All the Images of an Outdoor Scene

Srinivasa G. Narasimhan, Chi Wang, and Shree K. Nayar

Columbia University, New York NY 10027, USA
500 West 120th street, Rm 450 Computer Science
{srinivas, chi, nayar}@cs.columbia.edu

Abstract. The appearance of an outdoor scene depends on a variety of
factors such as viewing geometry, scene structure and reflectance (BRDF
or BTF), illumination (sun, moon, stars, street lamps), atmospheric con-
dition (clear air, fog, rain) and weathering (or aging) of materials. Over
time, these factors change, altering the way a scene appears. A large set
of images is required to study the entire variability in scene appearance.
In this paper, we present a database of high quality registered and cal-
ibrated images of a fixed outdoor scene captured every hour for over 5
months. The dataset covers a wide range of daylight and night illumi-
nation conditions, weather conditions and seasons. We describe in detail
the image acquisition and sensor calibration procedures. The images are
tagged with a variety of ground truth data such as weather and illumi-
nation conditions and actual scene depths. This database has potential
implications for vision, graphics, image processing and atmospheric sci-
ences and can be a testbed for many algorithms. We describe an example
application - image analysis in bad weather - and show how this method
can be evaluated using the images in the database. The database is avail-
able online at http://www.cs.columbia.edu/CAVE/. The data collection
is ongoing and we plan to acquire images for one year.

1 Variability in Scene Appearance

The appearance of a fixed scene depends on several factors - the viewing ge-
ometry, illumination geometry and spectrum, scene structure and reflectance
(BRDF or BTF) and the medium (say, atmosphere) in which the scene is im-
mersed. The estimation of one or more of these appearance parameters from one
or more images of the scene has been an important part of research in computer
vision. Several researchers have focused on solving this inverse problem under
specific conditions of illumination (constant or smoothly varying), scene struc-
ture (no discontinuities), BRDF (lambertian) and transparent media (pure air).
Images captured to evaluate their methods adhere to the specific conditions.
While understanding each of these specific cases is important, modeling scene
appearance in the most general setting is ultimately the goal of a vision system.
To model, develop and evaluate such a general vision system, it is critical to
collect a comprehensive set of images that describes the complete variability in
the appearance of a scene. Several research groups have collected images of a
scene (for example, faces, textures, objects) under varying lighting conditions

and/or viewpoints in *controlled lab environments*. The CMU-PIE database [1] has 40000 facial images under different poses, illumination directions and facial expressions. The FERET [2] database consists of 1196 images of faces with varying facial expressions. Similarly, the Yale Face database [3] has around 165 images taken under different lighting, pose and occlusion configurations. The SLAM database [4] provides a set of 1500 images of toy objects under different poses. The color constancy dataset collected by Funt et al. [5] provides a large set of images of objects (boxes, books and so on) acquired under different poses and with different illuminants (fluorescent, halogen, etc). The CURET database [6] provides a set of 12000 images of real world textures under 200 illumination and viewing configurations. It also provides an additional set of 14000 Bi-directional Texture Function (BTF) measurements of 61 real world surfaces.

Several databases of images of outdoor scenes have also been collected. The "natural stimuli collection" [7] has around 4000 images of natural scenes taken on clear, foggy and hazy days. Parraga et al. [8] provide a hyperspectral dataset of 29 natural scenes. The MIT city scanning project [9] provides a set of 10000 geo-referenced calibrated images acquired over a wide area of the MIT campus. These databases, however, do not cover the complete appearance variability (due to all outdoor illumination and weather conditions) in any one particular scene. Finally, web-cams [10] used for surveillance capture images regularly over long periods of time. However, they are usually low quality, non-calibrated, not tagged with ground truth data and focus only on activity in the scene. Note that the references we have provided for various databases are by no means complete. We refer the reader to [11] for a more comprehensive listing.

In this paper, we present a set of very high quality registered images of an outdoor scene, captured regularly for a period of 5 months. The viewpoint (or sensor) and the scene are fixed over time. Such a dataset is a comprehensive collection of images under a wide variety of seasons, weather and illumination conditions. This database serves a dual purpose; it provides an extensive testbed for the evaluation of existing appearance models, and at the same time can provide insight needed to develop new appearance models. To our knowledge, this is the first effort to collect such data in a principled manner, for an extended time period. The data collection is ongoing and we plan to acquire images for one year.

We begin by describing the image acquisition method, the sensor calibration procedures, and the ground truth data collected with each image. Next, we illustrate the various factors that effect scene appearance using images from our database captured over 5 months. We demonstrate thorough evaluation of an existing model for outdoor weather analysis, using the image database.

2   Data Acquisition

2.1   Scene and Sensor

The scene we image is an urban scene with buildings, trees and sky. The distances of these buildings range from about 20 meters to about 5 kilometers. The large

Fig. 1. Acquisition setup

distance range facilitates the observation of weather effects on scene appearance. See figure 5 for the entire field of view. The digital camera we use for image capture is a single CCD KODAK Professional DCS 315 (see figure 1). As usual, irradiance is measured using 3 broadband R, G, and B color filters. An AF Nikkor 24 mm - 70 mm zoom lens is attached to the camera.
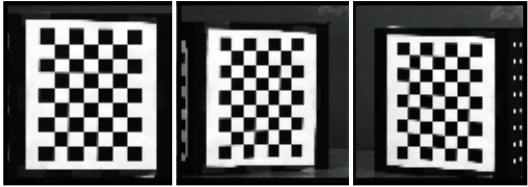
2.2   Acquisition Setup

The setup for acquiring images is shown in figure 1. The camera is rigidly mounted over a pan-tilt head which is fixed rigidly to a weather-proof box (see black box in figure 1). The weather-proof box is coated on the inside with two coats of black paint to prevent inter-reflections within the box. An anti-reflection glass plate is attached to the front of this box through which the camera views the scene. Between the camera and the anti-reflection plate, is a filter holder (for, say, narrow band spectral filters). The entire box with the camera and the anti-reflection glass plate is mounted on a panel rigidly attached to a window.

2.3   Image Quality and Quantity

Images are captured automatically every hour for 20 hours each day (on an average). The spatial resolution of each image is 1520×1008 pixels and the intensity resolution is 10 bits per pixel per color channel. Currently, we have acquired images for over 150 days. In total, the database has around 3000 images. Due to maintainance issues that arise from prolonged camera usage (camera power failures and mechanical problems in controlling camera shutter), we have had to remove the camera twice from the enclosure. We believe the resulting loss of few days in the database can be tolerated since the dataset has enormous redundancy. The new image sets are registered with existing ones using the matlab image registration utility.

To capture both subtle and large changes in illumination and weather, high dynamic range images are required. So, we acquire images with multiple exposures (by changing the camera shutter speed while keeping the aperture constant) and apply prior techniques to compute a high dynamic range image (≈ 12 bits per pixel) of the scene. Since the illumination intensity is expected to vary with time, the set of exposures are chosen adaptively. First, an auto-exposure image is taken and its shutter speed is noted. Then 4 more images are captured with exposures around this auto-exposure value. This type of adaptive exposure selection is commonly used by photographers and is called *exposure bracketing*.

Focal Length	2845 mm
$\triangle u_0$, $\triangle v_0$	7.2 , 9.5 pixels
Radial Dist. (C_1)	0.07
Reprojection Err	0.18 pixels

Fig. 2. Geometric calibration using planar checkerboard patterns. Table shows estimated intrinsic parameters. The distortion parameters not shown are set to zero.

3   Sensor Calibration

3.1   Geometric Calibration

Geometric calibration constitutes the estimation of the geometric mapping between 3D scene points and their image projections. Since the calibration was done in a location different from that used for image acquisition, we estimate only the intrinsic parameters of the camera. Intrinsic parameters include the effective focal length, f, skew, s, center of projection (u_0, v_0) and distortion parameters, $C_1...C_n$ (radial) and P_1, P_2 (tangential). Then, the relation between observed image coordinates and the 3D scene coordinates of a scene point is:

$$u_i = D_u s(x_i \frac{f}{z_i} + \delta u_i^{(r)} + \delta u_i^{(t)}) + u_0$$

$$v_i = D_v(y_i \frac{f}{z_i} + \delta v_i^{(r)} + \delta v_i^{(t)}) + v_0 , \tag{1}$$

where $(\delta u_i^{(r)}, \delta v_i^{(r)})$ are functions of the radial distortion parameters, and $(\delta u_i^{(t)}, \delta v_i^{(t)})$ are functions of the tangential distortion parameters. D_u, D_v are the conversion factors from pixels to millimeters. See [12] for more details. We captured the images of a planar checkerboard pattern under various orientations (see figure 2). The corresponding corners of the checkerboard patterns in these images were marked. These corresponding corner points were input to a calibration routine [13] to obtain the intrinsic parameters. Figure 2 shows the estimated intrinsic parameters. The CCD pixels are square and hence skew is assumed to be 1. The deviation of the principal point from the image center is given by $\Delta u_0, \Delta v_0$. Only the first radial distortion parameter, C_1, is shown. The remaining distortion parameters are set to zero.

3.2   Radiometric Calibration

Analysis of image irradiance using measured pixel brightness requires the radiometric response of the sensor. The radiometric response of a sensor is the mapping, g, from image irradiance, I, to measured pixel brightness, $M : M = g(I)$. Then, the process of obtaining I from M: $I = g^{-1}(M)$, upto a global scale factor, is termed as radiometric calibration.

(a) High Exposure   (b) Auto Exposure   (c) Low Exposure

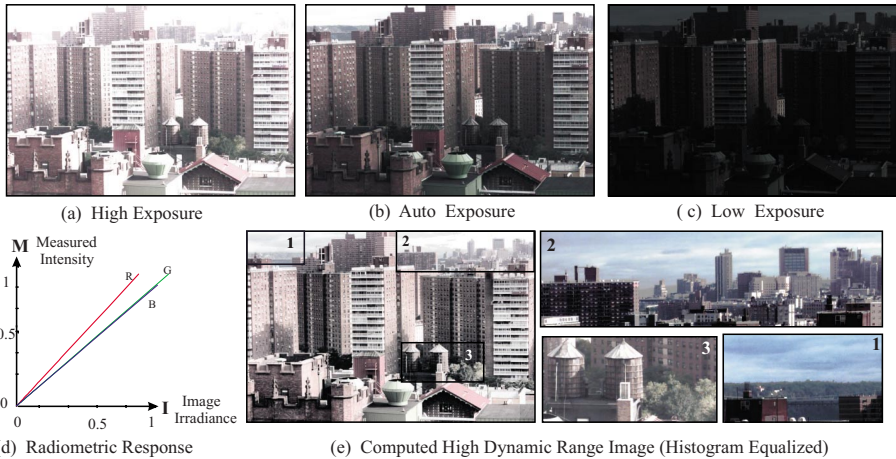(d) Radiometric Response   (e) Computed High Dynamic Range Image (Histogram Equalized)

Fig. 3. Radiometric Self-Calibration. (a) - (c) Three images (10 bits per pixel per RGB channel) captured with different camera exposures. (d) Computed radiometric response functions of the 3 RGB channels. The response functions are linear with slopes 1.5923, 1.005 and 0.982 for R, G, and B respectively. The colors can be balanced by normalizing the slope of each response function. (e) Histogram equalized high dynamic range image irradiance map obtained by combining images taken with multiple exposures. Insets indicate that the dynamic range in this image is much higher than the dynamic range in any image captured using a single exposure.

The response functions of CCD cameras (without considering the gamma or color corrections applied to the CCD readouts) are close to linear. We computed the response functions of the 3 RGB channels separately using Mitsunaga and Nayar's [14] radiometric self-calibration method. In this method, images captured with multiple exposures and the their relative exposure values are used to estimate the inverse response function in polynomial form. The results of the calibration are shown in the plots of figure 3(d). Notice that the response functions of R, G, and B are linear and they have different slopes - 1.5923, 1.005 and 0.982 respectively. To balance the colors, we normalize the response functions by the respective slopes. The images taken with different exposures are linearized using the computed response function. A high dynamic range image irradiance map (see figure 3) is obtained by using a weighted combination of the linearized images. This image has significantly more dynamic range then any of the original images taken with single exposures [15]. The high dynamic range can prove very useful when analyzing both subtle and large changes in weather and illumination.

4   Ground Truth Data

Any database is incomplete without the accompanying ground truth. We have tagged our images with a variety of ground truth information. Most important categories of the ground truth we collected are scene depth and weather information. The depths of scene points are mainly obtained using satellite digital
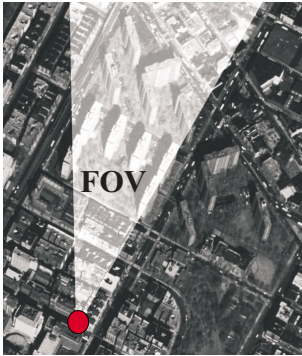
Conditions at 2001.03.06 11:51am	
Wind	NNW (340^O) 10 MPH
Visibility	1 1/4 mile(s)
Sky conditions	Overcast
Weather	Light snow, Mist
Precipitation last hour	A trace
Temperature	32.0 F (0.0 C)
Dew Point	32.0 F (0.0 C)
Relative Humidity	100%

Fig. 4. Sample ground truth data. [Left] A satellite digital orthophoto of a portion of the scene. The red spot indicates the position of the sensor and bright region indicates the field of view. Arcview [17] is used to measure the orthographic distance between any two scene points (seen in top view) with an accuracy of 1 meter. [Right] The weather data obtained from National Weather Service websites [18].

orthophotos supplied by the United States Geological Survey [16]. Arcview [17] is a mapping software that is used to measure the orthographic distance between two scene points (visible in the orthophotos) up to an accuracy of 1 meter. See figure 4 for an example of a satellite orthophoto. Note that accurate depth is not available at all pixels. However, since the field of view consists of mainly vertical buildings, rough planar models can be used. The position (longitude, latitude and altitude) of the sensor is included in the database. This information along with the date and time of day, can be used to accurately compute sun and moon orientation relative to the sensor. For exact equations, see [19,20].

Every hour we automatically collect standard weather information from the National Weather Service web sites [18]. This includes information about sky condition (sunny, cloudy), weather condition (clear, fog, haze, rain), visibility, temperature, pressure, humidity and wind (see figure 4). Such information can be used to estimate the scattering coefficient of the atmosphere [21].

5   WILD: Weather and ILlumination Database

We illustrate the variability in scene appearance due to weather, illumination, season changes, and surface weathering using images from our dataset captured over five months.

5.1   Variation in Illumination

The distribution of environmental illumination on a scene produces a wide variety of scene appearances. Commonly noticed effects include shadows, colors of sunrise and sunset, and illumination from stars and moon at night. The human visual system relies on illumination in the scene to perceive scene reflectance (retinex and color constancy [22]) and shape [23] correctly. As a result, rendering a scene with *consistent* illumination is critical for realism in graphics. Considerable amount of effort has been put into modeling outdoor illumination.

The book "Daylight and its Spectrum" [24] provides a compendium of color and intensity distributions of skylight for many years of the 20th century. Daylight spectrum has also been represented using a set of linear bases [25]. Works that model clouds and their effect on the ambient illumination also exist in literature [26,27]. In graphics, scenes have been rendered under different daylight [28] and night illuminations [29]. Shadows are a powerful cue for shape and illumination perception. Rendering shadows and extracting shape information from shadows [30,31] are also important problems.

Let us consider the various sources of illumination in any outdoor scene. The primary sources (self-luminous) include the sun during the day, the stars and lamps during night. There are numerous other secondary illumination sources such as skylight, ground light, moonlight, airlight [32] (due to scattering of light by the atmosphere), and scene points themselves (inter-reflections [33]). Our goal is to include the effects of all these sources in one comprehensive database. Figure 5 shows 6 images from our database illustrating the various shadow configurations on a sunny day. Figure 6 shows different illumination colors and intensities at sunrise, noon, sunset and night. Figure 7 depicts the variations in ambient lighting due to varying cloud covers.

When viewed up-close, rough surfaces appear to have $3D$ textures (due to surface height variations) rather than $2D$ textures. The appearance of $3D$ textures has been modeled in [6]. Figure 8 shows the appearance of a rooftop with ridges at different times of the day. Notice the change in appearance of cast shadows due to surface height variations.

The above variability in a single database facilitates research groups to study the illumination effects individually as well as simultaneously. For instance, one may use just sunny day images at one particular time of day, when the sun position remains constant and shadows do not change. In another instance, one can consider images captured only on cloudy days to model scenes under ambient lighting.

5.2   Variation in Weather Conditions

Most vision algorithms assume that light travels from a scene point to an observer unaltered. This is true only on clear days. In bad weather, however, the radiance from a scene point is severely scattered and thus, the appearance of a scene changes dramatically. The exact nature of scattering depends on a variety of factors such as shapes, orientations and densities of atmospheric particles and the colors, polarizations, and intensities of incident light [34]. Recently, there has been work on computing scene depths from bad weather images [35] as well as to remove weather effects from images. Some works extract structure and clear day scene colors using images of a scene taken under different weather conditions [36] or through different polarizer orientations [37]. Other works use pre-computed or measured distances to restore scene contrast [38,39]. In graphics, scenes (including skies) have been rendered considering multiple scattering of light in the atmosphere [28,40].

Fig. 5. Images illustrating different shadow configurations on a clear and sunny day. Shadows provide cues for illumination direction and the scene structure. Notice the positions of the sharp shadows on the buildings.
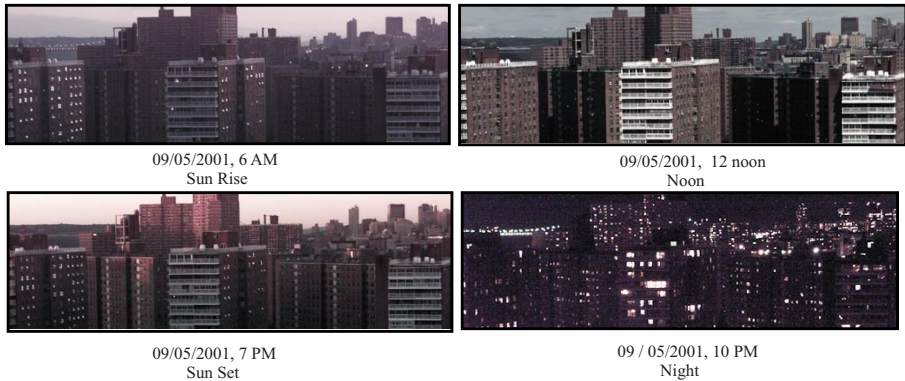

Fig. 6. Images illustrating the various colors and intensities of illumination at sunrise, noon, sunset and night. Notice the significant change in the colors of the sky.
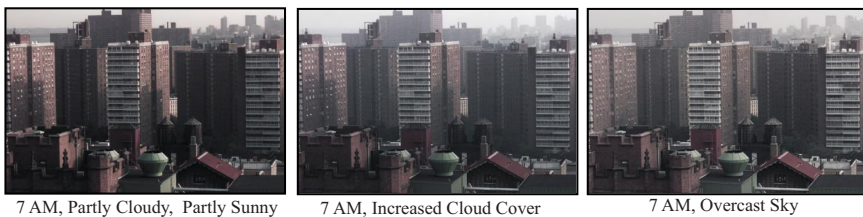

Fig. 7. Images showing various levels of cloud cover. The image on the left shows the appearance of the scene with a few scattered clouds. The two images on the right were taken under mostly cloudy and completely overcast conditions. Notice the soft shadows due to predominant ambient lighting.

9 AM	12 Noon	3 PM

Fig. 8. When viewed at close proximity (fine scale), the appearances of surfaces should be modeled using the bi-directional texture function (BTF) instead of the BRDF. Notice the change in cast shadows due to the ridges on the rooftop. All images are histogram equalized to aid visualization.
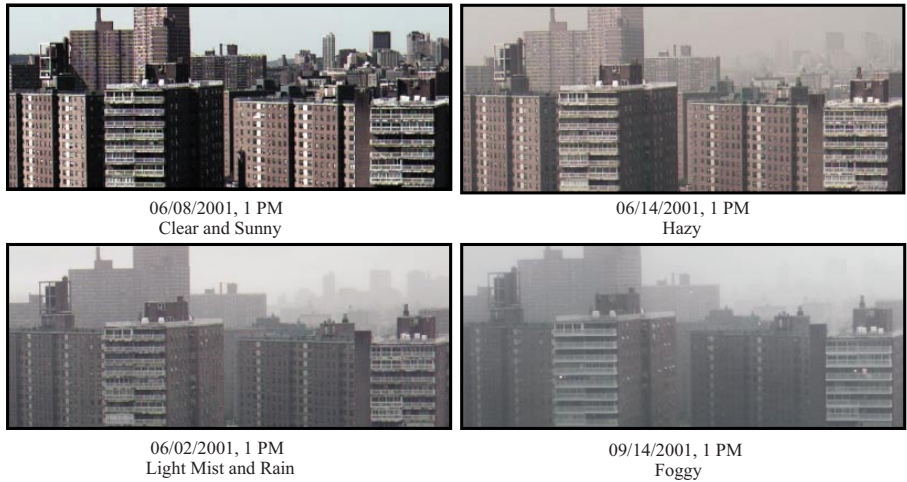

06/08/2001, 1 PM	06/14/2001, 1 PM
Clear and Sunny	Hazy
06/02/2001, 1 PM	09/14/2001, 1 PM
Light Mist and Rain	Foggy

Fig. 9. Images taken at the same time of day but under different weather conditions. Notice the degradation in visibility, especially, of far away scene points in bad weather.

05/28/01, 11 PM	05/26/01, 11 PM
Clear Night	Misty Night

Fig. 10. Night images showing light sources under clear and misty conditions. The sources appear like specks of light on a clear night. Notice the light spreading due to multiple scatterings on a misty night.

How can the above models for weather analysis be evaluated? Under what conditions do these models fail or perform well? To satisfactorily answer such questions, we need to evaluate the performance of such models and algorithms on images of a scene captured under a wide variety of illumination and weather

conditions. Our database includes images of the scene under many atmospheric conditions including clear and sunny, fog, haze, mist, rain and snow. Figure 9 shows 4 images of the same scene captured under different weather conditions. Notice the significant reduction in contrast (and increase in blurring) in far away buildings. Broadening of light beams due to multiple scatterings in the atmosphere is clearly illustrated by the lamps imaged at night (see figure 10).

Consider the event of mild fog setting in before sunrise, becoming dense as time progresses and finally clearing by noon. We believe that such lengthy, time varying processes can be studied better using our database. Study of such processes have potential implications for image based rendering.

5.3   Example Evaluation: Weather Analysis

Consider images taken under different weather conditions. The observed color, $\mathbf{E}$, of a scene point in bad weather is linearly related to its clear day color direction, $\hat{\mathbf{D}}$, and the color of the weather condition (say, fog or haze), $\hat{\mathbf{A}}$, by the dichromatic model of atmospheric scattering [35]:

$$\mathbf{E} = m\,\hat{\mathbf{D}} + n\,\hat{\mathbf{A}}\,. \tag{2}$$

So, $\mathbf{E}$, $\hat{\mathbf{A}}$, and $\hat{\mathbf{D}}$ lie on the same "dichromatic plane" (see figure 11). Here, $m = E_\infty \rho e^{-\beta d}$, $n = E_\infty(1 - e^{-\beta d})$, d is the depth of the scene point from the observer, β is called the scattering coefficient of the atmosphere [41], E_∞ is the brightness at the horizon, and ρ is a function of the scene point BRDF [36]. Under what weather conditions is the dichromatic model valid? How well does the dichromatic model describe the colors of scene points under a particular weather condition (say, mist)? Figure 11 shows the results of evaluating the model for fog, haze, mist and rain using multiple (5 in this case) images taken under each weather condition.

Based on the dichromatic model, Narasimhan and Nayar [36] developed constraints on changes in scene colors under different atmospheric conditions. Using these constraints, they developed algorithms to compute $3D$ structure and clear day colors of a scene from two or more images taken under different but unknown weather conditions. How do we evaluate the performance of such an algorithm? Figure 12 shows the comparison of the defogged image with the actual clear day image of the scene under similar illumination spectra (overcast skies). The accuracy of the computed scaled depth, βd, is compared against the ground truth relative depth values obtained from the satellite orthophotos. This demonstrates that models and algorithms pertaining to weather can be evaluated more thoroughly using images from this database.

Narasimhan and Nayar's algorithm described above shows that clear day images can be obtained using two images taken under different weather (say, foggy) conditions. Given two foggy images, can we generate novel foggy images? We show this is possible using the defogged color $\rho\hat{\mathbf{D}}$, fog color $\hat{\mathbf{A}}$, and optical depth βd, at each pixel. Note that these quantities can be computed using the above algorithm. Scattering coefficient β is a measure of the density of fog. Thus, the density of fog can either be increased or decreased by appropriately scaling

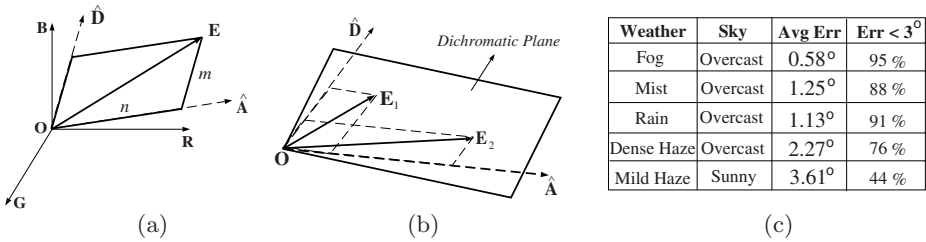| Weather	Sky	Avg Err	Err < 3°
Fog	Overcast	0.58°	95 %
Mist	Overcast	1.25°	88 %
Rain	Overcast	1.13°	91 %
Dense Haze	Overcast	2.27°	76 %
Mild Haze	Sunny	3.61°	44 %

(a)   (b)   (c)

Fig. 11. Dichromatic plane geometry and its evaluation. (a) Dichromatic Model [35]. (b) The observed color vectors $\mathbf{E}_i$ of a scene point under different (two in this case) weather conditions (say, mild fog and dense fog) lie on a plane called the dichromatic plane. (c) Experimental verification of the dichromatic model with the scene imaged 5 times under each of the different foggy, misty, rainy and hazy conditions. The third column is the mean angular deviation (in degrees) of the observed scene color vectors from the estimated dichromatic planes, over 1.5 megapixels in the images. The fourth column provides the percentage of pixels whose color vectors were within 3 degrees of the estimated dichromatic plane. Note that the dichromatic model works well for fog, mist, rain and dense haze under overcast skies. For mild haze conditions under sunny skies, the model does not perform well. Such evaluation is possible only since our database has several images under each weather condition.

the optical depth βd. Substituting the scaled optical depth $k \beta d$, clear day color $\rho \hat{\mathbf{D}}$, and fog color $\hat{\mathbf{A}}$, into equation 2, we compute the colors of scene points under novel fog conditions. The horizon brightness E_∞ is kept constant since it is just a global scale factor. Figure 13 shows 4 novel images generated with increasing fog densities.

5.4   Seasonal Variations

The types of outdoor illumination and weather conditions change with seasons. For instance, the intensity distribution of sunlight and skylight differ from summer to winter [24]. Similarly, the atmospheric conditions that manifest in fall are significantly different from those that occur in winter. Models of the atmosphere in different seasons can be found in [21] and other related papers [42]. Since we acquire the images for over a year, changes in scene appearance due to changes in seasons can be studied. For instance, one might easily compare the images taken on a clear day in spring with images taken on a clear day in winter under identical sensor settings. Figure 14 shows 2 images taken at the same time of day in summer and fall.

5.5   Surface Weathering

Over substantial periods of time, we commonly see oxidation of materials (rusting), deposition of dirt on materials and materials becoming wet or dry. These effects are important for realistic scene rendering and have been modeled by Dorsey and Hanrahan [43]. Since our images have high spatial resolution, portions of the image corresponding to small regions in the scene (say, a portion of

5 PM, Fog

6 PM, Fog
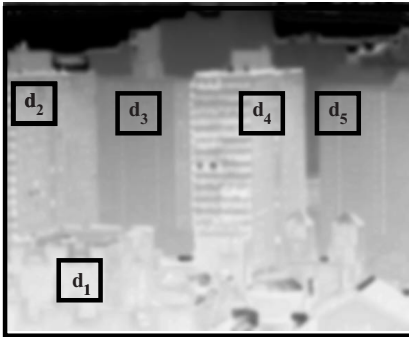
Computed Depth Map
(Brightened)

Relative Depth	Ground Truth	Computed Value
d3 / d1	9.55	10.457
d2 / d1	6.98	7.72
d4 / d1	6.41	6.49
d5 / d1	8.82	9.30

Relative Depth Verification
using Satellite Orthophoto Data

Computed Defogged Image

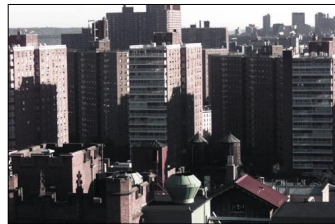Actual clear day image
(under mostly cloudy skies)

Fig. 12. Computing structure and defogging from two foggy images. Table comparing the computed relative depths with ground truth relative depths (obtained using satellite orthophotos) of 5 different regions, $d_1 - d_5$, in the scene. The relative depths are averaged over small neighborhoods. The window regions do not remain constant and thus produce erroneous depth values. All the images are contrast stretched for display purposes.

Fig. 13. Generation of novel images with increasing fog densities (or scattering coefficients). The relative scattering coefficients used in this case are β, 2β, 3β and 4β respectively.

Summer Season, 06/15/2001, 11 AM   Fall Season, 09/15/2001, 11 AM

Fig. 14. Images taken at the same time of day but on days in summer and fall. Both the images were taken on clear and sunny days. Notice the subtle differences in colors and the positions of shadows.

Fig. 15. Portions of a rooftop in the scene when it is dry, partly and completely wet.

a wall) can be analyzed. Figure 15 shows a small patch in the scene when it is dry and wet.

6   Summary

The general appearance of a scene depends on a variety of factors such as illumination, scene reflectance and structure, and the medium in which the scene is immersed. Several research groups have collected and analyzed images of scenes under different configurations of illuminations (both spectrum and direction), and viewpoints, in controlled lab environments. However, the processes that effect outdoor scene appearance such as climate, weather and illumination are very different from indoor situations. Ultimately, vision algorithms are expected to work robustly in outdoor environments. This necessitates a principled collection and study of images of an outdoor scene under all illumination and weather conditions. We have collected a large set of high quality registered images of an outdoor urban scene captured periodically for five months. We described the acquisition process, calibration processes and ground truth data collected. The utility of the database was demonstrated by evaluating an existing model and algorithm for image analysis in bad weather. This dataset has potential used in vision, graphics and atmospheric sciences.

References

1. Sim, T., Baker, S., Bsat, M.: The CMU pose, illumination and expression (PIE) database of faces. Tech Report CMU-RI-TR-01-02 (2001)
2. Phillips, P., Wechsler, H., Huang, J., Rauss, P.: The feret database and evaluation procedure for face-recognition algorithms. Image and Vision Computing **16** (1998)
3. Belhumeur, P.N., Hespanha, J.P., Kriegman, D.J.: Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. PAMI **19** (1997) 711–720
4. Murase, H., Nayar, S., Nene, S.: Software library for appearance matching (slam). In: ARPA94. (1994) I:733–737
5. Funt, B., Barnard, K., Martin, L.: Is machine colour constancy good enough? In: Proc ECCV. (1998)
6. Dana, K., Nayar, S., van Ginneken, B., Koenderink, J.: Reflectance and texture of real-world surfaces. In: Proc CVPR. (1997) 151–157
7. van Hateren, J.H., van der Schaaf, A.: Independent component filters of natural images compared with simple cells in primary visual cortex. In Proc. Royal Society of London **B 265** (1998) 359 – 366
8. Parraga, C.A., Brelstaff, G.J., Troscianko, T., Moorhead, I.: Color and illumination information in natural scenes. JOSA A **15** (1998) 563–569
9. Teller, S., Antone, M., Bodnar, Z., Bosse, M., Coorg, S., Jethwa, M., Master, N.: Calibrated registered images of an extended urban area. In: Proc. CVPR. (2001)
10. Hazecam: A live webcam. (In: http://www.hazecam.net)
11. CMUPage: The computer vision home page. (In: http://www.cs.cmu.edu/cil/vision.html)
12. Heikkila, J., Silven, O.: A four-step camera calibration procedure with implicit image correction. In: Proc. CVPR. (1997) 1106–1112

13. Bouguet, J.Y.: Camera calibration toolbox. (In: http://www.vision.caltech.edu/bouguetj/calib doc)
14. Mitsunaga, T., Nayar, S.: Radiometric self calibration. In: CVPR. (1999) I:374–380
15. Nayar, S., Mitsunaga, T.: High dynamic range imaging: Spatially varying pixel exposures. In: Proc. CVPR. (2000) I:472–479
16. USGS: U.S. Geological Survey Mapping home page. (In: http://mapping.usgs.gov)
17. ESRI: The ESRI home page. (In: http://www.esri.com)
18. NWS: The national weather service home page. (In: http://www.nws.noaa.gov)
19. Aurora: Sun angle basics. (In: http://aurora.crest.org/basics/solar/angle/)
20. Naval-Observatory: The Astronomical Alamanac for the Year 2001. US R. G. O Government Printing Office (2001)
21. Acharya, P.K., Berk, A., Anderson, G.P., Larsen, N.F., Tsay, S.C., Stamnes, K.H.: Modtran4: Multiple scattering and BRDF upgrades to modtran. SPIE Proc. Optical Spectroscopic Techniques and Instrumentation for Atmospheric and Space Research III **3756** (1999)
22. Land, E., McCann, J.: Lightness and retinex theory. JOSA **61** (1971)
23. Ramachandran, V.: Perceiving shape from shading. SciAmer **259** (1988) 76–83
24. Henderson, S.T.: Daylight and its Spectrum. New York : Wiley (1977)
25. Slater, D., Healey, G.: What is the spectral dimensionality of illumination functions in outdoor scenes? In: Proc. CVPR. (1998) 105–110
26. Moon, P., Spencer, D.: Illumination from a non-uniform sky. Illum Engg **37** (1942)
27. Gordon, J., Church, P.: Overcast sky luminances and directional luminous reflectances of objects and backgrounds under overcast skies. App. Optics **5** (1966)
28. Tadamura, K., Nakamae, E., Kaneda, K., Baba, M., Yamashita, H., Nishita, T.: Modeling of skylight and rendering of outdoor scenes. In: Eurographics. (1993)
29. Jensen, H.W., Durand, F., Stark, M.M., Premoze, S., Dorsey, J., Shirley, P.: A physically-based night sky model. In Proc. SIGGRAPH (2001)
30. Kender, J., Smith, E.: Shape from darkness. In: Proc. ICCV. (1987) 539–546
31. Kriegman, D., Belhumeur, P.: What shadows reveal about object structure. JOSA-A **18** (2001) 1804–1813
32. Koschmieder, H.: Theorie der horizontalen sichtweite. Beitr. Phys. freien Atm., **12** (1924) 33–53,171–181
33. Nayar, S., Ikeuchi, K., Kanade, T.: Shape from interreflections. IJCV **6** (1991)
34. Hulst, V.D.: Light Scattering by small Particles. John Wiley and Sons (1957)
35. Nayar, S., Narasimhan, S.: Vision in bad weather. In Proc. ICCV (1999)
36. Narasimhan, S., Nayar, S.: Chromatic framework for vision in bad weather. In Proc. CVPR (2000)
37. Schechner, Y., Narasimhan, S., Nayar, S.: Instant dehazing of images using polarization. In Proc. CVPR (2001)
38. Oakley, J., Satherley, B.: Improving image quality in poor visibility conditions using a physical model for degradation. IEEE Trans. on Image Processing **7** (1998)
39. Yitzhaky, Y., Dror, I., Kopeika, N.: Restoration of altmospherically blurred images according to weather-predicted atmospheric modulation transfer function. Optical Engineering **36** (1998)
40. T. Nishita, Y. Dobashi, E.N.: Display of clouds taking into account multiple anisotropic scattering and sky light. SIGGRAPH (1996)
41. McCartney, E.: Optics of the Atmosphere: Scattering by molecules and particles. John Wiley and Sons (1975)
42. AFRL/VSBM: Modtran. (In: http://www.vsbm.plh.af.mil/soft/modtran.html)
43. Dorsey, J., Hanrahan, P.: Digital materials and virtual weathering. Scientific American **282** (2000)