



Vision and the Atmosphere

SRINIVASA G. NARASIMHAN AND SHREE K. NAYAR
Computer Science Department, Columbia University, New York, NY, USA

srinivas@cs.columbia.edu

nayar@cs.columbia.edu

Received January 31, 2001; Revised October 16, 2001; Accepted December 4, 2001

Abstract. Current vision systems are designed to perform in clear weather. Needless to say, in any outdoor application, there is no escape from “bad” weather. Ultimately, computer vision systems must include mechanisms that enable them to function (even if somewhat less reliably) in the presence of haze, fog, rain, hail and snow.

We begin by studying the visual manifestations of different weather conditions. For this, we draw on what is already known about atmospheric optics, and identify effects caused by bad weather that can be turned to our advantage. Since the atmosphere modulates the information carried from a scene point to the observer, it can be viewed as a mechanism of visual information coding. We exploit two fundamental scattering models and develop methods for recovering pertinent scene properties, such as three-dimensional structure, from one or two images taken under poor weather conditions.

Next, we model the chromatic effects of the atmospheric scattering and verify it for fog and haze. Based on this chromatic model we derive several geometric constraints on scene color changes caused by varying atmospheric conditions. Finally, using these constraints we develop algorithms for computing fog or haze color, depth segmentation, extracting three-dimensional structure, and recovering “clear day” scene colors, from two or more images taken under different but unknown weather conditions.

Keywords: physics based vision, atmosphere, bad weather, fog, haze, visibility, scattering, attenuation, airlight, overcast sky, scene structure, defog, dehaze

1. Computer Vision and the Weather

Virtually all work in computer vision is based on the premise that the observer is immersed in a transparent medium (air). It is assumed that light rays reflected by scene objects travel to the observer without attenuation or alteration. Under this assumption, the brightness of an image point depends solely on the brightness of a single point in the scene. Quite simply, existing vision sensors and algorithms have been created only to function on “clear” days. A dependable vision system however must reckon with the entire spectrum of weather conditions, including, haze, fog, rain, hail and snow.

The study of the interaction of light with the atmosphere (and hence weather) is widely known as atmospheric optics. Atmospheric optics lies at the heart of

the most magnificent visual experiences known to man, including, the colors of sunrise and sunset, the blueness of the clear sky, and the rainbow (see Minnaert (1954) and Henderson (1977)). The literature on this topic has been written over the past two centuries. A summary of where the subject as a whole stands would be too ambitious a pursuit. Instead, our objective will be to sieve out of this vast body of work, models of atmospheric optics that are of direct relevance to computational vision. Our most prominent sources of background material are the works of McCartney (1975) and Middleton (1952) whose books, though dated, serve as excellent reviews of prior work.

The key characteristics of light, such as its intensity and color, are altered by its interactions with the atmosphere. These interactions can be broadly classified

into three categories, namely, *scattering*, *absorption* and *emission*. Of these, scattering due to suspended particles is the most pertinent to us. As can be expected, this phenomenon leads to complex visual effects. So, at first glance, atmospheric scattering may be viewed as no more than a hindrance to an observer. However, it turns out that bad weather can be put to good use. The farther light has to travel from its source (say, a surface) to its destination (say, a camera), the greater it will be effected by the weather. Hence, bad weather could serve as a powerful means for coding and conveying scene structure. This observation lies at the core of our investigation; we wish to understand not only what bad weather does *to* vision but also what it can do *for* vision.

Surprisingly little work has been done in computer vision on weather related issues. An exception is the work of Cozman and Krotkov (1997) which uses the scattering models in McCartney (1975) to compute depth cues. Their algorithm assumes that all scene points used for depth estimation have the same intensity on a clear day. Since scene points can have their own reflectances and illuminations, this assumption is hard to satisfy in practice.

Research in image processing has been geared towards restoring contrast of images degraded by bad weather. Note that bad weather effects depend strongly on the depths of scene points. Hence, simple contrast enhancement techniques such as histogram equalization and contrast stretching do not suffice here. Oakley and Satherley (1998) use separately measured range data and describe an algorithm to restore contrast of atmospherically degraded images based on the principles of scattering. However, they approximate the distribution of radiances in the scene by a single gaussian with known variance. Kopeika (1998) and Yitzhaky et al. (1998) restore image contrast using weather predicted atmospheric modulation transfer function and an a priori estimate of the distance from which the scene was imaged.

The goal of our work is to lay the foundation for interpreting scenes from one or more images taken under bad weather conditions. We discuss various types of weather conditions and their formation processes. We summarize two models of atmospheric scattering—attenuation and airlight—that are most pertinent to us. Using these models, we develop algorithms that recover complete depth maps of scenes without requiring any prior information about the properties of the scene points or atmospheric conditions.

Next, we study the color effects of atmospheric scattering. A new model that describes the appearance of scene colors under bad weather is presented and verified for fog and haze. Based on this color model, we develop several geometric constraints on scene-color changes, caused by varying atmospheric conditions. Using these constraints, we present methods to recover structure as well as “clear day” scene colors from images taken under poor weather conditions. All of these methods only require changes in weather conditions and accurate measurement of scene irradiance, and not any prior information about the scene points or weather conditions.

2. Bad Weather: Particles in Space

Weather conditions differ mainly in the types and sizes of the particles involved and their concentrations in space. A great deal of effort has gone into measuring particle sizes and concentrations for a variety of conditions (see Table 1). Given the small size of air molecules, relative to the wavelength of visible light, scattering due to air is rather minimal. We will refer to the event of pure air scattering as a *clear* day (or night). Larger particles produce a variety of weather conditions which we will briefly describe below.

Haze. Haze is constituted of *aerosol* which is a dispersed system of small particles suspended in gas. Haze has a diverse set of sources including volcanic ashes, foliage exudation, combustion products, and sea salt (see Hidy (1972)). The particles produced by these sources respond quickly to changes in relative humidity and act as nuclei (centers) of small water droplets when the humidity is high. Haze particles are larger than air molecules but smaller than fog droplets. Haze tends to produce a distinctive gray or bluish hue and is certain to effect visibility.

Table 1. Weather conditions and associated particle types, sizes and concentrations (adapted from McCartney (1975)).

| Condition | Particle type | Radius (μm) | Concentration (cm^{-3}) |
|-----------|---------------|--------------------------|------------------------------------|
| Air | Molecule | 10^{-4} | 10^{19} |
| Haze | Aerosol | 10^{-2} –1 | 10^3 – 10 |
| Fog | Water droplet | 1–10 | 100–10 |
| Cloud | Water droplet | 1–10 | 300–10 |
| Rain | Water drop | 10^2 – 10^4 | 10^{-2} – 10^{-5} |

Fog. Fog evolves when the relative humidity of an air parcel reaches saturation. Then, some of the nuclei grow by condensation into water droplets. Hence, fog and certain types of haze have similar origins and an increase in humidity is sufficient to turn haze into fog. This transition is quite gradual and an intermediate state is referred to as *mist*. While perceptible haze extends to an altitude of several kilometers, fog is typically just a few hundred feet thick. A practical distinction between fog and haze lies in the greatly reduced visibility induced by the former. There are many types of fog (ex., radiation fog, advection fog, etc.) which differ from each other in their formation processes (Myers, 1968).

Cloud. A cloud differs from fog only in existing at higher altitudes rather than sitting at ground level. While most clouds are made of water droplets like fog, some are composed of long ice crystals and ice-coated dust grains. Details on the physics of clouds and precipitation can be found in Mason (1975). For now, clouds are of less relevance to us as we restrict ourselves to vision at ground level rather than high altitudes.

Rain and snow. The process by which cloud droplets turn to rain is a complex one (Mason, 1975). When viewed up close, rain causes random spatial and temporal variations in images and hence must be dealt with differently from the more static weather conditions mentioned above. Similar arguments apply to snow, where the flakes are rough and have more complex shapes and optical properties (Koenderink and Richards, 1992; Ohtake, 1970). Snow too, we will set aside for now.

3. Mechanisms of Atmospheric Scattering

The manner in which a particle scatters incident light depends on its material properties, shape and size. The exact form and intensity of the scattering pattern varies dramatically with particle size (Minnaert, 1954). As




Figure 1. A particle in the path of an incident light wave abstracts and reradiates incident energy. It therefore behaves like a point source of light. The exact scattering function is closely related to the ratio of particle size to wavelength of incident light. (Adapted from Minnaert (1954)).

seen in Fig. 1, a small particle (about $1/10 \lambda$, where λ is the wavelength of light) scatters almost equally in the forward (incidence) and backward directions, a medium size particle (about $1/4 \lambda$) scatters more in the forward direction, and a large particle (larger than λ) scatters almost entirely in the forward direction. Substantial theory has been developed to derive scattering functions and their relations to particle size distributions (Mie, 1908; Hulst, 1957; Chandrasekhar, 1960; Chu and Hogg, 1968; Rensch and Long, 1970; Nieto-Vesperinas and Dainty, 1990).

Figure 1 illustrates scattering by a single particle. Clearly, particles are accompanied in close proximity by numerous other particles. However, the average separation between atmospheric particles is several times the particle size. Hence, the particles can be viewed as *independent* scatterers whose scattered intensities do not interfere with each other. This does not imply that the incident light is scattered only by a single particle. *Multiple* scatterings take place and any given particle is exposed not only to the incident light but also light scattered by other particles. A simple analogy is the inter-reflections between scene points. In effect, multiple scattering causes the single scattering functions in Fig. 1 to get smoother and less directional.

Now, consider the simple illumination and detection geometry shown in Fig. 2. A unit volume of scattering medium with suspended particles is illuminated with spectral irradiance $E(\lambda)$ per cross section area. The radiant intensity $I(\theta, \lambda)$ of the unit volume in the direction θ of the observer is (see McCartney (1975)):

$$I(\theta, \lambda) = \beta(\theta, \lambda)E(\lambda), \tag{1}$$

where, $\beta(\theta, \lambda)$ is the *angular scattering coefficient*. The radiant intensity $I(\theta, \lambda)$ is the flux radiated per unit solid angle, per unit volume of the medium. The irradiance $E(\lambda)$ is, as always, the flux incident on the volume




Figure 2. A unit volume of randomly oriented suspended particles illuminated and observed.

per unit cross-section area. The total flux scattered (in all directions) by this volume is obtained by integrating over the entire sphere:

$$\phi(\lambda) = \beta(\lambda)E(\lambda), \tag{2}$$

where, $\beta(\lambda)$ is the *total scattering coefficient*. It represents the ability of the volume to scatter flux of a given wavelength in all directions. It is generally assumed that the coefficient $\beta(\lambda)$ is constant (homogeneous medium) over horizontal paths. To satisfy this constraint, we will restrict ourselves to the case where the observer is at (or close to) ground level and is interested not in the sky but other objects on (or close to) ground level. Also, we will assume that the atmosphere is more or less homogeneous in the scene of interest.

3.1. Attenuation

The first mechanism that is relevant to us is the attenuation of a beam of light as it travels through the atmosphere. This causes the radiance of a scene point to fall as its depth from the observer increases. Here, we will summarize the derivation of the attenuation model given in McCartney(1975). Consider a collimated beam of light incident on the atmospheric medium, as shown in Fig. 3. The beam is assumed to have unit cross-sectional area. Consider the beam passing through an infinitesimally small sheet (lamina) of thickness dx . The fractional change in irradiance at location x can be written as:

$$\frac{dE(x, \lambda)}{E(x, \lambda)} = -\beta(\lambda) dx. \tag{3}$$




Figure 3. Attenuation of a collimated beam of light by suspended particles. The attenuation can be derived by viewing the medium as a continuum of thin sheets.

By integrating both sides between the limits $x = 0$ and $x = d$ we get:

$$E(d, \lambda) = E_o(\lambda)e^{-\beta(\lambda)d}, \tag{4}$$

where, $E_o(\lambda)$ is the irradiance at the source ($x = 0$). This is Bouguer's exponential law of attenuation (Bouguer, 1729). At times, attenuation due to scattering is expressed in terms of *optical thickness*, $T = \beta(\lambda)d$. The utility of Bouguer's law is somewhat limited as it assumes a collimated source of incident energy. This is easily remedied by incorporating the inverse-square law for diverging beams from point sources:

$$E(d, \lambda) = \frac{I_o(\lambda)e^{-\beta(\lambda)d}}{d^2}, \tag{5}$$

where, $I_o(\lambda)$ is the radiant intensity of the point source. This is Allard's law (Allard, 1876). (See Hardy (1967) for an analysis of the applicability of the inverse square criterion for sources of various sizes.)

In deriving Allard's law, we have assumed that all scattered flux is removed from the incident energy. The fraction of energy that remains is called *direct transmission* and is given by expression (5). We have ignored the flux scattered in the forward direction (towards the observer) by each particle. Fortunately, this component is small in vision applications since the solid angles subtended by the source and the observer with respect to each other are small (see Middleton (1949)). In the remainder of the paper, we refer to the terms direct transmission model and attenuation model interchangeably.

Finally, in some situations such as heavy fog, the exponential law may not hold due to significant multiple scatterings of light by atmospheric particles. We will assume here that once light flux is scattered out of a column of atmosphere (seen by a pixel, say), it does not re-enter the *same* column (or only an insignificant amount does). Multiple scattering can also cause blurring in the image of a scene. In other words, the flux scattered out of an atmospheric column (visible to a pixel) enters another column (seen by a *neighboring* pixel). In this work, we do not model the blurring effects of multiple scattering.

3.2. Airlight

A second mechanism causes the atmosphere to behave like a source of light. This phenomenon is called airlight (Koschmieder, 1924) and it is caused by the scattering of environmental illumination by particles in the

atmosphere. The environmental illumination can have several sources, including, direct sunlight, diffuse skylight and light reflected by the ground. While attenuation causes scene radiance to decrease with pathlength, airlight increases with pathlength. It therefore causes the apparent brightness of a scene point to increase with depth. We now build upon McCartney's (1975) derivation of airlight as a function of pathlength.

Consider the illumination and observation geometry shown in Fig. 4. The environmental illumination along the observer's line of sight is assumed to be constant but unknown in direction, intensity and spectrum. In effect, the cone of solid angle $d\omega$ subtended by a single receptor at the observer's end, and truncated by a physical object at distance d , can be viewed as a source of airlight. The infinitesimal volume dV at distance x from the observer may be written as the product of the cross section area, $d\omega x^2$, and thickness dx :

$$dV = d\omega x^2 dx. \quad (6)$$

Irrespective of the exact type of environmental illumination incident upon dV , its intensity due to scattering in the direction of the observer is:

$$dI(x, \lambda) = dV k\beta(\lambda) = d\omega x^2 dx k\beta(\lambda), \quad (7)$$

where, $\beta(\lambda)$ is the total scattering coefficient and the proportionality constant k accounts for the exact nature of the illumination and the form of the scattering function.




Figure 4. The cone of atmosphere between an observer and an object scatters environmental illumination in the direction of the observer. It therefore acts like a source of light, called airlight, whose brightness increases with pathlength.

If we view element dV as a source with intensity $dI(x, \lambda)$, the irradiance it produces at the observer's end, after attenuation due to the medium, is given by (5):

$$dE(x, \lambda) = \frac{dI(x, \lambda)e^{-\beta(\lambda)x}}{x^2}. \quad (8)$$

We can find the radiance of dV from its irradiance as:

$$dL(x, \lambda) = \frac{dE(x, \lambda)}{d\omega} = \frac{dI(x, \lambda)e^{-\beta(\lambda)x}}{d\omega x^2}. \quad (9)$$

By substituting (7) we get $dL(x, \lambda) = k\beta(\lambda)e^{-\beta(\lambda)x} dx$. Now, the total radiance of the pathlength d from the observer to the object is found by integrating this expression between $x = 0$ and $x = d$:

$$L(d, \lambda) = k(1 - e^{-\beta(\lambda)d}). \quad (10)$$

If the object is at an infinite distance (at the horizon), the radiance of airlight is maximum and is found by setting $d = \infty$ to get $L_\infty(\lambda) = k$. Therefore, the radiance of airlight for any given pathlength d is:

$$L(d, \lambda) = L_\infty(\lambda)(1 - e^{-\beta(\lambda)d}). \quad (11)$$

As expected, the radiance of airlight for an object right in front of the observer ($d = 0$) equals zero. Of great significance to us is the fact that the above expression no longer includes the unknown factor k . Instead, we have the airlight radiance $L_\infty(\lambda)$ at the horizon, which is an observable.

3.3. Overcast Sky Illumination

Allard's attenuation model in (5) is in terms of the radiant intensity of a point source. This formulation does not take into account the sky illumination and its reflection by scene points. We make two simplifying assumptions regarding the illumination received by a scene point. Then, we reformulate the attenuation model in terms of sky illumination and the BRDF of scene points.

Usually, the sky is overcast under foggy conditions. So, we use the overcast sky model for environmental illumination (Gordon and Church, 1966; Moon and Spencer, 1942). We also assume that the irradiance at each scene point is dominated by the radiance of the sky, and that the irradiance due to other scene points is not significant. In Appendix A, we show that the

attenuated irradiance at the observer is given by,

$$E(d, \lambda) = g \frac{L_\infty(\lambda)\rho(\lambda)e^{-\beta(\lambda)d}}{d^2}. \quad (12)$$

where $L_\infty(\lambda)$ is the horizon radiance. $\rho(\lambda)$ represents the sky aperture (the cone of sky visible to a scene point), and the reflectance of the scene point in the direction of the viewer. g represents the optical settings of the camera (aperture, for instance). Note that we refer to (5) as the direct transmission model while dealing with images of light sources taken at night. However, while dealing with images of scenes taken during daylight, we refer to (12) as the direct transmission model.

4. Depths of Light Sources from Attenuation

Consider the image of an urban setting taken at *night* (see Fig. 5). Environmental illumination of the scene due to sunlight, skylight and reflected ground light are minimal and hence airlight can be safely ignored. The bright points in the image are mainly sources of light such as street lamps and windows of lit rooms. On a clear night, these sources are visible to a distant observer in their brightest and clearest forms. As haze or fog sets in, the radiant intensities of the sources diminish due to attenuation. Our goal here is to recover the relative depths of the sources in the scene from two images taken under different (unknown) atmospheric conditions.

Since environmental illumination is negligible at night, the image irradiance of a light source in the scene can be expressed using the attenuation model (5) as:

$$E(d, \lambda) = g \frac{I_o(\lambda)e^{-\beta(\lambda)d}}{d^2}, \quad (13)$$




Figure 5. The relative depths of sources of unknown intensities can be recovered from two images taken under different but unknown atmospheric conditions.

where, $I_o(\lambda)$ is the radiant intensity of the source, d is the distance between the source and the camera and the constant gain g accounts for the optical parameters (aperture, for instance) of the camera. It is important to note that $\beta(\lambda)$ is the total scattering coefficient and not the angular one. We are assuming here that the lines of sight are not too inclined and hence all lines of sight pass through the same atmospheric conditions. This removes all dependence on the exact form of the scattering function; the attenuation is determined by a single coefficient $\beta(\lambda)$ that is independent of viewing direction.

If the detector of the camera has spectral response $s(\lambda)$, the final image brightness recorded is determined as:

$$E' = \int s(\lambda)E(d, \lambda) d\lambda = \int g s(\lambda) \frac{I_o(\lambda)e^{-\beta(\lambda)d}}{d^2} d\lambda. \quad (14)$$

For the visible light spectrum, the relationship between the scattering coefficient β , and the wavelength λ , is given by the inverse power law (analogous to Rayleigh's law for small air particles):

$$\beta(\lambda) = \frac{\text{Constant}}{\lambda^\gamma}, \quad (15)$$

where $\gamma \in [0, 4]$. Fortunately, for fog and dense haze, $\gamma \approx 0$ (see Middleton (1952) and McCartney (1975)). In these cases, β does *not* change appreciably with wavelength. Furthermore, since the spectral bandwidth of the camera is rather limited (visible light range for a gray-scale camera, and even narrower spectral bands when the camera is color), we will assume the scattering coefficient $\beta(\lambda)$ to be constant over this bandwidth. Then, we have:

$$E' = g \frac{e^{-\beta d}}{d^2} \int s(\lambda)I(\lambda) d\lambda = g \frac{e^{-\beta d}}{d^2} I'. \quad (16)$$

Now consider two different weather conditions, say, mild and dense fog. Or, one of the conditions could be clear with $\beta = 0$. In either case we have two different attenuation coefficients, β_1 and β_2 . If we take the ratio of the two resulting image brightness values, we get:

$$R = \frac{E'_1}{E'_2} = e^{-(\beta_1 - \beta_2)d}. \quad (17)$$

Using the natural log, we obtain: $R' = \ln R = -(\beta_1 - \beta_2)d$. This quantity is independent of the sensor gain and the radiant intensity of the source. In fact, it is

nothing but the *difference in optical thicknesses* (DOT) of the source for two weather conditions. In the atmospheric optics literature, the term DOT is used as a quantitative measure of the “change” in weather conditions. Now, if we compute the DOTs of two different light sources in the scene (see Fig. 5) and take their ratio, we determine the relative depths of the two source locations:

$$\frac{R'_i}{R'_j} = \frac{d_i}{d_j} \quad (18)$$

Hence, the relative depths of all sources (with unknown radiant intensities) in the scene can be computed from two images taken under unknown but different haze or fog conditions. Since we may not entirely trust the DOT computed for any single source, the above calculation may be made more robust by using:

$$\frac{R'_i}{\sum_{j=0}^{j=N} R'_j} = \frac{d_i}{\sum_{j=0}^{j=N} d_j} \quad (19)$$

By setting the denominator on the right hand side to an arbitrary constant we have computed the depths of all sources in the scene up to a scale factor.

Figure 6 shows experimental results on the recovery of light sources from night images. This experiment and all subsequent ones are based on images acquired using a Nikon N90s SLR camera and a Nikon LS-2000 slide scanner. All images are linearized using the radiometric response curve of the imaging system that is computed off-line using a color chart. Figure 6(a) shows a clear day image of a scene with five lamps. This image is provided only to give the reader an idea of where the lamps are located in the scene. Figures 6(b) and (c) are clear night and foggy night images of the same scene. The above algorithm for depth estimation was used to recover the locations of all five light sources up to a scale factor. Figure 6(d) shows different perspectives of the recovered coordinates of the lamps in three-dimensional space. The poles and the ground plane are added only to aid visualization of the results.

5. Structure from Airlight

Under dense fog and close by objects or mild fog and distant objects, attenuation of object brightness is severe and airlight is the main cause of image irradiance. Also, in the case of dense haze around noon, airlight dominates. In such cases, airlight causes object brightness to increase with distance from the observer. Here,

we present a simple method for computing scene structure from a single airlight image. A different but related method for computing depth cues was proposed by Cozman and Krotkov (1997).

Let a scene point at depth d produce airlight radiance $L(d, \lambda)$. If our camera has a spectral response $s(\lambda)$, the final brightness value recorded for the scene point is:

$$E'(d) = \int g s(\lambda) L(d, \lambda) d\lambda, \quad (20)$$

where, g accounts for the constant of proportionality between scene radiance and image irradiance (Horn, 1986). Substituting the model for airlight given by (11) we get:

$$E'(d) = \int g s(\lambda) L_\infty(\lambda) (1 - e^{-\beta(\lambda)d}) d\lambda \quad (21)$$

where, $L_\infty(\lambda)$ is again the radiance of airlight at the horizon. As before, we will assume that the scattering coefficient $\beta(\lambda)$ is more or less constant over the spectral band of the camera. This allows us to write:

$$E'(d) = E_\infty (1 - e^{-\beta d}). \quad (22)$$

Let us define:

$$S = \frac{E_\infty - E'(d)}{E_\infty}. \quad (23)$$

By substituting (22) in the above expression and taking the natural logarithm, we get:

$$S' = \ln S = -\beta d. \quad (24)$$

Hence, the three-dimensional structure of the scene can be recovered up to a scale factor (the scattering coefficient β) from a single image. Clearly, at least a small part of the horizon must be visible to obtain E_∞ . If so, this part is easily identified as the brightest region of the image. If there is a strong (directional) sunlight component to the illumination, scattering would be greater in some directions and airlight could be dependent on viewing direction. This problem can be alleviated by using the horizon brightness E_∞ that lies closest to the scene point under consideration. Figure 7 shows the structure of an urban setting computed from a hazy image taken around noon, and the structure of a mountain range computed using a foggy image. Given that some of the objects are miles away




Figure 6. Relative depths from brightness decay of point sources at night. (a) A scene with five light sources (street lamps). This image is shown only to convey the relative locations of the sources to the reader. (b) An image of the scene taken on a clear night. (c) An image of the scene taken on a foggy night. The three-dimensional coordinates of the five sources were computed from images (b) and (c). (d) Rotated graphical illustrations used to demonstrate the computed lamp coordinates (small bright spheres). The lamp poles and the ground plane are added only to aid visualization.

from the camera, such scene structures are hard to compute using stereo or structure from motion. An interesting study of the visibility of distant mountains taking into account earth's curvature can be found in Porch (1975).

6. Dichromatic Atmospheric Scattering

Thus far, we have not taken into account the chromatic effects of atmospheric scattering. Furthermore, we have described attenuation and airlight separately.

However, in most situations the effects of both attenuation and airlight coexist. In the remainder of the paper, we discuss the chromatic effects of atmospheric scattering that include both attenuation and airlight, and hence develop a general framework for analyzing color images taken in bad weather. For this, we first present a new model that describes the appearance of scene colors in poor visibility conditions.

As we know, attenuation causes the radiance of the surface to decay as it travels to the observer. In addition, if the particle sizes are comparable to the wavelengths

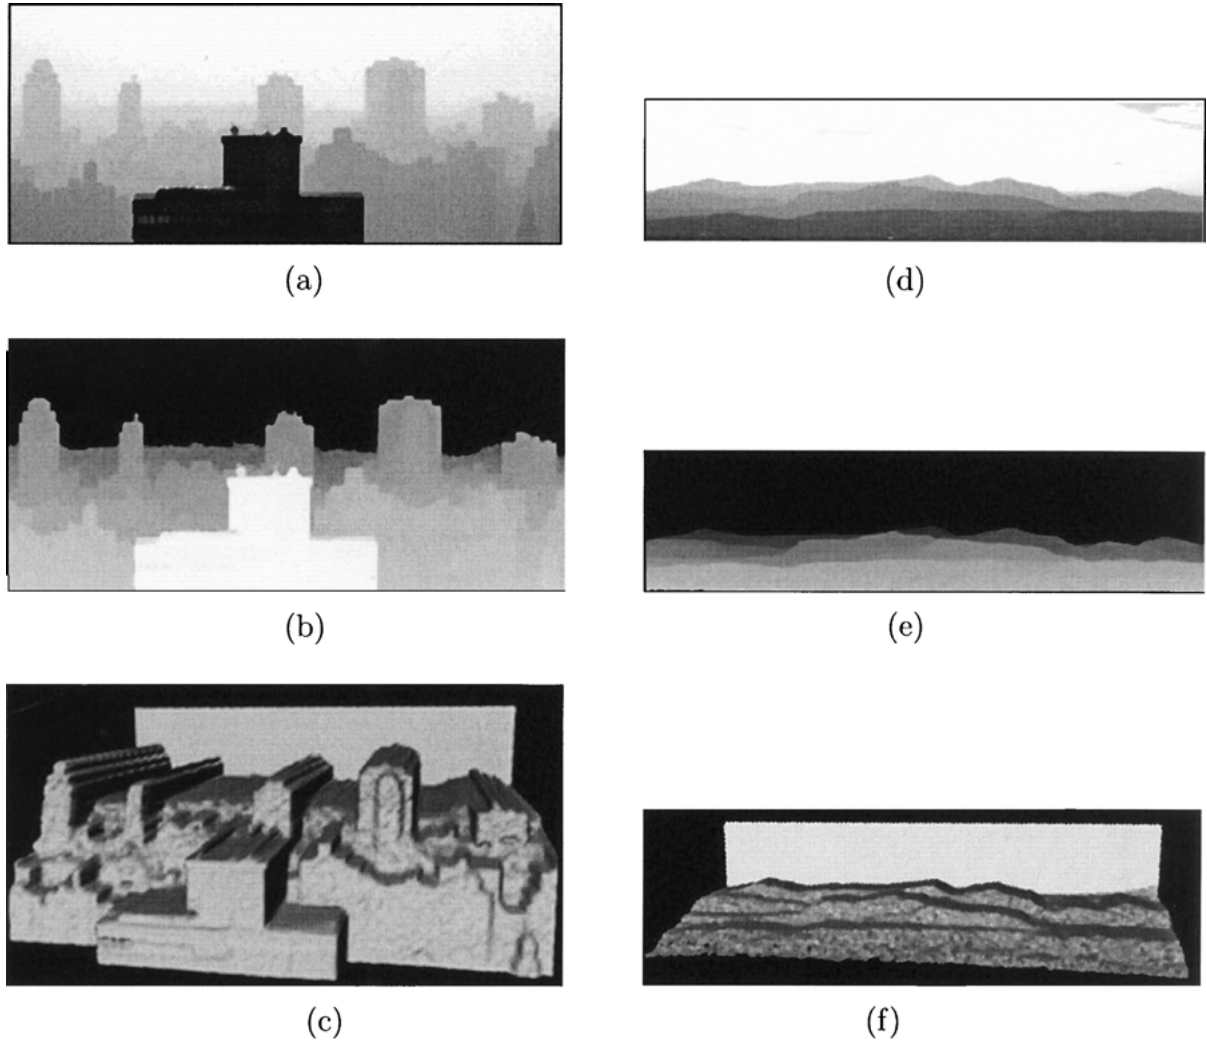


Figure 7. Structure from one image taken under dense fog/haze. (Left column) (a) Image of an urban scene taken under noon haze. (b) Depth map of the scene computed using the image in (a). (c) A three-dimensional rendering of the scene. (Right column) (d) Image of a mountain range taken under foggy conditions. (e) Depth map computed from the image in (d). (f) A three-dimensional rendering of the scene. Some of the objects in these scenes are several kilometers away from the camera.

of the reflected light, the spectral composition of the reflected light can be expected to vary as it passes through the medium. For fog and dense haze, these shifts in the spectral composition are minimal (see Middleton (1952) for details), and hence we may assume the hue of direct transmission to be *independent of the depth* of the reflecting surface. The hue of airlight depends on the particle size distribution and tends to be gray or light blue in the case of haze and fog. Therefore, the final spectral distribution $E(d, \lambda)$ received by the observer is a sum of the distributions $D(d, \lambda)$ of directly transmitted light and $A(d, \lambda)$ of airlight, which are de-

termined by the attenuation model (12) and the airlight model (11) respectively:

$$\begin{aligned} E(d, \lambda) &= D(d, \lambda) + A(d, \lambda), \\ D(d, \lambda) &= g \frac{e^{-\beta(\lambda)d}}{d^2} L_\infty(\lambda) \rho(\lambda), \\ A(d, \lambda) &= g(1 - e^{-\beta(\lambda)d}) L_\infty(\lambda). \end{aligned} \quad (25)$$

Here, $L_\infty(\lambda)$ is the radiance of the horizon ($d = \infty$), and g is a constant that accounts for the optical settings of the imaging system. $\rho(\lambda)$ represents the reflectance

properties and sky aperture of the scene point. We refer to the above expression as the *dichromatic atmospheric scattering model*. It is similar in its spirit to the dichromatic reflectance model (Shafer, 1985) that describes the spectral effects of diffuse and specular surface reflections. A fundamental difference here is that one of our chromatic components is due to surface and volume scattering (transmission of reflected light) while the other is due to pure volume scattering (airlight). If a chromatic filter with a spectral response $f(\lambda)$ is incorporated into the imaging system, image irradiance is obtained by multiplying (25) by $f(\lambda)$ and integrating over λ :

$$E^{(f)}(d) = D^{(f)}(d) + A^{(f)}(d). \quad (26)$$

In the case of a color image detector several such filters (say, red, green and blue) with different sensitivities are used to obtain a color measurement vector. The dichromatic model can then be written as:

$$\mathbf{E}(d) = \mathbf{D}(d) + \mathbf{A}(d) \quad (27)$$

where, $\mathbf{E} = [E^{(f_1)}, E^{(f_2)}, \dots, E^{(f_n)}]^T$. As we mentioned earlier (see (15)), for fog and haze, the dependence of the scattering coefficient $\beta(\lambda)$ on the wavelength (within the small bandwidth of the camera) of light tends to be rather small. Therefore, except in the case of certain types of metropolitan haze, we may assume the scattering coefficient to be constant with respect to wavelength ($\beta(\lambda) = \beta$). Then, expression (26) may be simplified as:

$$E^{(f)}(d) = p'(d)D^{(f)} + q'(d)A^{(f)}, \quad (28)$$

where:

$$\begin{aligned} D^{(f)} &= \int gf(\lambda)L_\infty(\lambda)\rho(\lambda)d\lambda, \\ A^{(f)} &= \int gf(\lambda)L_\infty(\lambda)d\lambda, \\ p'(d) &= \frac{e^{-\beta d}}{d^2}, \quad q'(d) = (1 - e^{-\beta d}). \end{aligned} \quad (29)$$

Here, $D^{(f)}$ is the image irradiance due to the scene point *without* atmospheric attenuation and $A^{(f)}$ is the image irradiance at the horizon in the presence of bad weather. We are assuming here that the clear and bad weather have illuminations with similar spectral distributions. Hence, the color measurement given by (27) can be rewritten as: $\mathbf{E}(d) = p'(d)\mathbf{D} + q'(d)\mathbf{A}$. Since the intensity of illumination (or magnitude of the illumination




Figure 8. Dichromatic atmospheric scattering model. The color \mathbf{E} of a scene point on a foggy or hazy day, is a linear combination of the direction $\hat{\mathbf{D}}$ of direct transmission color, and the direction $\hat{\mathbf{A}}$ of airlight color.

spectrum) at a scene point is expected to vary between clear and bad weather, it is more convenient to write:

$$\mathbf{E}(d) = m|L_\infty|p'(d)\hat{\mathbf{D}} + n|L_\infty|q'(d)\hat{\mathbf{A}} \quad (30)$$

where $\hat{\mathbf{D}}$ and $\hat{\mathbf{A}}$ are unit vectors and m and n are scalars. $|L_\infty|$ is the magnitude of the illumination spectrum. For convenience, the dichromatic model is re-written as:

$$\mathbf{E} = p\hat{\mathbf{D}} + q\hat{\mathbf{A}}, \quad (31)$$

where p is the magnitude of direct transmission, and q is the magnitude of airlight (see Fig. 8). From (30) we have,

$$p = \frac{E_\infty r e^{-\beta d}}{d^2}, \quad q = E_\infty (1 - e^{-\beta d}). \quad (32)$$

where $E_\infty = n|L_\infty|$, is termed as the sky intensity and $r = m/n$ is a function that depends on the properties of the scene point (reflectance and sky aperture). For our analysis, the exact nature of r is not important; it suffices to note that r does not depend on the weather condition β .¹ This simplified dichromatic scattering model will prove useful in the coming sections when we attempt to recover scene structure and remove weather effects from images.

It is easy to see that the simplified dichromatic model (31) is *linear* in color space. In other words, $\hat{\mathbf{D}}$, $\hat{\mathbf{A}}$ and \mathbf{E} lie on the same *dichromatic plane* in color space. As




Figure 9. For fog and haze, the transmittance ($e^{-\beta(\lambda)d}$) does not vary appreciably with wavelength within the visible spectrum. The plots were generated using the atmospheric transmission software MODTRAN 4.0, with a fixed viewing geometry (distance, d and viewing directions are fixed).

stated earlier, we impose the restriction that the hue of illumination under various weather conditions remains the same although its intensity can vary. It follows that the unit vectors $\hat{\mathbf{D}}$ and $\hat{\mathbf{A}}$ do not change due to different atmospheric conditions (say, mild fog and dense fog). Therefore, the colors of any scene point, observed under different atmospheric conditions, lie on a single dichromatic plane (see Fig. 10(a)).

We performed simulations using the atmospheric transmission software MODTRAN 4.0 (Acharya et al., 1999) to verify that the scattering coefficient does not vary with wavelength within the visible spectrum ($0.4\mu\text{--}0.7\mu$). Figure 9 shows plots of transmittance ($e^{-\beta(\lambda)d}$) for a particular viewing geometry in fog and haze respectively. The distance from the observer to the scene was fixed at $d = 0.2$ km and the viewing direction was fixed at 5 degrees off the ground plane. The plots show that the variation in β is very small within the visible spectrum.

Experiments with real scenes (shown in Fig. 17) were performed to verify this model under three different fog and haze conditions. The sky was overcast in all these conditions. The images used contained around half a million pixels. The dichromatic plane for each pixel was computed by fitting a plane to the colors of that pixel, observed under the three atmospheric conditions. The error of the plane-fit was computed in terms of the angle between the observed color vectors and the estimated plane. The average absolute error (in de-




Figure 10. (a) Dichromatic plane geometry and its verification. The observed color vectors E_i of a scene point under different (two in this case) foggy or hazy conditions lie on a plane called the dichromatic plane. (b) Experimental verification of the dichromatic model with two scenes imaged under three different foggy and hazy conditions, respectively. The error was computed as the mean angular deviation (in degrees) of the observed scene color vectors from the estimated dichromatic planes, over half a million pixels in the images.

grees) for all the pixels in each of the two scenes is shown in Fig. 10(b). The small error values indicate that the dichromatic model indeed works well for fog and haze.

7. Weather Removal and Structure using Chromatic Decomposition

Consider color images of a scene taken under clear weather and foggy or hazy weather. Assume that the clear day image is taken under environmental illumination with similar spectral characteristics as the bad weather image. If not, a white patch in the scene may be used to apply the needed color corrections. The sky in the bad weather image reveals the *direction* of the airlight color $\hat{\mathbf{A}}$. The *direction* of the color $\hat{\mathbf{D}}$ of each scene point is revealed by the clear weather image. Therefore, Eq. (31) can be used to decompose the bad weather color \mathbf{E} at each pixel into its two components and determine the airlight magnitude $q(d)$. The resulting airlight image is then used to compute a depth map as described in Section 5. Figure 11 shows experimental results obtained using the above decomposition method. Figure 12 demonstrates a simple form of weather removal by defogging windows of buildings.

In computing depth from the airlight component, we have assumed that the atmosphere itself is uniformly illuminated. Consider a pathlength that extends from a point on a building to an observer. Clearly, atmospheric points closer to the building see less of the sky due to occlusion by the building. This effect increases towards the foot of the building. Some of the errors in the computed structure can be attributed to this *illumination*




Figure 11. Structure from chromatic decomposition. (a) Clear day image of buildings. (b) Foggy day image of the same scene. (c) The direct transmission component (brightened) estimated by the chromatic decomposition algorithm. Black and gray points (windows) are discarded due to lack of color. (d) Depth map of the scene computed from the airlight component (depths of window areas are interpolated). (e) A three-dimensional rendering of the computed depth map.



Figure 12. Demonstration of fog removal. (a) A clear day image of a building taken under an overcast sky. The color directions (and not magnitudes) of scene points (non-window regions) are recorded as true colors or clear day colors. (b) A foggy day image of the same scene again captured under an overcast sky. Note: Even though both images in (a) and (b) were taken on overcast days (i.e., spectral composition of the daylight on both days are more or less identical), the horizon brightnesses (and/or camera exposure parameters) can vary. (c) The true colors recorded were used to decompose the foggy image into direct transmission and airlight components. The airlight component was subtracted from the window regions to demonstrate a simple form of weather removal.

occlusion effect (see Appendix B for a more detailed treatment).

Finally, there are certain limitations to this type of decomposition. First, we cannot decompose (31) if both the airlight and scene points have the same color. Also, this algorithm for chromatic decomposition is restrictive since it requires a clear day image of the scene. In the remainder of the paper, we develop more general constraints and algorithms to compute structure as well as recover “clear day” colors, without requiring a clear day image of the scene.

8. Computing the Direction of Airlight Color

The direction of airlight (fog or haze) color can be simply computed by averaging a patch of the sky on a foggy or hazy day (as was done in Section 7), or from scene points whose direct transmission color is *black*.² However, these methods necessitate either (a) the inclusion of a part of the sky (which is more prone to color saturation or clipping) in the image or (b) a clear day image of the scene with *sufficient* black points to yield a robust estimate of the direction of airlight color. Here, we present a method that does not require either the sky or a clear day image, to compute the direction of airlight color.

Figure 13 illustrates the dichromatic planes for two scene points P_i and P_j , with different direct transmission colors $\hat{\mathbf{D}}^{(i)}$ and $\hat{\mathbf{D}}^{(j)}$. The dichromatic planes Q_i




Figure 13. Intersection of two different dichromatic planes yields the direction $\hat{\mathbf{A}}$ of airlight color.

and Q_j are given by their normals,

$$\begin{aligned}\mathbf{N}_i &= \mathbf{E}_1^{(i)} \times \mathbf{E}_2^{(i)}, \\ \mathbf{N}_j &= \mathbf{E}_1^{(j)} \times \mathbf{E}_2^{(j)}.\end{aligned}\quad (33)$$

Since the direction $\hat{\mathbf{A}}$ of the airlight color is the same for the entire scene, it must lie on the dichromatic planes of all scene points. Hence, $\hat{\mathbf{A}}$ is given by the intersection of the two planes Q_i and Q_j ,

$$\hat{\mathbf{A}} = \frac{\mathbf{N}_i \times \mathbf{N}_j}{\|\mathbf{N}_i \times \mathbf{N}_j\|}.\quad (34)$$

In practice, scenes have several points with different colors. Therefore, we can compute a robust intersection of several dichromatic planes by minimizing the objective function

$$\epsilon = \sum_i (\mathbf{N}_i \cdot \hat{\mathbf{A}})^2.\quad (35)$$

Thus, we are able to compute the color of fog or haze using only the observed colors of the scene points under two atmospheric conditions, and not relying on a patch of the sky being visible in the image.

We verified the above method for the two scenes shown in Fig. 17. First, the direction of airlight color was computed using (35). Then, we compared it with the direction of the airlight color obtained by averaging an unsaturated patch of the sky. For the two scenes, the angular deviations were found to be 1.2° and 1.6° respectively. These small errors in the computed directions of airlight color indicate the robustness of the method.

9. Dichromatic Constraints for Iso-depth Scene Points

In this section, we derive a simple constraint for scene points that are at the same depth from the observer. This constraint can then be used to segment the scene based on depth, *without* knowing the actual reflectances of the scene points and their sky apertures. For this, we first prove the following lemma.

Lemma. *Ratios of the direct transmission magnitudes for points under two different weather conditions are equal, if and only if the scene points are at equal depths from the observer.*

Proof: Let β_1 and β_2 be two unknown weather conditions with horizon brightness values E_{∞_1} and E_{∞_2} .

Let P_i and P_j be two scene points at depths d_i and d_j , from the observer. Also, let $r^{(i)}$ and $r^{(j)}$ represent sky apertures and reflectances of these points.

From (32), the direct transmission magnitudes of P_i under β_1 and β_2 , can be written as

$$p_1^{(i)} = \frac{E_{\infty_1} r^{(i)} e^{-\beta_1 d_i}}{d_i^2}, \quad p_2^{(i)} = \frac{E_{\infty_2} r^{(i)} e^{-\beta_2 d_i}}{d_i^2}.$$

Similarly, the direct transmission magnitudes of P_j under β_1 and β_2 , are

$$p_1^{(j)} = \frac{E_{\infty_1} r^{(j)} e^{-\beta_1 d_j}}{d_j^2}, \quad p_2^{(j)} = \frac{E_{\infty_2} r^{(j)} e^{-\beta_2 d_j}}{d_j^2}.$$

Then, we immediately see that the relation:

$$\frac{p_2^{(i)}}{p_1^{(i)}} = \frac{p_2^{(j)}}{p_1^{(j)}} = \left(\frac{E_{\infty_2}}{E_{\infty_1}} \right) e^{-(\beta_2 - \beta_1)d}, \quad (36)$$

holds if and only if $d_i = d_j = d$. □

So, if we have the ratio of direct transmissions for each pixel in the image, we can group the scene points according to their depths from the observer. But how do we compute this ratio for any scene point without knowing the actual direct transmission magnitudes?

Consider the dichromatic plane geometry for a scene point P , as shown in Fig. 14. Here, we denote a vector by the line segment between its end points. Let p_1 and p_2 be the unknown direct transmission magnitudes of P




Figure 14. Geometric constraint for iso-depth scene points. The ratio p_2/p_1 of the direct transmissions for a scene point under two different atmospheric conditions is equal to the ratio $|E_2 A_t|/|E_i O|$ of the parallel sides. Shaded triangles are similar.

under β_1 and β_2 , respectively. Similarly let q_1 and q_2 be the unknown airlight magnitudes for P under β_1 and β_2 .

We define a magnitude $|O A_t|$ on the airlight vector such that $\overline{E_2 A_t} \parallel \overline{E_1 O}$. Also, since the *direction* of direct transmission color for a scene point does not vary due to different atmospheric conditions, $\overline{E_1 A_1} \parallel \overline{E_2 A_2}$. Here A_1 and A_2 correspond to the end points of the airlight magnitudes of P under β_1 and β_2 , as shown in Fig. 14. Thus, $\triangle E_1 O A_1 \sim \triangle E_2 A_t A_2$. This implies,

$$\frac{p_2}{p_1} = \frac{q_2 - |O A_t|}{q_1} = \frac{|E_2 A_t|}{|E_1 O|}. \quad (37)$$

Since the right hand side of (37) can be computed using the observed color vectors of the scene point P , we can compute the ratio (p_2/p_1) of direct transmission magnitudes for P under two atmospheric conditions. Therefore, from (36), we have a simple method to find points at the same depth, *without* having to know their reflectances and sky apertures.

Let us now consider the numerical stability of the direct transmission ratio (37). Under heavy fog/haze (or when the dynamic range of the sensor is low), the direct transmission magnitudes are low and their ratio could be unstable. In such cases, the ratio constraint can be supported by another constraint for depth segmentation we describe briefly. Consider the dichromatic planes of two different scene points as illustrated in Fig. 15. It




Figure 15. Another geometric constraint for two scene points to be equidistant from the observer. The dichromatic planes for the two points P_1 and P_2 are shown. Note that superscripts denote scene points while subscripts denote weather conditions. Shaded triangles are similar if and only if P_1 and P_2 are equidistant from the observer. Hence the iso-depth constraint is $\overline{E_1^{(1)} E_1^{(2)}} \parallel \overline{E_2^{(1)} E_2^{(2)}}$.

can be shown (using the geometric analysis in Fig. 14) that the shaded triangles are similar if and only if the two scene points are at equal depths from the observer. Therefore, the constraint for two scene points to be iso-depth is given in terms of observables,

$$\overline{E_1^{(1)} E_1^{(2)}} \parallel \overline{E_2^{(1)} E_2^{(2)}}. \quad (38)$$

Using the constraints in (36) and (38) a sequential labeling like algorithm can be used to efficiently segment scenes into regions of equal depth.

10. Scene Structure from Two Bad Weather Images

We extend the direct transmission ratio constraint given in (36) one step further and present a method to construct the complete structure of an arbitrary scene, from two images taken under poor weather conditions.

From (36), the ratio of direct transmissions of a scene point P under two atmospheric conditions, is given by

$$\frac{p_2}{p_1} = \frac{E_{\infty_2}}{E_{\infty_1}} e^{-(\beta_2 - \beta_1)d}. \quad (39)$$

Note that we have already computed the left hand side of the above equation using (37). Taking natural logarithms on both sides, we get

$$(\beta_2 - \beta_1)d = \ln\left(\frac{E_{\infty_2}}{E_{\infty_1}}\right) - \ln\left(\frac{p_2}{p_1}\right). \quad (40)$$

So, if we know the horizon brightness values, E_{∞_1} and E_{∞_2} , then we can compute the scaled depth $(\beta_2 - \beta_1)d$ at P . As before, $(\beta_2 - \beta_1)d$ is just the difference in optical thicknesses (DOT) for the pathlength d , under the two weather conditions.

10.1. Estimation of E_{∞_1} and E_{∞_2}

The expression for scaled depth give in (40), includes the horizon brightness values, E_{∞_1} and E_{∞_2} . These two terms are *observables* only if some part of the sky is visible in the image. However, the brightness values within the region of the image corresponding to the sky, cannot be trusted since they are prone to intensity saturation and color clipping. Here, we estimate E_{∞_1}

and E_{∞_2} using only points in the “non-sky” region of the scene.

Let q_1 and q_2 denote the magnitudes of airlight for a scene point P under atmospheric conditions β_1 and β_2 . Using (32), we have

$$q_1 = E_{\infty_1}(1 - e^{-\beta_1 d}), \quad q_2 = E_{\infty_2}(1 - e^{-\beta_2 d}). \quad (41)$$

Therefore,

$$\frac{E_{\infty_2} - q_2}{E_{\infty_1} - q_1} = \frac{E_{\infty_2}}{E_{\infty_1}} e^{-(\beta_2 - \beta_1)d}. \quad (42)$$

Substituting (39), we can rewrite the previous equation as

$$\left(\frac{p_2}{p_1}\right) = \frac{q_2 - c}{q_1}, \quad \text{where, } c = E_{\infty_2} - \left(\frac{p_2}{p_1}\right)E_{\infty_1}. \quad (43)$$

Comparing (43) and (37), we get $c = |OA_i|$ (see Fig. 14). hence, the expression for c in (43) represents a straight line equation in the unknown parameters, E_{∞_1} and E_{∞_2} . Now consider several pairs of $\{c^{(i)}, (p_2^{(i)}/p_1^{(i)})\}$ corresponding to scene points P_i , at different depths. Then, the estimation of E_{∞_1} and E_{∞_2} is reduced to a *line fitting* problem. Quite simply, we have shown that the horizon brightnesses under different weather conditions can be computed using *only* non-sky scene points.

Since both the terms on the right hand side of (40) can be computed for every scene point, we have a simple algorithm for computing the scaled depth at each scene point, and hence the complete scene structure, from two images taken under different atmospheric conditions.

10.2. Experimental Results

We now present results showing scene structure recovered from both synthetic and real images. The synthetic scene we used is shown on the left side of Fig. 16(a) as a 200×200 pixel image with 16 color patches. The colors in this image represent the direct transmission or “clear day” colors of the scene. We assigned a random depth value to each color patch. The rotated 3D structure of the scene is shown on the right side of Fig. 16(a). Then, two different levels of fog ($\beta_1/\beta_2 = 0.67$) were added to the synthetic scene according to the dichromatic model. To test robustness, we added noise to the foggy images. The noise was randomly selected from a uniformly distributed color cube of dimension 10. The




Figure 16. Experiments with a synthetic scene. (a) On the left, a 200×200 pixel image representing a synthetic scene with 16 color patches, and on the right, its rotated 3D structure. (b) Two levels of fog ($\beta_1/\beta_2 = 0.67$) are added to the synthetic image according to the dichromatic model. To test robustness, noise is added by random selection from a uniformly distributed color cube of dimension 10. (c) The recovered structure (3×3 median filtered).

resulting two foggy (and noisy) images are shown in Fig. 16(b). The structure shown in 16(c) is recovered from the two foggy images using the technique we described above.

Simulations were repeated for the scene in Fig. 16(a) for two relative scattering coefficient values (β_1/β_2), and three different noise levels. Once again, the noise was randomly selected from a uniformly distributed color cube of dimension η . Table 2 shows results of simulations for two parameter sets $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.5, 100, 255\}$ and $\{0.67, 200, 400\}$. The computed




Figure 17. Structure from two bad weather images. (a) A scene imaged under two different foggy conditions. (b) Depth map computed from images in (a). (c) Another scene imaged under two different hazy conditions. (d) Depth map computed from images in (c). All these images were captured under overcast sky conditions.

values for E_{∞_1} , E_{∞_2} , and the percentage RMS error in the recovered scaled depths, computed over all 200×200 pixels are given. These results show that our method for recovering structure is robust for reasonable amounts of noise.

Experiments with two real scenes under foggy and hazy conditions are shown in Fig. 17. The first of the two scenes was imaged under two foggy conditions, and is shown in 17(a). The second scene was imaged

Table 2. Simulations were repeated for the scene in Fig. 16(a), for two sets of parameter values, and three different noise levels. Noise was randomly selected from a uniformly distributed color cube of dimension η .

| Noise (η) | 0 | 5 | 10 | 15 |
|--|-----|-------|-------|-------|
| Actual values $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.5, 100, 255\}$ | | | | |
| Estimated E_{∞_1} | 100 | 108.7 | 109.2 | 119.0 |
| Estimated E_{∞_2} | 255 | 262.7 | 263.6 | 274.0 |
| Depth error (%) | 0.0 | 7.14 | 11.7 | 15.3 |
| Actual values $\{\beta_1/\beta_2, E_{\infty_1}, E_{\infty_2}\} = \{0.67, 200, 400\}$ | | | | |
| Estimated E_{∞_1} | 200 | 204.3 | 223.7 | 249.5 |
| Estimated E_{∞_2} | 400 | 403.8 | 417.5 | 444.2 |
| Depth error (%) | 0.0 | 12.3 | 15.3 | 17.8 |

under two hazy conditions as shown in 17(c). Figure 17(b) and (d) shows the corresponding recovered depth maps.

11. Clear Day Scene Colors

As we stated in the beginning of the paper, most outdoor vision applications perform well only under clear weather. Any discernible amount of scattering due to fog or haze in the atmosphere, hinders a clear view of the scene. Earlier we presented a simple form of weather removal that requires a clear day image of the scene (see Figs. 11 and 12 in Section 7). In this section, we compute the scene colors as they would appear on a clear but overcast day from two bad weather images. More precisely, we compute the direct transmission colors of the entire scene using minimal a priori scene information. For this, we first show that, given additional scene information (airlight or direct transmission vector) at a *single* point in the scene, we can compute the clear day colors of the *entire* scene from two bad weather images.

Consider the dichromatic model given in (31). The color of a scene point P_i under weather condition β is,

$$\mathbf{E}^{(i)} = p^{(i)}\hat{\mathbf{D}}^{(i)} + q^{(i)}\hat{\mathbf{A}}, \quad (44)$$

where $p^{(i)}$ is the direct transmission magnitude, and $q^{(i)}$ is the airlight magnitude of P_i . Suppose that the direction $\hat{\mathbf{D}}^{(i)}$ of direct transmission color for a *single* point P_i is given. Besides, the direction $\hat{\mathbf{A}}$ of airlight color for the entire scene can be estimated using (35). Therefore, the coefficients $p^{(i)}$ and $q^{(i)}$ can be computed using (44). Furthermore, the optical thickness βd_i of P_i can be computed from (32).

Since we have already shown how to compute the scaled depth of every scene point (see (40)), the relative depth d_j/d_i of any other scene point P_j with respect to P_i can be computed using the ratio of scaled depths. Hence, the optical thickness and airlight for the scene point P_j , under the same atmospheric condition are given by

$$\begin{aligned} \beta d_j &= \beta d_i (d_j/d_i), \\ q^{(j)} &= E_{\infty} (1 - e^{-\beta d_j}). \end{aligned} \quad (45)$$

Finally, the direct transmission color vector of P_j can be computed as

$$p^{(j)}\hat{\mathbf{D}}^{(j)} = \mathbf{E}^{(j)} - q^{(j)}\hat{\mathbf{A}}. \quad (46)$$

Thus, given a *single* measurement (in this case, the direction of direct transmission color of a single scene point), we have shown that the direct transmission and airlight color vectors of any other point, and hence the entire scene can be computed. But how do we specify the clear day color of any scene point without actually capturing the clear day image?

For this, we assume that there exists at least one scene point whose direct transmission color \mathbf{D} lies on




Figure 18. The observed color \mathbf{E} of a scene point, its airlight direction $\hat{\mathbf{A}}$ and clear day color direction $\hat{\mathbf{D}}$ are shown in the R-G-B color cube. \tilde{q} is the distance from \mathbf{E} to a surface of the cube along negative $\hat{\mathbf{A}}$. For scene points whose clear day colors do not lie on the cube surface, \tilde{q} is greater than the true airlight magnitude q .

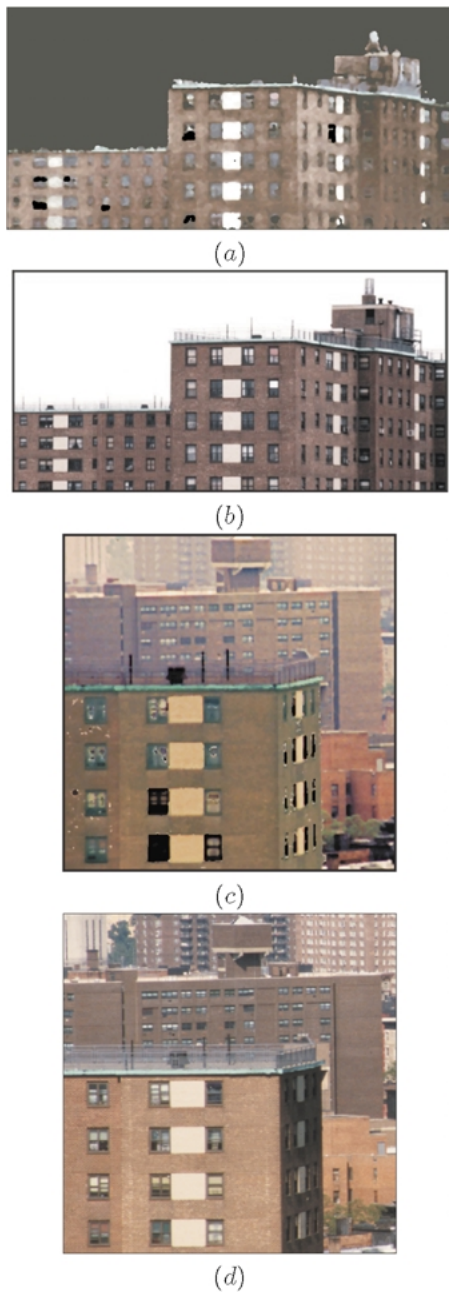


Figure 19. [(a) and (c)] Clear day scene colors recovered from the two foggy and hazy images shown in Fig. 17(a) and (c) respectively. The colors in some of the dark window interiors are dominated by airlight and thus their clear day colors are computed to be black. The images are median filtered to reduce noise and brightened for display purposes. [(b) and (d)] Actual clear day images of the scenes are shown for qualitative comparison. Note: The clear day images on the right and the bad weather images (Fig. 17) were captured on different days. Some differences between actual and recovered clear day colors are due to the different spectral distributions of illumination in the scene, during image acquisition.

the surface of the color cube (including origin or black) and we wish to identify such point(s) in the scene automatically. Consider the R-G-B color cube in Fig. 18. If the clear day color of a scene point lies *on* the surface

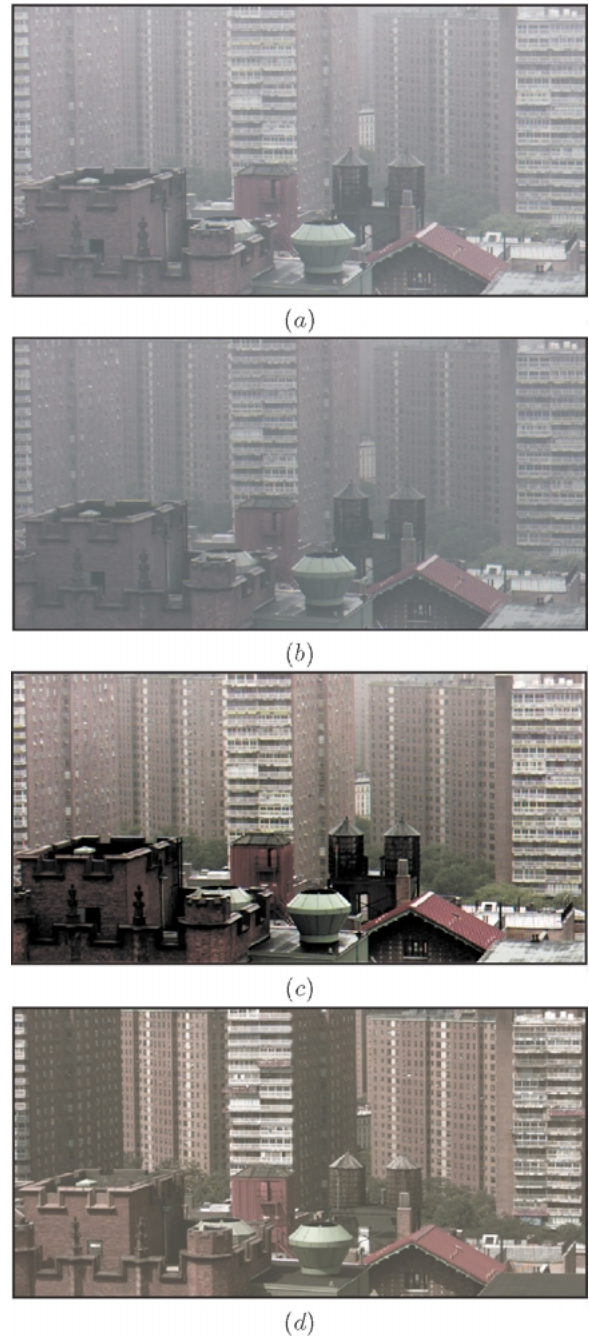



Figure 20. (a) and (b) Foggy images of a scene under an overcast sky. (c) Defogged image. (d) Actual clear day image taken under a partly cloudy sky.



(a)

| Relative Depth | Ground Truth | Computed Value |
|----------------|--------------|----------------|
| d_3 / d_1 | 9.55 | 10.457 |
| d_3 / d_2 | 1.3687 | 1.354 |
| d_3 / d_4 | 1.489 | 1.61 |
| d_3 / d_5 | 1.083 | 1.124 |

(b)

Figure 21. (a) Depth map computed using images in Fig. 20(a) and (b). Depth map is brightened for display purposes. (b) Table comparing the computed relative depths with ground truth relative depths of 5 different regions, $d_1 - d_4$, in the scene. The relative depths are averaged over small neighborhoods. Note that scaled depth can be computed only approximately due to the illumination occlusion problem (see Appendix B for more details). The depths in some window interiors are not reliable since they changed during acquisition of images over time.

of the color cube, then the computed \tilde{q} is equal to the airlight magnitude q of that point. However, if it lies *within* the color cube, then clearly $\tilde{q} > q$. For each point P_i , we compute $\tilde{q}^{(i)}$ and optical thickness $\tilde{\beta}_1 d_i$. Note that $\tilde{\beta}_1 d_i$ may or may not be the correct optical thickness. We normalize the optical thicknesses of the scene points by their scaled depths (DOTs) to get

$$\tilde{\alpha}_i = \frac{\tilde{\beta}_1 d_i}{(\beta_2 - \beta_1) d_i}. \quad (47)$$

For scene points that do not lie on the color cube surface, $\tilde{\alpha}_i$ is greater than what it should be. Since we have assumed that there exists at least one scene point whose clear day color is on the surface of the cube, it must be the point that has the minimum $\tilde{\alpha}_i$. So, $\tilde{q}^{(i)}$ of that point is its true airlight. Hence, from (45), the airlights and direct transmission colors of the entire scene can be computed without using a clear day image. For robustness, we use k least $\tilde{\alpha}_i$'s. We call this the *Color Cube Boundary Algorithm*.

Figure 19 illustrates experiments with real scenes. Usually in urban scenes, window interiors have very little color of their own. Their intensities are solely due to airlight and not due to direct transmission. In other words, their direct transmission color is black (the origin of the color cube). We detected such points in the scene using the above technique and recovered the clear day colors of foggy and hazy scenes. A second result is shown in Figs. 20 and 21.

12. Summary

Research in atmospheric optics has been around for over two centuries. The physical processes that govern the effects of atmospheric scattering on scene appearance are well established. This article is just an initial attempt at understanding and exploiting the manifestations of weather in order to interpret, recover and render scenes under various atmospheric conditions. We summarized existing models in atmospheric optics and proposed new ones, keeping in mind the constraints faced by most vision applications. We presented several simple algorithms for recovering scene structure from one or two bad weather images and demonstrated that bad weather can be put to good use. Using scene structure, algorithms to remove weather effects were developed. We intend to use these results as building blocks for developing more advanced weather-tolerant vision techniques. Potential applications of this work are in outdoor surveillance, navigation, underwater explorations and image based rendering.

Appendix A: Direct Transmission Under Overcast Skies

We present an analysis of the effect of sky illumination and its reflection by a scene point, on the direct transmission from the scene point. For this, we make two simplifying assumptions on the illumination received by scene points. Usually, the sky is overcast under foggy conditions. So we use the overcast sky model




Figure 22. The illumination geometry of a scene point P with surface normal $\hat{\mathbf{n}}$. The irradiance of P is due to the airlight radiance of its sky aperture Ω .

(Gordon and Church, 1966; IRIA, 1978) for environmental illumination. We also assume that the irradiance of each scene point is dominated by the radiance of the sky, and that the irradiance due to other scene points is not significant. See Langer and Zucker's work (1994) for a related analysis.

Consider the illumination geometry shown in Fig. 22. Let P be a point on a surface and $\hat{\mathbf{n}}$ be its normal. We define the *sky aperture* Ω of point P , as the cone of sky visible from P . Consider an infinitesimal patch of the sky, of size $\delta\theta$ in polar angle and $\delta\phi$ in azimuth as shown in Fig. 22. Let this patch subtend a solid angle $\delta\omega$ at P . For overcast skies, Moon (Moon and Spencer, 1942) and Gordon (Gordon and Church, 1966) have shown that the radiance of the infinitesimal cone $\delta\Omega$, in the direction (θ, ϕ) is given by $L(\theta, \phi) = L_\infty(\lambda)(1 + 2 \cos \theta)\delta\omega$, where $\delta\omega = \sin \theta \delta\theta \delta\phi$. Hence, the irradiance at P due to the entire aperture Ω , is given by

$$E(\lambda) = \iint_{\Omega} L_\infty(\lambda)(1 + 2 \cos \theta) \cos \theta \sin \theta d\theta d\phi, \quad (48)$$

where $\cos \theta$ accounts for foreshortening (Horn, 1986). If R is the *BRDF* of P , then the radiance from P toward the observer can be written as

$$L_o(\lambda) = \iint_{\Omega} L_\infty(\lambda)f(\theta)R(\theta, \phi, \lambda) d\theta d\phi, \quad (49)$$

where $f(\theta) = (1 + 2 \cos \theta) \cos \theta \sin \theta$. Let σ be the projection of a unit patch around P , on a plane perpendicular to the viewing direction. Then, the radiant intensity of P is given by $I_o(\lambda) = \sigma L_o(\lambda)$. Since $L_\infty(\lambda)$ is a constant with respect to θ and ϕ , we can factor it

out of the integral and write concisely as

$$I_o(\lambda) = L_\infty(\lambda)\rho(\lambda), \quad (50)$$

where

$$\rho(\lambda) = \sigma \iint_{\Omega} f(\theta)R(\theta, \phi, \lambda) d\theta d\phi. \quad (51)$$

The term $\rho(\lambda)$ represents the sky aperture and the reflectance in the direction of the viewer. Substituting for $I_o(\lambda)$ in the direct transmission model in (5), we obtain

$$E(d, \lambda) = g \frac{L_\infty(\lambda)\rho(\lambda)e^{-\beta(\lambda)d}}{d^2}, \quad (52)$$

where g represents the optical setting of the camera (exposure, for instance). We have thus formulated the direct transmission model in terms of overcast sky illumination and the reflectance of the scene points.

Appendix B: Illumination Occlusion Problem

In deriving the expression for the radiance due to airlight in Section 3.2, we assumed that the atmosphere is illuminated uniformly regardless of the type of illumination. This is not always true since not all points in the atmosphere “see” the same solid angle of the sky. In fact, the scene itself occludes part of the sky hemisphere visible to a point in the atmosphere. For explanation purposes, consider a scene with a single building. The solid angle subtended at any point in the atmosphere by the sky is called its sky aperture. As seen in Fig. 23, this solid angle decreases as the distance increases from the observer for any given pathlength. Similarly, the solid angle is smaller for points near the bottom of the building.

We now present a simplified analysis of this effect. We assume that the atmosphere is illuminated mainly by overcast skylight (ground light is ignored here). Then, the irradiance received by any point in the atmosphere is given by (see Eq. (48)),

$$\begin{aligned} E &= E^{(\text{hemisphere})} - E^{(\text{occluded})}, \\ E^{(\text{occluded})} &= \int_{-\phi}^{\phi} \int_0^{\theta} L_\infty(1 + 2 \cos \theta) \\ &\quad \times \cos \theta \sin \theta d\theta d\phi, \\ E^{(\text{hemisphere})} &= \int_{-\pi}^{\pi} \int_0^{\pi/2} L_\infty(1 + 2 \cos \theta) \\ &\quad \times \cos \theta \sin \theta d\theta d\phi, \end{aligned} \quad (53)$$




Figure 23. The scene occludes the sky aperture of points in the atmosphere. As a result points in the atmosphere are not uniformly illuminated by the sky.

where $E^{(hemisphere)}$ is the irradiance the point would receive from the entire sky hemisphere (as if there were no occlusions). $E^{occluded}$ is the irradiance the point would have received from the occluded part. θ and ϕ denote the polar and azimuth of the occluded region. The above equation simplifies to

$$E = L_{\infty} \frac{7\pi - 7\phi \cos^2 \theta (3 + 4 \cos \theta)}{3}. \quad (54)$$

To correct for the radiance of airlight in Section 3.2, we multiply by the fraction of irradiance received by each point and rewrite the airlight radiance (10) of a pathlength d as

$$\begin{aligned} L(d, \lambda) &= k(1 - e^{-\beta(\lambda)d}) - \int_0^d k \left(\frac{\phi \cos^2 \theta (3 + 4 \cos \theta)}{\pi} \right) \\ &\times \beta(\lambda) e^{-\beta(\lambda)x} dx. \end{aligned} \quad (55)$$

Note here that both θ and ϕ depend on the depth from the observer x (see Fig. 23). In other words, the integral in the previous equation depends on the exact extent of occlusion by the scene. In our experiments, we have assumed uniform illumination of the atmosphere and thus some of the errors in the depth maps can be attributed to this effect.

Acknowledgments

This work was supported in parts by a DARPA/ONR HumanID Contract (N00014-00-1-0916), an NSF Award (IIS-99-87979), and a DARPA/ONR MURI Grant (N00014-95-1-0601). The authors thank Jan Koenderink of Utrecht University for pointers to early work on atmospheric optics. The authors also thank Yoav Schechner for the discussions on this topic that helped improve the paper. Some of the results presented in this paper have appeared in the proceedings of the IEEE International Conference on Computer Vision'1999 (Nayar and Narasimhan, 1999), IEEE Conference on Computer Vision and Pattern Recognition'2000 (Narasimhan and Nayar, 2000) and SPIE Conference on Human Vision and Electronic Imaging'2001 (Narasimhan and Nayar, 2001).

Notes

1. We do not handle situations where wet materials may appear darker than dry materials.
2. Sky and black points take on the color of airlight on a bad weather day.

References

- Acharya, P.K., Berk, A., Anderson, G.P., Larsen, N.F., Tsay, S.C., and Stamnes, K.H. 1999. Modtran4: Multiple scattering and BRDF upgrades to modtran. In *SPIE Proc. Optical Spectroscopic Techniques and Instrumentation for Atmospheric and Space Research III*, p. 3756.
- Allard, E. 1876. *Memoire sur l'intensite' et la portee des phares*. Dunod: Paris.
- Bouguer, P. 1729. *Traite' d'optique sur la gradation de la lumiere*.
- Chandrasekhar, S. 1960. *Radiative Transfer*. Dover Publications: New York.
- Chu, T.S. and Hogg, D.C. 1968. Effects of precipitation on propagation at 0.63, 3.5 and 10.6 microns. *The Bell System Technical Journal*.
- Cozman, F. and Krotkov, E. 1997. Depth from scattering. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, vol. 31, pp. 801–806.
- Gordon, J. and Church, P. 1966. Overcast sky luminances and directional luminous reflectances of objects and backgrounds under overcast skies. *Applied Optics*, 5:919.
- Hardy, A.C. 1967. How large is a point source? *Journal of Optical Society of America*, 57(1).
- Henderson, S.T. 1977. *Daylight and its Spectrum*. Wiley: New York.
- Hidy, G.M. 1972. *Aerosols and Atmospheric Chemistry*. Academic Press: New York.
- Horn, B.K.P. 1986. *Robot Vision*. The MIT Press: Cambridge, MA.
- IRIA. 1978. *The Infrared Handbook*. Infrared Information and Analysis Center, Environmental Research Institute of Michigan.

- Koenderink, J.J. and Richards, W.A. 1992. Why is snow so bright? *Journal of Optical Society of America*, 9(5):643–648.
- Kopeika, N.S. 1998. *A System Engineering Approach to Imaging*. SPIE Press.
- Koschmieder, H. 1924. Theorie der horizontalen sichtweite. *Beitr. Phys. Freien Atm.*, 12:33–53, 171–181.
- Langer, M.S. and Zucker, S.W. 1994. Shape from shading on a cloudy day. *JOSA-A*, 11(2):467–478.
- Mason, B.J. 1975. *Clouds, Rain, and Rainmaking*. Cambridge University Press: Cambridge.
- McCartney, E.J. 1975. *Optics of the Atmosphere: Scattering by Molecules and Particles*. John Wiley and Sons: New York.
- Middleton, W.E.K. 1949. The effect of the angular aperture of a telephotometer on the telephotometry of collimated and non-collimated beams. *Journal of Optical Society of America*, 39:576–581.
- Middleton, W.E.K. 1952. *Vision Through the Atmosphere*. University of Toronto Press.
- Mie, G. 1908. A contribution to the optics of turbid media, especially colloidal metallic suspensions. *Ann. of Physics*, 25(4):377–445.
- Minnaert, M. 1954. *The Nature of Light and Color in the Open Air*. Dover: New York.
- Moon, P. and Spencer, D.E. 1942. Illumination from a non-uniform sky. *Illum Eng.*, 37:707–726.
- Myers, J.N. 1968. Fog. *Scientific American*, pp. 75–82.
- Narasimhan, S.G. and Nayar, S.K. 2000. Chromatic framework for vision in bad weather. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Narasimhan, S.G. and Nayar, S.K. 2001. Vision and the weather. In *Proceedings of SPIE Conference on Human Vision and Electronic Imaging VI*, p. 4299.
- Nayar, S.K. and Narasimhan, S.G. 1999. Vision in bad weather. In *Proceedings of the 7th International Conference on Computer Vision*.
- Nieto-Vesperinas, M. and Dainty, J.C. 1990. *Scattering in Volumes and Surfaces*. North-Holland: New York.
- Oakley, J.P. and Satherley, B.L. 1998. Improving image quality in poor visibility conditions using a physical model for degradation. *IEEE Trans. on Image Processing*, 7.
- Ohtake, T. 1970. Factors affecting the size distribution of raindrops and snowflakes. *Journal of Atmospheric Science*, 27:804–813.
- Porch, W.M. 1975. Visibility of distant mountains as a measure of background aerosol pollution. *Applied Optics*, 14.
- Rensch, D.B. and Long, R.K. 1970. Comparative studies of extinction and backscattering by aerosols, fog, and rain at 10.6 and 0.63 microns. *Applied Optics*, 9(7).
- Shafer, S. 1985. Using color to separate reflection components. *Color Research and Applications*, pp. 210–218.
- Van De Hulst. 1957. *Light Scattering by Small Particles*. John Wiley and Sons: New York.
- Yitzhaky, Y., Dror, I., and Kopeika, N.S. 1998. Restoration of atmospherically blurred images according to weather-predicted atmospheric modulation transfer function. *Optical Engineering*, 36.