

Vision in Bad Weather *

Shree K. Nayar and Srinivasa G. Narasimhan

Department of Computer Science, Columbia University
New York, New York 10027

Email: {nayar, srinivas}@cs.columbia.edu

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. Next, we 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. Based on this observation, we develop models and methods for recovering pertinent scene properties, such as three-dimensional structure, from images taken under poor weather conditions.

1 Vision and the Atmosphere

Virtually all work in 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]). 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 will be the works of McCartney [McCartney, 1976] and Middleton [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 [Cozman and Krotkov, 1997] which uses the scattering models in [McCartney, 1976] 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.

In this paper, we develop algorithms that recover complete depth maps of scenes without making assumptions about the properties of the scene points or the atmospheric conditions. How do such scene recovery methods compare with existing ones? Unlike binocular stereo, they do not suffer from the problems of correspondence and discontinuities. Nor do they require tracking of image features as in structure from motion. Furthermore, they are particularly useful for scenes with distant objects (even miles away) which pose problems for stereo and motion. The techniques we present here only require changes in weather conditions and accurate measurement of image irradiance.

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 con-

*This work was supported by the DARPA/ONR MURI Program under Grant N00014-95-1-0601 and the David and Lucile Packard Foundation. The authors thank Jan Koenderink of Utrecht University for pointers to early work on atmospheric optics.

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	$10^2 - 10^4$	$10^{-2} - 10^{-5}$

Table 1: Weather conditions and associated particle types, sizes and concentrations (adapted from McCartney[1976]).

centrations 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 exudations, 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 hue and is certain to effect visibility.

Fog: Fog evolves when the relative humidity of an air parcel approaches a saturation level. Then, some of the nuclei grow by condensation into water droplets. Hence, fog and 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 miles, 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 which differ from each other in their formation processes [Myers, 1968].

Cloud: A cloud differs from fog only in existing at high altitudes (troposphere) 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 stable weather conditions mentioned above. Similar arguments apply to snow, which, at a simple level may be

viewed as frozen rain where the drops are solid, rougher 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 seen in Figure 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 [Mie, 1908] (see [Nieto-Vesperinas and Dainty, 1990] for more recent advances).




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]).

Figure 1 illustrates scattering by a single particle. Clearly, particles are accompanied in close proximity by numerous other particles. However, the average separation between weather 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. In effect, this causes the single scattering functions in Figure 1 to get smoother and less directional.

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

$$I(\theta, \lambda) = \beta(\theta, \lambda) E(\lambda), \quad (1)$$

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 per unit cross-section area.




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

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, 1976]. Consider a collimated beam of light incident on the atmospheric medium, as shown in Figure 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 intensity scattered by the lamina is $I(\theta, \lambda) = \beta(\theta, \lambda) E(\lambda) dx$. The total flux scattered (in all directions) by this lamina is obtained by integrating over the entire sphere:

$$\phi(\lambda) = \beta(\lambda) E(\lambda) dx, \quad (2)$$

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. Hence, the fractional change in irradiance at location x can be written as:

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

By integrating both sides between the limits $x = 0$ and $x = d$ we get: $E(d, \lambda) = E_o(\lambda) e^{-\beta(\lambda) d}$, where, $E_o(\lambda)$ is the irradiance at the source ($x = 0$). This is Bouguer's exponential law of attenuation, derived in 1729. Its utility 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:

$$E(d, \lambda) = \frac{I_o(\lambda) e^{-\beta(\lambda) d}}{d^2}, \quad (4)$$

where, $I_o(\lambda)$ is the radiant intensity of the point source. This is Allard's law developed in 1876.

At times, attenuation due to scattering is expressed in terms of *optical thickness* which is $T = \beta(\lambda)d$. It is generally




Figure 3: Attenuation of a collimated beam of light by suspended particles. The attenuation can be derived by viewing the medium as a collection of thin sheets (laminae).

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. Finally, 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 (4). 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]).

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 [McCartney, 1976] derivation of airlight as a function of pathlength.

Consider the illumination and observation geometry shown in Figure 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 $dV = d\omega x^2 dx$. Irrespective of the exact type of envi-




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.

Environmental illumination incident upon dV , its intensity due to scattering in the direction of the observer is:

$$dI(x, \lambda) = dV k \beta(\lambda) = d\omega x^2 dx k \beta(\lambda), \quad (5)$$

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.

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 (4):

$$dE(x, \lambda) = \frac{dI(x, \lambda) e^{-\beta(\lambda) x}}{x^2}. \quad (6)$$

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

$$dL(x, \lambda) = \frac{dE(x, \lambda)}{d\omega} = \frac{dI(x, \lambda) e^{-\beta(\lambda) x}}{d\omega x^2}. \quad (7)$$

By substituting (5) 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 $x = d$:

$$L(d, \lambda) = k (1 - e^{-\beta(\lambda) d}). \quad (8)$$

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 $L_h(\infty, \lambda) = k$. Therefore, the radiance of airlight for any given pathlength d is:

$$L(d, \lambda) = L_h(\infty, \lambda) (1 - e^{-\beta(\lambda) d}). \quad (9)$$

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 angular factor k . Instead, we have the airlight radiance $L_h(\infty, \lambda)$ at the horizon, which is an observable.

4 Depths of Light Sources from Attenuation

Consider the image of an urban setting taken at *night* (see Figure 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.




Figure 5: The relative depths of sources of unknown intensities can be recovered from two images taken under different but 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 (4) as:

$$E(d, \lambda) = g \frac{I(\lambda) e^{-\beta(\lambda) d}}{d^2}, \quad (10)$$

where, $I(\lambda)$ 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 $\beta(\lambda)$ which is independent of viewing direction.

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

$$E' = \int s(\lambda) E(d, \lambda) d\lambda = \int g s(\lambda) \frac{I(\lambda) e^{-\beta(\lambda) d}}{d^2} d\lambda. \quad (11)$$

Since the spectral bandwidth of the camera is rather limited (visible light range when camera is black and white, and even narrower spectral bands when the camera is color), we will

assume the total scattering coefficient $\beta(\lambda)$ to be constant over this bandwidth. Then, we have:

$$E' = g \frac{e^{-\beta d}}{d^2} \int s(\lambda) I(\lambda) d\lambda = g \frac{e^{-\beta d}}{d^2} I'. \quad (12)$$

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:

$$R = \frac{E'_1}{E'_2} = e^{-(\beta_1 - \beta_2) d}. \quad (13)$$

Using the natural log, we obtain:

$$R' = \ln R = -(\beta_1 - \beta_2) d. \quad (14)$$

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 different conditions. Now, if we compute the DOTs of two different light sources in the scene (see Figure 5) and take their ratio, we determine the relative depths of the two source locations:

$$\frac{R'_i}{R'_j} = \frac{d_i}{d_j} \quad (15)$$

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:


$$\frac{R'_i}{\sum_{j=0}^{j=N} R'_j} = \frac{d_i}{\sum_{j=0}^{j=N} d_j} \quad (16)$$

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.




(a)



(b)



(c)



(d)

Figure 6: (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 accuracy of the computed lamp coordinates (small bright spheres). The lamp poles and the ground plane are added only to aid visualization.

5 Structure from Airlight

When we have 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, most visible scene points are not illuminated and airlight dominates. In both 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 (see [Cozman and Krotkov, 1997]).

Let a scene point with 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:

$$E'(d) = \int g s(\lambda) L(d, \lambda) d\lambda, \quad (17)$$

where, g accounts for the constant of proportionality between scene radiance and image irradiance. Substituting the model for airlight given by (9) we get:

$$E'(d) = \int g s(\lambda) L_h(\infty, \lambda) (1 - e^{-\beta(\lambda) d}) d\lambda \quad (18)$$

where, $L_h(\infty, \lambda)$ 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:

$$E'(d) = E_h'(\infty) (1 - e^{-\beta d}). \quad (19)$$

Let us define:

$$S = \frac{E_h'(\infty) - E'(d)}{E_h'(\infty)}. \quad (20)$$

By substituting (19) in the above expression and taking the natural logarithm, we get:

$$S' = \ln S = -\beta d. \quad (21)$$

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 $E_h'(\infty)$. 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 $E_h'(\infty)$ 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 from the camera, such scenes are hard to compute using stereo or structure from motion.




Figure 7: (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 (rotated) rendering of the scene. (d) Image of a mountain range taken under foggy conditions. (e) Depth map computed from the image in (d). (f) A three-dimensional (rotated) rendering of the scene. Some of the objects in these scenes are several miles away from the camera.

6 Dichromatic Atmospheric Scattering

Thus far, we have not exploited the chromatic effects of atmospheric scattering. 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 of the reflected light, the spectral composition of the reflected light can be expected to vary as it passes through the medium. Fortunately, for fog and dense haze, these shifts in the spectral composition are minimal (see [Middleton, 1952] and [Nayar and Narasimhan, 1999] 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 $E_{dt}(d, \lambda)$ of directly transmitted light and $E_a(d, \lambda)$ of airlight, which are determined by the attenuation model (10) and the airlight model (9):

$$\begin{aligned} E(d, \lambda) &= E_{dt}(d, \lambda) + E_a(d, \lambda), & (22) \\ E_{dt}(d, \lambda) &= g \frac{e^{-\beta(\lambda) d}}{d^2} L_r(\lambda), \\ E_a(d, \lambda) &= g (1 - e^{-\beta(\lambda) d}) L_h(\lambda). \end{aligned}$$

Here, $L_r(\lambda)$ is the surface radiance prior to attenuation, $L_h(\lambda)$ is the radiance of the horizon ($d = \infty$), and g is a constant that accounts for the optical settings of the imaging system. 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 (22) by $f(\lambda)$ and integrating over λ :

$$E^{(f)}(d) = E_{dt}^{(f)}(d) + E_a^{(f)}(d). \quad (23)$$

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 (see Figure 8):

$$\mathbf{E}(d) = \mathbf{E}_{dt}(d) + \mathbf{E}_a(d) \quad (24)$$

where, $\mathbf{E} = [E^{(f_1)}, E^{(f_2)}, \dots, E^{(f_n)}]^T$. As we mentioned earlier, the dependence of the scattering coefficient $\beta(\lambda)$ on the wavelength of light tends to be rather small. Therefore, except in the case of certain types of metropolitan haze, we may assume scattering to be constant with respect to wavelength ($\beta(\lambda) = \beta$). Then, expression (23) may be simplified as:

$$E^{(f)}(d) = p(d) E_r^{(f)} + q(d) E_h^{(f)}, \quad (25)$$




Figure 8: The color at an image point is the sum of two vectors, namely, the color due to transmission of light reflected by the scene point and the color due to airlight.

$$\begin{aligned} E_r^{(f)} &= \int g f(\lambda) L_r(\lambda) d\lambda, \quad E_h^{(f)} = \int g f(\lambda) L_h(\lambda) d\lambda, \\ p(d) &= \frac{e^{-\beta d}}{d^2}, \quad q(d) = (1 - e^{-\beta d}). \end{aligned} \quad (26)$$

Here, $E_r^{(f)}$ is the image irradiance due to the scene point *without* atmospheric attenuation and $E_h^{(f)}$ 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 final color measurement given by (24) can be rewritten as: $\mathbf{E}(d) = p(d) \mathbf{E}_r + q(d) \mathbf{E}_h$. Since the intensity of illumination at a scene point is expected to vary between clear and bad weather, it is more convenient to write:

$$\mathbf{E}(d) = r p(d) \hat{\mathbf{E}}_r + s q(d) \hat{\mathbf{E}}_h \quad (27)$$

where $\hat{\mathbf{E}}_r$ and $\hat{\mathbf{E}}_h$ are unit vectors and r and s are scalars.

7 Structure from 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 horizon in the bad weather image reveals the *direction* of the airlight color $\hat{\mathbf{E}}_h$. The *direction* of the color $\hat{\mathbf{E}}_r$ of each scene point is revealed by the clear weather image. Therefore, equation (27) can be used to decompose the bad weather color $\mathbf{E}(d)$ at each pixel into its two components and determine the scaled airlight magnitude $s q(d)$. The resulting airlight image is then used to compute a depth map as in section 5. Figure 9 shows experimental results obtained using the above decomposition method. 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 our computed depth maps can be attributed to this effect (see [Nayar and Narasimhan, 1999] for details).




(a)




(b)



(c)



(d)



(e)

Figure 9: Removal of fog and depth estimation using the dichromatic atmospheric scattering model. (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. (See CDROM version of the proceedings for color images.)

8 Conclusion

Ultimately, vision systems must be able to handle problems posed by bad weather. This article is no more than an initial attempt at understanding and exploiting the manifestations of weather. We summarized existing models in atmospheric optics and proposed new ones, keeping in mind the constraints faced by most vision applications. In addition, we presented three simple algorithms for recovering scene structure from one or two images, without requiring prior knowledge of atmospheric conditions. We intend to use these results as building blocks for developing more advanced weather-tolerant vision techniques.

References

- [Cozman and Krotkov, 1997] F. Cozman and E. Krotkov. Depth from scattering. *Proc. of IEEE Conf. on Comp. Vision and Pattern Recog.*, 31:801–806, 1997.
- [Hidy, 1972] ed. G. M. Hidy. *Aerosols and Atmospheric Chemistry*. Academic Press, New York, 1972.
- [Koenderink and Richards, 1992] J. J. Koenderink and W. A. Richards. Why is snow so bright? *Journal of Optical Society of America*, 9(5):643–648, 1992.
- [Koschmieder, 1924] H. Koschmieder. Theorie der horizontalen sichtweite. *Beitr. Phys. freien Atm.*, 12:33–53, 171–181, 1924.
- [Mason, 1975] B. J. Mason. *Clouds, Rain, and Rainmaking*. Cambridge University Press, Cambridge, 1975.
- [McCartney, 1976] E. J. McCartney. *Optics of the Atmosphere: Scattering by Molecules and Particles*. John Wiley and Sons, New York, 1976.
- [Middleton, 1949] W. E. K. Middleton. 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, 1949.
- [Middleton, 1952] W. E. K. Middleton. *Vision through the atmosphere*. University of Toronto Press, 1952.
- [Mie, 1908] G. Mie. A contribution to the optics of turbid media, especially colloidal metallic suspensions. *Ann. of Physics*, 25(4):377–445, 1908.
- [Minnaert, 1954] M. Minnaert. *The Nature of Light and Color in the Open Air*. Dover, New York, 1954.
- [Myers, 1968] J. N. Myers. Fog. *Scientific American*, pages 75–82, December 1968.
- [Nayar and Narasimhan, 1999] S. K. Nayar and S. G. Narasimhan. Computer Vision in Bad Weather. Technical report, Dept. of Computer Science, Columbia University, New York, (in prep.), July 1999.
- [Nieto-Vesperinas and Dainty, 1990] M. Nieto-Vesperinas and J. C. Dainty. *Scattering in volumes and surfaces*. North-Holland, New York, 1990.
- [Ohtake, 1970] T. Ohtake. Factors affecting the size distribution of raindrops and snowflakes. *Journal of Atmospheric Science*, 27:804–813, 1970.
- [Shafer, 1985] S. Shafer. Using color to separate reflection components. *Color Research and Applications*, pages 210–218, 1985.