, \text{img_height}) > 1500$ then $J = \text{guidedFilter}(\text{guide} = I, \text{src} = \frac{I-A}{t} + A);$ else $J = \frac{I-A}{t} + A;$ end</p>	<p>Algorithm 2: Image Sharpening, complex</p> <hr/> <p>Input: I (input fundus image) Result: J (sharpened image) $\tilde{t} = \text{Algo_1}(\text{morphological_laplace}(I, (2, 2, 1)));$ $\tilde{t} = 1 - \frac{\tilde{t} - \min(\tilde{t})}{\max(\tilde{t}) - \min(\tilde{t})};$ $\epsilon = \max(10^{-8}, \frac{\min(\tilde{t})}{2});$ $t = \text{elementwise_max}(\tilde{t}, \epsilon);$ $J = \text{Algo_1}(I, t = t);$</p>
---	---




Fig. 6. The result of sharpening each image in Fig. 4 using Algorithm 2. (Color figure online)

Table 2. IDRiD Dataset, an unbalanced class distribution.

Category	Pos/ \sum Pos		Pos/(Pos + Neg)	
	Train	Test	Train	Test
MA	0.027	0.024	0.0007	0.0003
HE	0.253	0.256	0.0066	0.0036
EX	0.207	0.261	0.0054	0.0036
SE	0.049	0.043	0.0013	0.0006
OD	0.464	0.416	0.0120	0.0058

Table 3. Competing method results, best per category of A, B, D, or X.

Task	Method	Dice (delta)
EX	A	0.496 (0.175)
HE	A	0.102 (0.102)
MA	A	0.122 (0.122)
OD	A	0.849 (0.332)
SE	A	0.423 (0.264)

Table 4. Main results, pre-processing yields large improvements.

Task	Method	Dice (delta)
EX	avg4:sA + sC + sX + sZ	0.728 (0.407)
	avg2:sA + sZ	0.615 (0.295)
HE	avg3:sA + sC + sX	0.491 (0.491)
	avg3:sB + sC + sX	0.368 (0.368)
MA	avg4:A + B + C + X	0.251 (0.251)
	avg2:sB + sX	0.219 (0.219)
OD	avg4:sA + sC + sX + sZ	0.876 (0.359)
	avg2:sA + sZ	0.860 (0.343)
SE	avg4:sA + sC + sX + sZ	0.491 (0.332)
	avg3:B + C + X	0.481 (0.322)

Blackbox Evaluation: Does an Enhancement Method Improve Performance? We implement and train a standard U-Net model [10] and evaluate change in performance via the Dice coefficient. We apply this model simultaneously to five segmentation tasks (MA, HE, SE, EX, OD) on the IDRiD dataset; the model has five corresponding output channels. We use a binary cross entropy loss summed over all pixels and output channels. We apply task balancing weights to ensure equal contribution of positive pixels to the loss. The weights are computed via $\frac{\max_i w_i}{\mathbf{w}}$, where the vector \mathbf{w} contains counts of positive pixels across all training images for each of the 5 task categories (see left column of Table 2). Without the weighting, the model did not learn to segment MA, HE, and EX even with our enhancements. For the purpose of our experiment, we show in the results that this weighting is suboptimal as it does not balance bright and dark categories. The Adam Optimizer has a learning rate 0.001 and weight decay 0.0001. We also applied the following pre-processing: center crop the fundus to minimize background, resize to (512×512) pixels, apply the pre-processing enhancement method (independent variable), clip pixel values into $[0, 1]$, randomly rotate and flip. Rotations and flipping were only applied on the training set; we excluded them from validation and test sets. We randomly hold out two training images as the validation set in order to apply early stopping with a patience of 30 epochs. We evaluate test set segmentation performance with the Sørensen-Dice coefficient, which is commonly used for medical image segmentation.

4.2 Pre-processing Enhancement Methods for Retinal Fundus Images

We combine the brightening, darkening and sharpening methods together and perform an ablation study. We assign the eight methods in Fig. 4 a letter. The brightening methods, from top left to bottom right are A, B, C, D. The corresponding darkening methods are W, X, Y, Z. We also apply sharpening via Algorithm 2. Combined methods assume the following notation: $A + X$ is the average of A and X , which is then sharpened; $sA + sX$ is the average of sharpened A with sharpened X . A standalone letter X is a sharpened X . All methods have the same hyperparameters, which we chose using a subset of IDRiD training images. When solving for t , the size of the neighborhood Ω is (5×5) ; the guided filter for t has radius = 100, $\epsilon = 1e^{-8}$. When solving for J , the max operator in the denominator is $\max(\min(t)/2, 1e-8)$. For sharpening, we blur using a guided filter (radius = 30, $\epsilon = 1e^{-8}$), and we do not use a guided filter to denoise as the images are previously resized to (512×512) .

5 Results

5.1 Our Pre-processing Enhancement Methods Significantly Improve Performance on All Tasks

We show the top two models with highest test performance in each category in Table 4. The delta values show the pre-processing enhancement methods significantly improve performance over the no-enhancement (identity function) baseline for all tasks, underscoring the value of our theory and methods.

Enhancement Improves Detection of Rare Classes. The smallest delta improvement in Table 4 is 0.219 for MA, the rarest category across and within categories (as shown in Table 2). Our smallest improvement is a large increase considering the largest possible Dice score is one.

Enhancement can be Class Balancing. The IDRiD results support the primary hypothesis that enhancement makes the segmentation task easier. The delta values show the baseline identity model did not learn to segment MA or HE. Indeed, during the implementation, we initially found that the model learned to segment only the optic disc (OD). Of the categories, OD has the most extremal intensities (brightest) and is typically the largest feature by pixel count in a fundus image. In our multi-task setting, the gradients from other tasks were therefore overshadowed by OD gradients. After we implemented a category balancing weight, the no-enhancement baseline model was still unable to learn MA and HE. As an explanation for this phenomenon, we observe that EX, SE and OD are bright features while MA and HE are dark features. Considering the class balancing weights, the bright features outnumber the dark features three to two. This need to carefully weigh the loss function suggests differences in color intensity values *cause* differences in performance. It is therefore particularly interesting that the enhancement methods were able to learn




Fig. 7. Visualization of our enhancement methods. Each row is an image. Each column is an enhancement method. Last two rows compare our Algorithm 1 with CLAHE. (Color figure online)

despite also being subject to these issues. In fact, we can observe that the best enhancements in the table incorporate the Z method, which performs a strong darkening of bright regions. We interpret this result as strong evidence that our enhancement methods make the segmentation task easier, and in fact, that they can be used as a form of class balancing by image color augmentation.

5.2 Comparison to Existing Work

The methods A , D , X and arguably B correspond to existing work and were visualized (with sharpening) in Fig. 6. The A method outperforms B , D and X on all tasks. Its values, reported in Table 3, are substantially lower than the values in Table 4. We attribute the low scores to our intentional category imbalance.

Contrast Limited Adaptive Histogram Equalization (CLAHE) applied to the luminance channel of LAB colorspace is useful for retinal fundus enhancement [1]. We compare it to Algorithm 1 in bottom rows of Fig. 7, using the LAB conversion for both methods. We observe that CLAHE preserves less detail, and both methods overemphasize uneven illumination. CLAHE is faster to compute and could serve as a simple drop-in replacement, with clip limit as a proxy for the scalar t .

5.3 Qualitative Analysis

We visualize a subset of our image enhancement methods in the top three rows of Fig. 7. Each row presents a different fundus image from a private collection. We observe that the input images are difficult to see and appear to have little detail, while the enhanced images are quite colorful and very detailed. The halo effect around the fundus is caused by the guided filter (with $\epsilon = 1e^{-8}$) rather than the theory. The differences in bright and dark regions across each row provide intuitive sense of how averaging the models (Fig. 4 and 6) can yield a variety of different colorings.

6 Conclusion

In this paper, we re-interpret a theory of image distortion as pixel color amplification and utilize the theory to develop a family of enhancement methods for retinal fundus images. We expose a relationship between three existing priors commonly used for image dehazing with a fourth novel prior. We apply our theory to whole image brightening and darkening, resulting in eight enhancement methods, five of which are also novel (methods B, C, W, Y, and Z). We also show a derivation of the Unsharp Masking algorithm for image sharpening and develop a sharpening algorithm for retinal fundus images. Finally, we evaluate our enhancement methods as pre-processing steps for multi-task deep network retinal fundus image segmentation. We show the enhancement methods give strong improvements and can perform class balancing. Our pixel color amplification theory applied to retinal fundus images yields a variety of rich and colorful enhancements, as shown by our compositions of methods A-D and W-Z, and the theory shows great promise for wider adoption by the community.

Acknowledgements. We thank Dr. Alexander R. Gaudio, a retinal specialist and expert in degenerative retinal diseases, for his positive feedback and education of fundus images.

Supported in part by the National Funds through the Fundação para a Ciência e a Tecnologia within under Project CMUPERI/TIC/0028/2014.

References

1. Cao, L., Li, H., Zhang, Y.: Retinal image enhancement using low-pass filtering and alpha-rooting. *Signal Process.* **170**, 107445 (2020). <https://doi.org/10.1016/j.sigpro.2019.107445>. <http://www.sciencedirect.com/science/article/pii/S0165168419304967>
2. Fattal, R.: Single image dehazing. *ACM Trans. Graph.* **27**(3), 72:1–72:9 (2008). <https://doi.org/10.1145/1360612.1360671>
3. Galdran, A., Bria, A., Alvarez-Gila, A., Vazquez-Corral, J., Bertalmio, M.: On the duality between retinex and image dehazing. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (June 2018). <https://doi.org/10.1109/cvpr.2018.00857>
4. Gaudio, A.: Open source code (2020). <https://github.com/adgaudio/ietk-ret>
5. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(12), 2341–2353 (2011). <https://doi.org/10.1109/TPAMI.2010.168>
6. He, K., Sun, J., Tang, X.: Guided image filtering. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(6), 1397–1409 (2013). <https://doi.org/10.1109/TPAMI.2012.213>
7. Lee, S., Yun, S., Nam, J.H., Won, C.S., Jung, S.W.: A review on dark channel prior based image dehazing algorithms. *EURASIP J. Image Video Process.* **2016**(1), 4 (2016). <https://doi.org/10.1186/s13640-016-0104-y>
8. Narasimhan, S.G., Nayar, S.K.: Vision and the atmosphere. *IJCV* **48**(3), 233–254 (2002)
9. Petrou, M., Petrou, C.: *Image Processing: The Fundamentals*, pp. 357–360. Wiley, Chichester (2011)
10. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
11. Sahasrabudhe, V., et al.: Indian diabetic retinopathy image dataset (IDRID) (2018). <https://doi.org/10.21227/H25W98>
12. Savelli, B., et al.: Illumination correction by dehazing for retinal vessel segmentation. In: 2017 IEEE 30th International Symposium on Computer-Based Medical Systems (CBMS), pp. 219–224 (June 2017). <https://doi.org/10.1109/CBMS.2017.28>
13. Smailagic, A., Sharan, A., Costa, P., Galdran, A., Gaudio, A., Campilho, A.: Learned pre-processing for automatic diabetic retinopathy detection on eye fundus images. In: Karray, F., Campilho, A., Yu, A. (eds.) ICIAR 2019. LNCS, vol. 11663, pp. 362–368. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-27272-2_32
14. Tan, R.T.: Visibility in bad weather from a single image. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8 (June 2008). <https://doi.org/10.1109/CVPR.2008.4587643>
15. Wang, Y., Zhuo, S., Tao, D., Bu, J., Li, N.: Automatic local exposure correction using bright channel prior for under-exposed images. *Signal Process.* **93**(11), 3227–3238 (2013). <https://doi.org/10.1016/j.sigpro.2013.04.025>. <http://www.sciencedirect.com/science/article/pii/S0165168413001680>