Blind Dehazing Using Internal Patch Recurrence

Yuval Bahat Michal Irani
Dept. of Computer Science and Applied Mathematics
The Weizmann Institute of Science*, ISRAEL

Abstract

Images of outdoor scenes are often degraded by haze, fog and other scattering phenomena. In this paper we show how such images can be dehazed using internal patch recurrence. Small image patches tend to repeat abundantly inside a natural image, both within the same scale, as well as across different scales. This behavior has been used as a strong prior for image denoising, super-resolution, image completion and more. Nevertheless, this strong recurrence property significantly diminishes when the imaging conditions are not ideal, as is the case in images taken under bad weather conditions (haze, fog, underwater scattering, etc.). In this paper we show how we can exploit the deviations from the ideal patch recurrence for "Blind Dehazing" - namely, recovering the unknown haze parameters and reconstructing a haze-free image. We seek the haze parameters that, when used for dehazing the input image, will maximize the patch recurrence in the dehazed output image. More specifically, pairs of co-occurring patches at different depths (hence undergoing different degrees of haze) allow recovery of the airlight color, as well as the relative-transmission of each such pair of patches. This in turn leads to dense recovery of the scene structure, and to full image dehazing.

1. Introduction

Images of outdoor scenes are often degraded by a scattering medium (e.g., aerosols, dust particles and water droplets). Haze, fog and even underwater scattering are such phenomena, whose degradation effect on the resulting images grows with scene depth. This degradation is double folded: (i) the irradiance $L(x)$ emitted from scene points is attenuated due to scattering caused by haze particles along the line of sight, and (ii) ambient *airlight* A is also scattered by haze particles, causing some of it to be reflected into the line of sight and reach the camera. The resulting image $I(x)$ exhibits reduced contrast and distorted colors,

*This work was funded in part by the Israel Science Foundation (Grant 931/14).

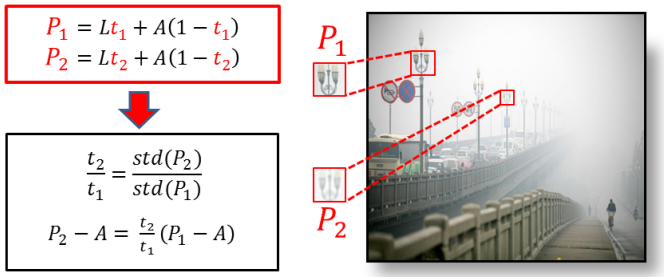

Figure 1. **Constraints induced by co-occurring patches:** *While a pair of co-occurring patches P_1 and P_2 look very different in the hazy image $I(x)$, they originate from the same (unknown) dehazed patch $L(x)$. This provides a strong prior for image dehazing, constraining the relative transmission parameters t_2/t_1 and the shared airlight A of such pairs of patches (see Sec. 2).*

and is typically modeled [7, 15, 17] by:

$$I(x) = t(x)\, L(x) + (1 - t(x))\, A, \tag{1}$$

where $t(x)$ is the corresponding attenuation factor, known as the *transmission*

$$t(x) = e^{-\beta Z(x)}, \tag{2}$$

where β is a scattering coefficient and $Z(x)$ is the distance to the scene point. $t(x)$ is typically assumed to be the same for all three color channels (R,G,B) [4, 7].

Blind image dehazing, namely, recovering the haze parameters A and $t(x)$ and inverting Eq. (1) to recover a haze-free image, is an under-constrained problem. Different dehazing methods proposed different ways to cope with this problem. Some assume having multiple images of the same scene (e.g., taken under different polarizations [15, 16] or under different weather conditions [12, 13, 14]).

More recently, methods based on a single image have been proposed, which tackle the lack of constraints by incorporating various priors. Tan [18] assumed maximal local contrast in the dehazed image (which occasionally leads to an exaggerated contrast). He et al. [7] recovered the t-map $t(x)$ by using the *Dark Channel Prior*, and Tang et al. [19] combined these notions and others into a learning framework. Zhu et al. [20] performed image dehazing by

Figure 2. **Dehazing by maximizing patch recurrence.** *Our algorithm seeks the haze parameters such that, when used for dehazing the input image $I(x)$, will* maximize the patch recurrence *in the output haze-free image $L(x)$. This maximization results in successful dehazing even of extremely distant forest regions (best viewed on screen).*

learning a linear color attenuation model. All four methods [7, 18, 19, 20] estimate A by implicitly assuming that regions at infinity (e.g., the sky) are visible in the image, and that the airlight color A is the brightest color among these regions. Fattal [4] assumed that the scene irradiance $L(x)$ and transmission map $t(x)$ are locally uncorrelated, and used a user assisted approach to estimate A. In his later works [5, 17], haze parameters were recovered by assuming that the distribution of pixels inside constant albedo patches can be described as lines in RGB space.

In this paper we propose to use the *internal patch recurrence property* as a strong prior for single-image blind dehazing. Small image patches (e.g., 5×5, 7×7) tend to repeat abundantly inside a single natural image, both within the same scale, as well as across different scales of the image. This patch recurrence property has been used as a strong prior to solve a variety of ill-posed vision problems, including super-resolution [6], image denoising [1,2], image completion [3], and more. More recently, it was shown [9, 10] that under significant camera blur, the recurrence property significantly diminishes, thus encoding information about the unknown global blur kernel. This was used for blind super-resolution [9] and blind deblurring [10].

Nevertheless, degradation is not limited to blur; other types of degradations in the imaging process may also lead to diminished patch recurrence. Deviations from the ideal patch recurrence encode valuable information about the unknown degradation process, in general. In particular, images taken under bad weather conditions (haze, fog, etc.) suffer from diminished patch recurrence. Recurring patches at different depths undergo different amounts of haze, hence no longer look the same (e.g., see the patches P_1 and P_2 in Fig. 1). Nevertheless, these *differences* between such "co-occurring patches" (patches with high normalized-correlation) allow recovery of their shared airlight color A and their *relative* transmission parameters. Combining the information from a *sparse* set of co-occurring pairs of patches in the image, yields the global airlight color A and a *dense* t-map $t(x)$. This in turn allows for complete image dehazing. Generally speaking, we seek the haze parameters $t(x)$ and A such that, when used for dehazing the input image $I(x)$, will maximize the patch recurrence property in the output haze-free image $L(x)$ (See Fig. 2).

Unlike [7, 18, 19, 20], our method does not require scene points at infinite distance to appear in the image, nor does it make any prior assumption about the brightness of A.

A brief overview of our algorithm is outlined below, and is detailed in Sections 2,3,4.

Overview of the Algorithm:

Input: Hazy image $I(x)$
Output: Airlight $\hat{A}$, t-map $\hat{t}(x)$, dehazed image $\hat{L}(x)$

1. **Detect "co-occurring pairs":** (Sec. 3)
   (a) Extract structured (high-variance) patches from image $I(x)$.
   (b) Search for matching patches (with high normalized-correlation).
2. **Extract Pairwise haze parameters for each pair:**
   (a) Estimate relative t-values t_2/t_1 (Sec. 2.1).
   (b) Estimate their shared airlight (Sec. 2.2).
3. **Estimate Global haze parameters:**
   (a) Recover global airlight $\hat{\mathbf{A}}$ using all pairwise airlight estimates (Sec. 4.1).
   (b) Recover *dense* t-map $\hat{\mathbf{t}}(\mathbf{x})$ which: (i) is smooth, and (ii) satisfies the *sparse* pairwise constraints (Sec. 4.2).
4. **Recover haze-free $\hat{L}(x)$:** $\hat{L}(x) = \frac{I(x)-\hat{A}}{\hat{t}(x)} + \hat{A}$.

The rest of this paper is organized as follows. Sec. 2 introduces the constraints induced by pairs of co-occurring patches. Sec. 3 explains how such pairs of patches can be

found reliably in a hazy image. Sec. 4 describes our algorithm for recovering global dense haze parameters. Experimental results are presented in Sec. 5, obtaining state-of-the-art performance on established benchmark datasets.

2. Constraints Induced by Pairs of Patches

In this section we show how *a single pair* of "co-occurring patches" at different depths (i.e., the same local structure undergoing different degrees of haze) encodes valuable information about the unknown haze parameters: the *airlight* color, as well as the *relative-transmission* of the two patches. Most existing methods (e.g., [7, 15, 17, 18, 20]) assume a uniform global airlight A for the entire image. Our current algorithm (Sec. 4) also follows this assumption. Nevertheless, in this section we show that pairs of patches can potentially provide information for the more general case, when the airlight is spatially varying (which may happen e.g., when the sun is low in the sky, or when the scene is illuminated by multiple light sources).

Let $P_1[x]$ and $P_2[x]$ denote a pair of small co-occurring patches[1] (7×7) that emanate from the *same underlying haze-free patch* $L[x]$, but are located at different scene depths, $Z_1[x]$ and $Z_2[x]$. Since these patches are very small, we can assume constant depths in each patch, $Z_1[x] \equiv Z_1$ and $Z_2[x] \equiv Z_2$, hence also constant transmission values, t_1 and t_2 (since $t = e^{-\beta Z}$). We further assume that the airlight is locally uniform within each patch ($A_1[x] \equiv A_1$ and $A_2[x] \equiv A_2$), even if not uniform in the entire image. According to Eq. (1):

$$\begin{aligned} P_1[x] &= L[x]\, t_1 + A_1\, (1 - t_1) \\ P_2[x] &= L[x]\, t_2 + A_2\, (1 - t_2) \end{aligned} \tag{3}$$

Note that under ideal (haze-free) imaging conditions, the patches P_1 and P_2 should be identical (and equal to their shared underlying haze-free patch L). However, due to the haze and their different depths, they look quite different (See Fig. 1).

Eliminating the Airlight Component: When the transmission t and airlight A in Eq. (1) are locally uniform, subtracting the local *mean-color* eliminates the airlight component. This observation was made by Narasimhan [11] for uniformly colored regions. However, it is generally *true for almost any small image patch* (uniform or textured), as long as A and t are locally uniform within the patch (an assumption violated only for patches at depth discontinuities). The elimination of the airlight is obtained by subtracting the mean value of each side of Eq. (3), which yields:

$$\begin{aligned} \tilde{P}_1[x] &= \tilde{L}[x]\, t_1 \\ \tilde{P}_2[x] &= \tilde{L}[x]\, t_2 \end{aligned} \tag{4}$$

where $\tilde{P}_i[x] = P_i[x]$-mean(P_i) and $\tilde{L}[x] = L[x]$-mean(L) (the mean is computed separately for each color channel within each patch). Note that the elimination of the airlight A_i is done independently for each patch, and *can be done even if different patches have different airlight colors.*

2.1. Obtaining Pairwise Relative Transmission

Given a pair of co-occurring patches, we can now proceed to obtain their *relative transmission parameters* t_2/t_1. Eq. (4) entails:

$$\begin{aligned} ||\tilde{P}_1|| &= ||\tilde{L}\, t_1|| = ||\tilde{L}||\, t_1 \\ ||\tilde{P}_2|| &= ||\tilde{L}\, t_2|| = ||\tilde{L}||\, t_2 \end{aligned} \tag{5}$$

Recall that t_1 and t_2 are scalars (between 0 and 1). Thus:

$$\frac{t_2}{t_1} = \frac{||\,\tilde{P}_2||}{||\,\tilde{P}_1||}, \tag{6}$$

Assuming l_2 norm, the ratio between the transmission parameters of two co-occurring patches, P_1 and P_2, reduces to a *simple ratio of their standard-deviations*:

$$\frac{t_2}{t_1} = \frac{\text{std}(P_2)}{\text{std}(P_1)}. \tag{7}$$

Note that the relative transmission parameters t_2/t_1 can be recovered *even if the airlight colors of the two patches* P_1 *and* P_2 *are different* ($A_1 \neq A_2$).

Without loss of generality, we will assume from here on that $0 \leq t_2/t_1 \leq 1$, since we can always denote by P_2 the patch with the smaller standard-deviation, and by P_1 the patch with the larger standard-deviation.

2.2. Obtaining Pairwise Shared Airlight

Assuming that a pair of co-occurring patches share the same airlight color, $A_1 = A_2 = A$, we can proceed to obtain this *shared airlight* (even if not shared by all other image patches). Eq. (3) can be rewritten as:

$$\begin{aligned} P_1[x] - A &= (L[x] - A)\, t_1 \\ P_2[x] - A &= (L[x] - A)\, t_2 \end{aligned} \tag{8}$$

which entails that:

$$(P_2[x] - A) = \frac{t_2}{t_1}\, (P_1[x] - A) \tag{9}$$

Combining the constraints from Eqs. (4) and (9) yields:

$$(P_2[x] - A)\, \tilde{P}_1[x] - (P_1[x] - A)\, \tilde{P}_2[x] = 0, \tag{10}$$

[1] We use square brackets $[x]$ to denote coordinates within a columned-stacked patch P, and regular parentheses (x) to denote coordinates in the entire image.

This last constraint holds for all pixels x in those two patches, thus providing an overdetermined set of linear constraints on each of the three airlight color components, $A = (A^R, A^G, A^B)$. The shared airlight A can thus be estimated, e.g., using Least-Squares:

$$A = \frac{\left[\tilde{P}_2 - \tilde{P}_1\right]^T \left[P_1 \circ \tilde{P}_2 - P_2 \circ \tilde{P}_1\right]}{\| \tilde{P}_2 - \tilde{P}_1 \|^2}, \qquad (11)$$

where $\circ$ denotes the element-wise vector multiplication (also known as the Hadamard product).

Note that our recovery of the airlight A does *not* require having scene points at infinity in the image. This is in contrast to [7, 18, 19, 20], which rely on this assumption.

The pairwise constraints derived above typically provide information about the t-ratios and shared airlight only for a sparse set of patches in the image (those reliably detected as co-occurring patches). In Sec. 4 we show how these can be used to recover *dense* and global haze parameters for the entire image.

3. Unveiling Patch Recurrence

Informative co-occurring patches image the same local structure $L[x]$, but at different scene depths. The variations between such patches is what allows recovery of their relative transmission parameters, but also pose a challenge on reliable detection of such co-occurring pairs. In this section we explain how to "unveil" patch recurrences, despite being "fogged" by the haze.

3.1. Searching for Normalized-Nearest Neighbors

The obscuring effect of the airlight A is removed in the mean-free patches $\tilde{P}_1$ and $\tilde{P}_2$ (Eq. (4)). However, their different transmissions still obscure their similarity. Therefore, we normalize the mean-free patches, $\tilde{P}_1$ and $\tilde{P}_2$, by their norms (which is equivalent to normalizing by the standard-deviation of each patch, in the case of l_2 norm):

$$\frac{\tilde{P}_1[x]}{||\tilde{P}_1||} = \frac{t_1\tilde{L}[x]}{t_1||\tilde{L}||} = \frac{\tilde{L}[x]}{||\tilde{L}||} \qquad (12)$$
$$\frac{\tilde{P}_2[x]}{||\tilde{P}_2||} = \frac{t_2\tilde{L}[x]}{t_2||\tilde{L}||} = \frac{\tilde{L}[x]}{||\tilde{L}||}.$$

Yielding
$$\tilde{P}_1/||\tilde{P}_1|| = \tilde{P}_2/||\tilde{P}_2||. \qquad (13)$$

Thus, normalizing the mean-free version of all hazy image patches unveils their recurrence property. Pairs of co-occurring patches can now be detected by applying Nearest-Neighbors (NN) search on the normalized patches.[2]

[2]This is equivalent to finding pairs of patches with high normalized correlation. The advantage of first normalizing each patch and then searching for NNs, is that it allows to use also l_2-based NN search algorithms.

Since our final goal is to recover a haze-free image with maximal patch recurrence across multiple scales, we search for normalized-NNs (co-occurring patches) across multiple scales of the input hazy image $I(x)$. Note that scaling down the hazy image does not change the physical parameters of the scene (the airlight A at infinity, the depth of scene points Z, or the haze scattering parameters β); it only has a zoom-out effect. Thus, scaling down the hazy image only coarsens the *transmission-map* (t-map), but preserves its absolute range of values.

Given a hazy input image $I(x)$, we generate multi-scale versions of it: $\{I^{sc}\}$, with scale factors, e.g., $sc = 1, 0.75, 0.5, 0.25$. We apply NN-search only to patches which have high std (above 25 grayscale levels). Seeking NNs only for high-std patches prevents overfitting the noise and increases the numerical stability. For each such patch, we seek k-NNs [3] in multiple image scales. We discard pairs of patches whose normalized-correlation is lower than 0.7.

3.2. Handling Outliers and Pairwise Errors

While the normalization in Eq. (13) unveils the similarity of co-occurring patches, it also yields many false matches. A pair of patches with *different haze-free patches* $L_1[x]$ and $L_2[x]$, which satisfy $L_2[x] = aL_1[x] + b$ (for any scalars a and b), will be falsely identified as a co-occurring pair by our NN search (e.g., two *arbitrary* vertical edges).

We next show that the *physical lower bound on the transmission $t(x)$ suffices for detecting such false matches*.

A Lower-Bound on $\mathbf{t(x)}$: Transmission values must satisfy $t(x) \geq 0, \ \forall x$. The haze free image $L(x)$ must further satisfy $L(x) \geq 0, \ \forall x$. Incorporating these two constraints in Eq. (1) yields:

$$I(x) = L(x)\, t(x) + A\,(1 - t(x)) \geq A\,(1 - t(x)). \quad (14)$$

The constraint in Eq. (14) holds for each pixel, and for each of its color channels. Hence we can define for every pixel a **Lower-Bound $\mathbf{t_{LB}}(x)$** on $t(x)$, by rearranging Eq. (14) and taking the maximum over R, G, B:

$$t(x) \geq \max_{c \in R,G,B} \left\{1 - \frac{I_c(x)}{A_c}\right\} \ \triangleq \ \mathbf{t_{LB}}(\mathbf{x}). \qquad (15)$$

This lower bound shares a similar reasoning with the *Dark Channel Prior* (DCP) exploited by He et al. [7], but unlike the DCP, it holds for every pixel x in the image.

Note that computing $t_{LB}[x]$, given the estimated $\hat{A}$ of a pair of patches, is very easy, and does not require complex estimation of the true absolute $t[x]$. It simply involves computing the term in (15) for every pixel in the patch.

We later use this lower bound in Sec. 4 for constrained optimization, when recovering the t-map $t(x)$. Nevertheless, this lower bound can also be used to identify outlier

[3]In our implementation we used a KD-tree search (its approximate-NN version with $\epsilon = 1$), with $k = 9$ (i.e., 9 NNs for each patch).

pairs of patches, as shown next.

Claim: (t_{LB} preserves t-ratios)
A true (inlier) co-occurring pair of patches (that share the same $L[x]$) satisfies

$$\frac{t_{LB_2}[x]}{t_{LB_1}[x]} = \frac{t_2[x]}{t_1[x]}. \tag{16}$$

Proof: Rewriting $t_{LB}[x]$ of Eq. (15) in terms of $L[x]$, $t[x]$ and A, yields

$$t_{LB}[x] = \max_{c\in R,G,B}\left\{1 - \frac{P_c[x]}{A_c}\right\} =$$
$$= \max_{c\in R,G,B}\left\{1 - \frac{L_c[x]t[x] + (1-t[x])A_c}{A_c}\right\} =$$
$$= t[x]\left(1 - \min_{c\in R,G,B}\left\{\frac{L_c[x]}{A_c}\right\}\right). \tag{17}$$

Two true co-occurring patches share the same underlying $L[x]$, hence the ratio of their $t_{LB}[x]$ reduces to their true t-ratio:

$$\frac{t_{LB_2}[x]}{t_{LB_1}[x]} = \frac{t_2[x]}{t_1[x]}\frac{\left[1 - \min_{c\in R,G,B}\left\{\frac{L_c[x]}{A_c}\right\}\right]}{\left[1 - \min_{c\in R,G,B}\left\{\frac{L_c[x]}{A_c}\right\}\right]} = \frac{t_2[x]}{t_1[x]}. \tag{18}$$

■

In contrast, a false (outlier) pair rarely satisfies Eq. (16), since

$$\frac{t_{LB_2}[x]}{t_{LB_1}[x]} = \frac{t_2[x]}{t_1[x]}\frac{1 - \min_{c\in R,G,B}\left\{\frac{aL_c[x]+b_c}{A_c}\right\}}{1 - \min_{c\in R,G,B}\left\{\frac{L_c[x]}{A_c}\right\}} \neq \frac{t_2[x]}{t_1[x]}. \tag{19}$$

Thus, given a pair of patches with high normalized correlation, we estimate their likelihood to be an inlier pair as follows: we estimate their shared airlight $\hat{A}$, and use that to estimate $t_{LB_1}[x]$ and $t_{LB_2}[x]$ (for each pixel in the patches). We further estimate their t-ratio from the hazy image $\frac{t_2}{t_1} = \frac{\text{std}(P_2)}{\text{std}(P_1)}$ (Eq. (7)). Then, for each pixel in the patch, we compute the difference $\left|\frac{t_{LB_2}[x]}{t_{LB_1}[x]} - \widehat{\frac{t_2}{t_1}}\right|$ between the *estimated* t-ratio $\widehat{\frac{t_2}{t_1}}$ and the t_{LB}-ratio. The average over all pixels in the pair of patches, constitutes a good indicator for the inlier/outlier reliability of the pair. Patches for which this average difference exceeds a certain threshold (0.07 in our current implementation) are detected as outliers.

4. The Algorithm

So far we saw how the recurrence of patches undergoing different amount of haze induces constraints on their shared airlight A and their relative transmission parameters t_2/t_1.

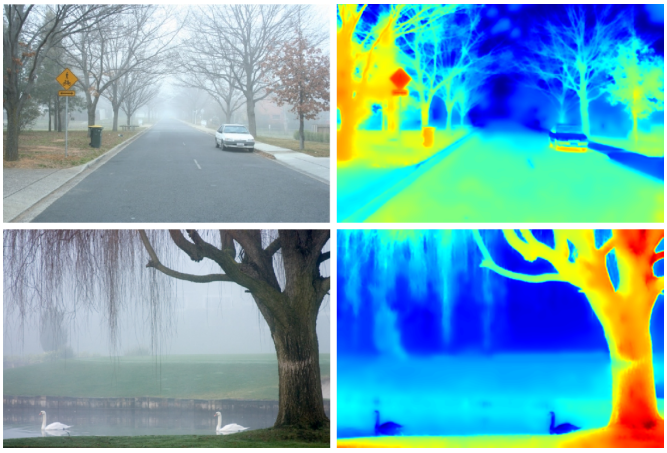

Figure 3. **Recovered t-maps.** *Red/blue indicate large/small $t(x)$, which correspond to near/far scene points.*

These constraints, however, are imposed only at sparse image locations (where co-occurring patches were reliably detected), and the t-ratios provide only relative quantities, and not absolute transmission values. In this section we explain how to exploit those sparse relative pairwise constraints to obtain a global airlight A and a dense absolute t-map $t(x)$. Once these are recovered, they lead to the recovery of the haze-free image $L(x)$ using Eq. (1):

$$L(x) = \frac{I(x) - A}{t(x)} + A. \tag{20}$$

4.1. Recovering the Global Airlight

Let $\{(P_{k_1}, P_{k_2})\}_{k=1}^{K}$ be all the pairs of reliably detected co-occurring patches (after outlier rejection) in the input image $I(x)$. Let $\left\{\hat{A}_k\right\}_{k=1}^{K}$ be their pairwise estimated airlights, computed using Eq. (11). The single global airlight estimate $\hat{A}$ for the entire image could in principle be computed as the average $\hat{A} = \frac{1}{K}\sum_k \hat{A}_k$. However, not all pairs of patches are equally informative for recovering the global A. Pairs of patches originating from significantly different depths (hence undergoing different haze) are more informative than those with similar amounts of haze. Thus our global airlight estimate $\hat{A}$ is a weighted average

$$\hat{A} = \frac{\sum_k w_k \hat{A}_k}{\sum_k w_k}, \tag{21}$$

giving higher weights w_k to more informative pairs of patches. Informative pairs should have a very large pairwise t-ratio ($\frac{t_1}{t_2} \gg 1$). But this on its own is not enough; two very distant patches with very small t_1 and t_2 may have a large t-ratio, but are not very informative. Therefore we also want a large pairwise t difference $(t_1 - t_2)$. Since we do not know the absolute t at this stage, we substitute it with weights based on the lower bound t_{LB}, defined in Eq. (15).

We found that setting the weights to

$$w_k = \left[(t_{LB_1} - t_{LB_2}) \left(\frac{t_{LB_1}}{t_{LB_2}} - 1 \right) \right]^2 \qquad (22)$$

achieves the desired weighting effect, where t_{LB_i} of patch P_i $(i = 1, 2)$ is defined as $t_{LB_i} = \max_{x \in P_i} t_{LB_i}[x]$.

Since the weights w_k depend on t_{LB}, which in turn depend on the desired estimate $\hat{A}$, we compute (21) and (22) iteratively. Our initial guess for $\hat{A}$ is the regular average of all $\hat{A}_k$'s, and we iterate until the rate of change in $\hat{A}$ is very small (typically up to 20 iterations). Our experiments show that this iterative weighing process reduces the estimation error of A by $\sim$25% compared to the initial guess.

4.2. Recovering the Transmission Map $\mathbf{t(x)}$

Given the recovered airlight $\hat{A}$, we seek a dense transmission map $t(x)$ such that, when dehazing the input image $I(x)$ with this transmission map, the resulting *haze-free* image $L(x)$ will have maximal internal patch recurrence (both within and across different scales of $L(x)$). In particular, we seek the t-map $t(x)$ that will minimize:

$$\arg\min_{t(x)} \{ \underbrace{\rho_L(t(x))}_{\text{patch-recurrence-term}} + \underbrace{\lambda \rho_s(t(x))}_{\text{smoothness-term}} \} \ \ s.t. \ \underbrace{t_{LB}(x) \le t(x) \le 1}_{\text{constrained range}} \qquad (23)$$

where $\rho_L(t)$ is the patch recurrence constraint on the resulting dehazed image $L(x)$ (expressed in terms of the unknown $t(x)$), $\rho_s(t)$ is a smoothness term on the recovered t-map, and λ was empirically set to 0.5. The t-map $t(x)$ is further constrained, at every pixel x, to lie between the lower bound $t_{LB}(x)$ (defined in Eq. (15)) and the upper bound 1.

The patch-recurrence data term: $\rho_L(t)$ is computed over all reliably detected pairs of co-occurring patches (*within* and *across* scales of the input image $I(x)$). A pair of co-occurring patches P_1 and P_2 (that share the same $L[x]$), should be equal to each other after dehazing, namely: $L_1[x] = L_2[x]$. This can be rewritten in terms of $t(x)$ as: $\frac{P_1 - \hat{A}}{t_1} + \hat{A} = \frac{P_2 - \hat{A}}{t_2} + \hat{A}$, and can be further turned into a linear expression in the unknown $t(x)$: $(P_1 - \hat{A})t_2 = (P_2 - \hat{A})t_1$. Therefore, we would like for the t-map $t(x)$ to minimize the following patch-recurrence penalty term:

$$\rho_L(t(x)) = \sum_{k \in \text{pairs}} \| \left(P_{k_1} - \hat{A} \right) t_{k_2}(x) - \left(P_{k_2} - \hat{A} \right) t_{k_1}(x) \|^2 . \qquad (24)$$

where $t_{k_1}(x)$ and $t_{k_2}(x)$ are the transmission values corresponding to the center pixels of patches P_{k_1} and P_{k_2}, respectively.

In order to have a spatially denser patch-recurrence constraint, we augment the set of co-occurring patches exhibiting high normalized-correlation (described in Sec. 3), by searching for pairs of patches exhibiting low *Sum of Squared Differences* (SSD). These low-SSD pairs of patches not only co-occur (share the same underlying $L[x]$), but also undergo the same degradation by haze, which means they share the same $t(x)$. Hence for them we minimize this modified version of the term in Eq. (24):

$$\sum_{k \in \substack{low-SSD \\ pairs}} \| t_{k_2}(x) - t_{k_1}(x) \|^2 . \qquad (25)$$

Combining the two sets of pairs increases the spatial density of the patch recurrence constraint. Our experiments show that more than 20% of the t-map pixels are directly effected by it in any image we have experimented with to-date (more than 50 images).

The smoothness term: $\rho_s(t)$ is computed over all pixels in the image:

$$\rho_s(t(x)) = \sum_{\text{all pixels } x} w(x) \ \| \nabla \log(t(x)) \|^2 \qquad (26)$$

where $w(x)$ is a decreasing sigmoid function[4] of $\nabla L(x)$. Recall from Eq. (2) that $\log(t(x)) \propto Z(x)$. Thus, the "depth map" $\log(t(x))$ is penalized for having large derivatives (large depth discontinuities), *except* at strong edges in the dehazed image $L(x)$. Since we do not know $L(x)$ in advance (as it depends on the estimated $t(x)$), we solve the objective function in Eq. (23) using iterative minimization.

Iterative minimization: We solve Eq. (23) iteratively, where each iteration is a constrained non-convex optimization problem. We start with an initial guess $t(x) = t_{LB}(x)$. Note that $t_{LB}(x)$ already satisfies the constrained range of $t(x)$, as well as the t-ratio constrains (see claim in Sec. 3.2), but is usually not smooth.

At each iteration we approximate $L(x)$ by using Eq. (20) with $t(x)$ from the previous iteration. We re-estimate the weights $w(x) = \text{sigmoid}(\nabla L(x))$ of the smoothness term using the current $L(x)$, and re-estimate $t(x)$ by minimizing Eq. (23). This process is iterated until there is no significant change in the error term (typically around 20 iterations).

The above iterative process results in our final t-map $t(x)$ and our dehazed image $L(x)$. Examples of recovered t-maps are shown in Fig. 3. Since small amount of haze is present even on clear days, totally haze-free images tend to look unnatural [7]. Therefore, we allow for a small amount of haze to be left in our final dehazed output images $\hat{L}(x)$, namely: $\hat{L}(x) = \alpha\, L(x) + (1 - \alpha)\, I(x)$. We used $\alpha = 0.85$ in our displayed images.

5. Experiments

We evaluated our results on a large database of fifty images (the database used by [5], enriched with images we

[4] We used $w(x) = \text{sigmoid}(\nabla L(x)) = \left(1 + e^{48(\|\nabla L(x)\| - 0.1)}\right)^{-1}$.

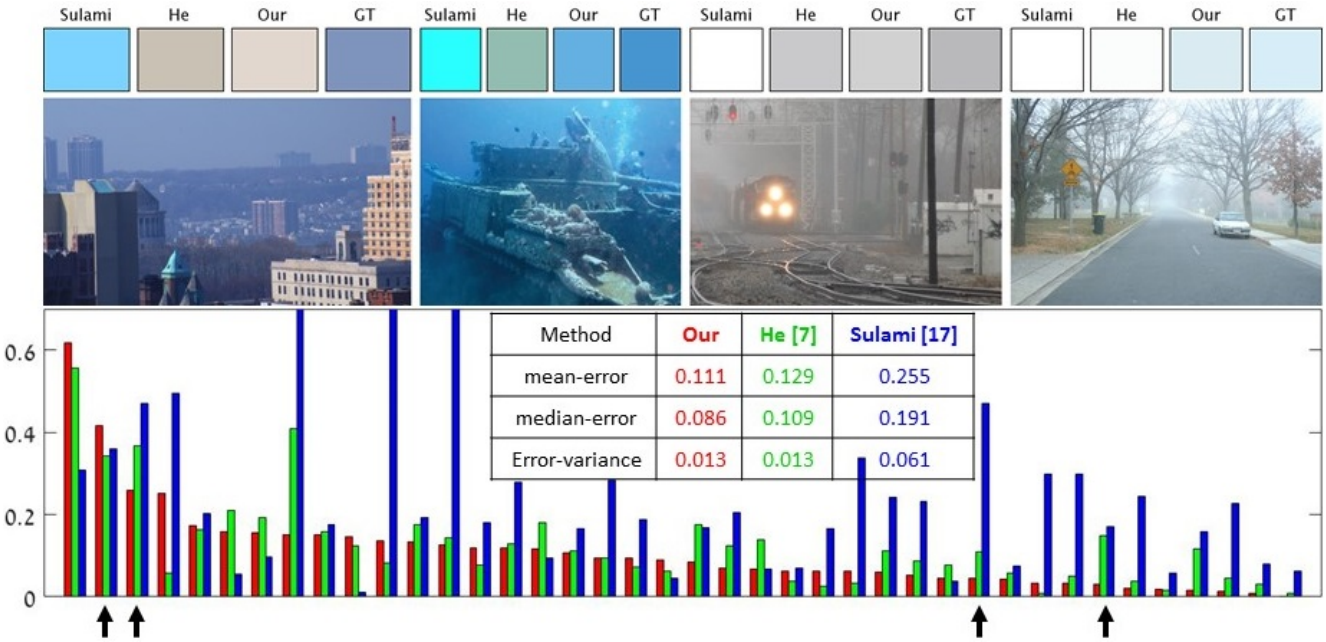

Figure 4. **Evaluating the accuracy of recovered airlight.** *Top: Examples of hazy images, along with their manually extracted ground truth airlight color (**GT**), and the estimated airlights of **Sulami** et al [17], **He** et al [7], and **Ours**. Bottom: Errors estimated on* 40 *heavily hazed images (from which GT airlight could be reliably extracted). Black arrows indicate bars corresponding to the example images at the top (same order of appearance). The table shows that our method reduces the mean and median errors by* 14% *and* 21%*, respectively, compared to the state-of-the-art method of [7].*

added). Please refer to our project website www.wisdom.weizmann.ac.il/~vision/BlindDehazing/ to see all our results.

In order to quantitatively evaluate the quality of the recovered airlight A, we performed the following experiment. Ground truth airlight is typically not available in real-world hazy images. Nevertheless, under the assumption of a single global airlight, if the image exhibits strong haze, and contains points at infinite distance (e.g., the sky), then their color should coincide with the airlight color A [15]. Therefore, in severely hazed images where the sky or distant scene points are visible, we can manually obtain a good approximation of the ground-truth airlight.

Forty (heavily-hazed) images in the database contained a visible piece of sky or a very distant scene point. We manually marked that region, and computed its average color. This served as the ground-truth airlight A for our global airlight estimation. The graph in Fig. 4 plots the errors between the estimated airlight and the ground-truth airlight, computed for 3 methods (using the online codes published by the authors): Ours (red bars), He *et al.* [7] (green bars), and Sulami *et al.*[5] [17] (blue bars). The results of the *mean* and *median* errors are further reported in the table in Fig. 4. The results show that our estimated global airlight A for these images is more accurate than both methods. Note that those images contain distant scene points, thus satisfy the assumption of [7, 18, 19, 20] (an assumption that neither our method, nor [17], rely on). Despite that, the results show that our method provides a decrease in the average and median estimation errors by 14% and 21%, respectively, compared to the current state-of-the-art method by He et al. [7].

Figs. 2,5,6 show results of our dehazing algorithm. Fig. 2 examplifies the power of the patch recurrence property. It shows how the maximization of patch recurrence in the haze-free image results in successful dehazing even of extremely distant and very hazy forest regions.

Figs. 5,6, as well as many more images on our project page, show a comparison of dehazing results, along with the corresponding recovered airlight colors A, between our method and that of [7], on all the images in the dataset. We chose to compare to [7] because it is the state-of-the-art *end-to-end* blind dehazing algorithm (both recovering the airlight A and the dehazed image, in a single framework). Moreover, its online available code allowed the comparison also on the challenging images we added to the dataset. In addition, we present the results using the recently released commercial dehazing tool[6] of Adobe Lightroom [8].

The images in Fig. 5 contain heavily hazed regions (either the sky or just dense haze regions), thus satisfy the underlying assumptions of [7, 18, 19, 20]. Interestingly, the airlight color A recovered by our method does not necessar-

[5]The online code available on the project page of [17] is not the full implementation of [17], but rather a more limited one which assumes a fixed albedo image (based on personal communication with the authors).

[6]We applied the dehazing effect of [8] with maximal (+100) dehazing settings on all images.

Figure 5. **Dehazing results on images containing heavily-hazed distant scene points:** *The hazy input image (first column), our blind dehazing result (second column), result by [7] (third column) and the results using [8] (last column). The recovered airlight color $\hat{A}$ of our and He's dehazing results is presented in a framed rectangle above it. (Best viewed on screen.)*

ily coincide with the brightest color of these regions (e.g. in image (a)), yet using it for dehazing yields better results.

Fig. 6 presents images that do not satisfy the assumptions in [7, 18, 19, 20]. It shows the advantage of our method in cases where the image does not contain regions of dense haze (images (a) and (b)), or in images taken underwater (image (c)), in which case the airlight is not bright, nor does it coincide with the brightest image region.

6. Conclusion

We showed that the unknown haze parameters can be recovered from a single hazy image, by seeking the haze parameters that, when "removed", maximize the patch recurrence in the output haze-free image. Pairs of co-occurring patches at different depths (undergoing different haze) allow recovery of their relative transmission and their shared airlight color. Unlike most blind dehazing methods, our method does not rely on imaging points at infinity, nor does it assume that the airlight is necessarily bright. While in this paper we assumed a single uniform airlight A, we showed that pairs of co-occurring patches encode information of their shared airlight color, even if not uniform in the image. Extending our algorithm to handle non-uniform airlight $A(x)$ is part of our future work.

References

[1] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, volume 2, pages 60–65, 2005.

[2] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *T-IP*, 16(8):2080–2095, 2007.

[3] A. Efros, T. K. Leung, et al. Texture synthesis by non-parametric sampling. In *ICCV*, pages 1033–1038, 1999.

[4] R. Fattal. Single image dehazing. In *ACM Transactions on Graphics (TOG)*, volume 27, page 72, 2008.

[5] R. Fattal. Dehazing using color-lines. *ACM Transaction on Graphics*, 34(13), 2014.

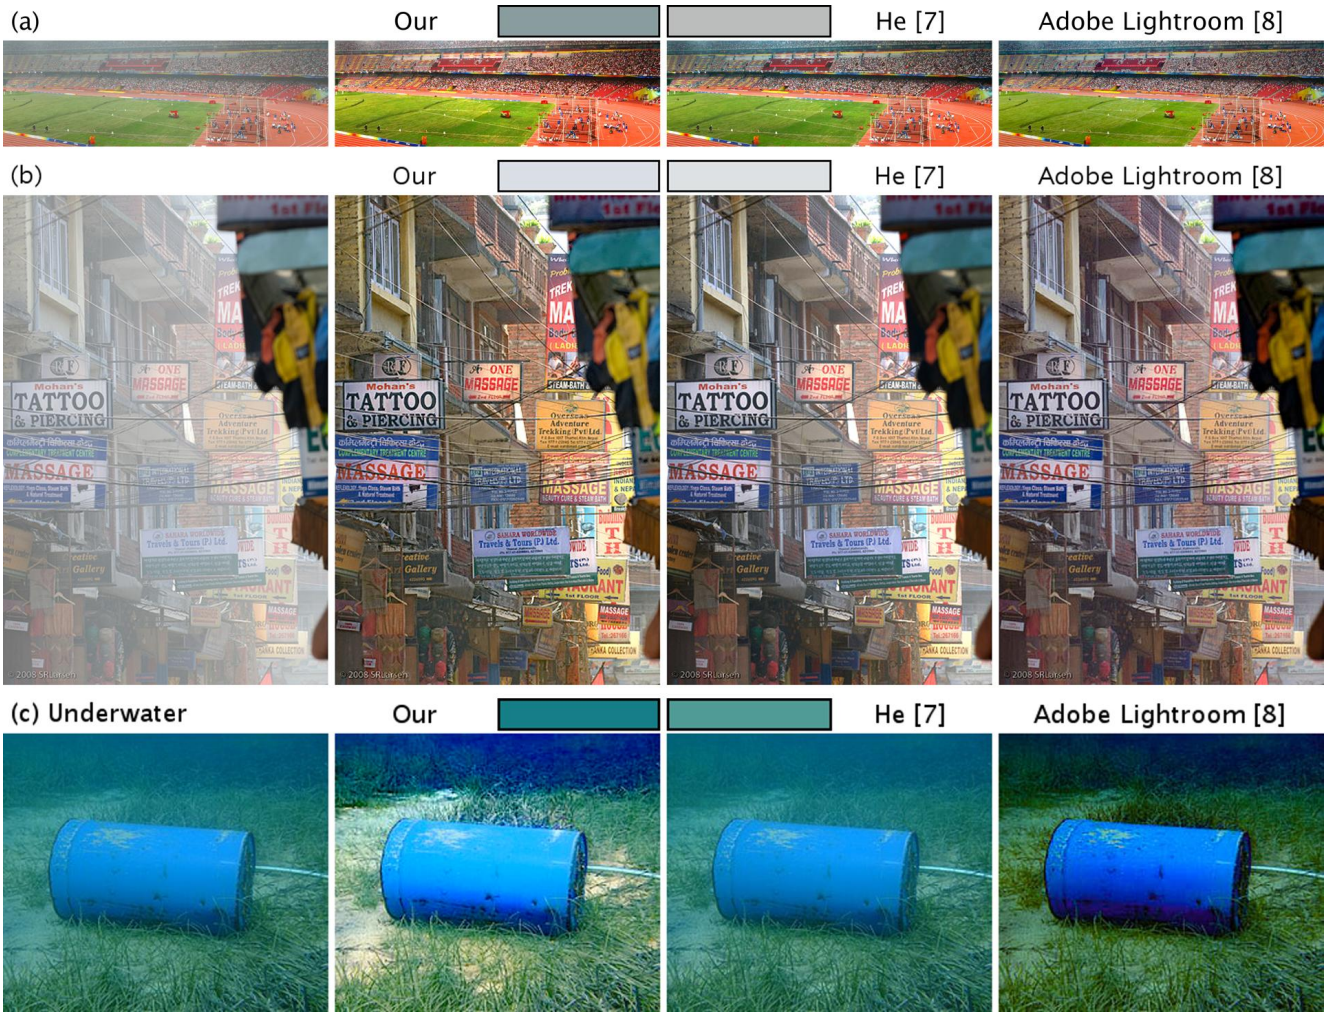

Figure 6. **Results on images with challenging A-recovery conditions:** *The hazy inputs (first column), our dehazing results (second column), results of [7] (third column) and the results using [8] (last column). The recovered airlight $\hat{A}$ of our and He's dehazing results is displayed in a rectangle above it. These images do not satisfy the assumptions of [7, 18, 19, 20]. Images (a) and (b) do not contain points at infinite distance. Image (c) is an underwater image, in which case the airlight is not bright, nor does it coincide with the brightest image regions. The advantage of our method is prominent on these images. Note the recovered distant red and yellow gallery seats in (a). Note the realistic colors of the grass and sand recovered underwater in (c). (Best viewed on screen.)*

[6] D. Glasner, S. Bagon, and M. Irani. Super-resolution from a single image. In *ICCV*, pages 349–356, 2009.

[7] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. *PAMI*, 33(12):2341–2353, 2011.

[8] Adobe Systems Incorporated. Adobe Photoshop Lightroom CC.

[9] T. Michaeli and M. Irani. Nonparametric blind super-resolution. In *ICCV*, pages 945–952, 2013.

[10] T. Michaeli and M. Irani. Blind deblurring using internal patch recurrence. In *ECCV*, pages 783–798. 2014.

[11] S. G. Narasimhan. *Models and algorithms for vision through the atmosphere*. PhD thesis, Citeseer, 2003.

[12] S. G. Narasimhan and S. K. Nayar. Chromatic framework for vision in bad weather. In *CVPR*, pages 598–605, 2000.

[13] S. G. Narasimhan and S. K. Nayar. Contrast restoration of weather degraded images. *PAMI*, 25(6):713–724, 2003.

[14] S. K. Nayar and S. G. Narasimhan. Vision in bad weather. In *ICCV*, volume 2, pages 820–827, 1999.

[15] Y. Y. Schechner, S. G. Narasimhan, and S. K. Nayar. Instant dehazing of images using polarization. In *CVPR*, 2001.

[16] S. Shwartz, E. Namer, and Y. Y. Schechner. Blind haze separation. In *CVPR*, volume 2, pages 1984–1991, 2006.

[17] M. Sulami, I. Glatzer, R. Fattal, and M. Werman. Automatic recovery of the atmospheric light in hazy images. In *ICCP*, pages 1–11, 2014.

[18] R. T. Tan. Visibility in bad weather from a single image. In *CVPR*, pages 1–8, 2008.

[19] K. Tang, J. Yang, and J. Wang. Investigating haze-relevant features in a learning framework for image dehazing. In *CVPR*, pages 2995–3002, 2014.

[20] Q. Zhu, J. Mai, and L. Shao. A fast single image haze removal algorithm using color attenuation prior. *IEEE Transactions on Image Processing*, 24(11):3522–3533, 2015.