

Visibility Cameras: Where and How to Look

Nathan Graves
Electrical Engineering & Computer Science
University of California, Merced
Merced, CA 95343
ngraves@ucmerced.edu

Shawn Newsam
Electrical Engineering & Computer Science
University of California, Merced
Merced, CA 95343
snewsam@ucmerced.edu

ABSTRACT

This paper investigates image processing and pattern recognition techniques to estimate light extinction based on the visual content of images from static cameras. We propose two predictive models that incorporate multiple scene regions into the estimation: regression trees and multivariate linear regression. Incorporating multiple regions is important since regions at different distances are effective for estimating light extinction under different visibility regimes. We evaluate our models using a sizable dataset of images and ground truth light extinction values from a visibility camera system in Phoenix, Arizona.

Categories and Subject Descriptors

I.4.8 [Image Processing and Computer Vision]: Scene Analysis; I.5.4 [Pattern Recognition]: Applications

Keywords

Atmospheric visibility, multivariate linear regression, regression trees

1. INTRODUCTION

Atmospheric visibility can be a useful indicator of atmospheric pollution resulting from suspended particulates especially in drier climates. This coupled with the rapidly growing number of cameras in our ecosystem motivates image-based visibility estimation as an appealing complement to traditional means of monitoring air pollution since specialized equipment for measuring pollution is comparatively expensive. So-called visibility camera systems are already seeing widespread deployment. For example, the Interagency Monitoring of Protected Visual Environments (IMPROVE) program has installed and maintains cameras in over two dozen national parks. In addition, regional air quality agencies have deployed visibility camera systems in over 30 cities. More broadly, though, there are potentially tens of thousands of web, surveillance, traffic, and other cameras which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MAED'12, November 2, 2012, Nara, Japan.

Copyright 2012 ACM 978-1-4503-1588-3/12/11 ...\$15.00.



(a) $b_{ext} = 25 \text{ Mm}^{-1}$

(b) $b_{ext} = 174 \text{ Mm}^{-1}$

Figure 1: We investigate methods to estimate light extinction b_{ext} using visibility cameras. Shown above are images corresponding to good and poor conditions from a visibility camera in Phoenix, Arizona. Ground truth readings from a transmissometer appear in the captions.

could be used to monitor atmospheric visibility and thus air pollution.

The work in this paper represents a step towards using multimedia data, in particular images from static cameras, to perform quantitative estimation of atmospheric visibility. We investigate image processing and pattern recognition techniques to derive predictive models of light extinction based on image content. Light extinction captures the joint effects of light scattering and absorption that result from particulates in the atmosphere. We relate light extinction to atmospheric transmission and then finally to measures of image contrast using a log-linear model. We evaluate our approach using a sizable dataset of images and ground truth light extinction values.

We demonstrate that models which incorporate scene regions located at different distances from the camera are more effective than models which incorporate only a single region. This result is due to the fact that far regions are not useful when visibility is relatively poor since they are not observable at all and close regions are not useful when visibility is relatively good since there is not enough intervening atmosphere to reduce visual acuity by a measurable amount. We investigate two methods for incorporating multiple regions: linear regression trees and multivariate linear regression. We also investigate different numbers and sizes of image regions. Our results improve upon our previous approach [7] which uses only a single region by over 30%.

2. RELATED WORK

There is a sizable body of work on the related problem of improving the fidelity of images taken under hazy or otherwise atmospherically degraded conditions. This includes work by Narasimhan and Nayar on using physics-based models to improve a single image [20, 21] and using multiple images of the same scene but under different conditions [19, 18, 17]; work by Schechner and colleagues on using polarization to improve one or more images [24, 25, 14, 27, 15]; and work by He et al. on using a dark channel prior to dehaze a single image [8]. The objective of this paper, however, is to derive quantitative estimates of atmospheric visibility and so these works are not directly applicable.

There is a much smaller body of work on using images to measure atmospheric visibility. Caimi et al. [5] review the theoretical foundations of visibility estimation using image features such as contrast, and describe a Digital Camera Visibility Sensor system, but they do not apply their technique to real data. Kim and Kim [9] investigate the correlation between hue, saturation, and intensity, and visual range in traditional slide photographs. They conclude that atmospheric haze does not significantly affect the hue of the sky but strongly affects the saturation of the sky, but they do not use the image features to estimate visibility. Baumer et al. [3] use an image gradient based approach to estimate visual range using digital cameras but their technique requires the detection of a large number of targets, some only a few pixels in size. This detection step is sensitive to parameter settings and is not robust to camera movement. Also, for ranges over 10 km, they only compare their estimates to human observations which have limited accuracy. Luo et al. [12] use Fourier analysis as well as the image gradient to estimate visibility but they also only compare their estimates to human observations. Raina et al. [23] do compare their estimates to measurements taken using a transmissometer-like device but their approach requires the manual extraction of visual targets. The work by Molenar et al. [13] is closest to the proposed technique in that it is fully automated and the results are compared to transmissometer readings. However, their technique uses a single distant and thus small mountain peak to estimate contrast and thus is very sensitive to camera movement and is unlikely to perform well under varying visibility regimes.

In contrast to the works above, our approach is fully automated, does not rely on the detection and segmentation of small targets, is robust to modest camera movement, and performs favorably when compared to ground truth measurements acquired using specialized equipment.

3. THE PROBLEM

Our goal is to estimate visibility from a static image. Reduced visibility by the intervening atmosphere is mainly due to three factors: 1) light radiating from the scene is absorbed before it reaches an observer; 2) light radiating from the scene is scattered out of the visual pathway of an observer; and 3) ambient light is scattered into the visual pathway of an observer. The combined effect of the absorption and scattering is referred to as the total *light extinction*. The higher the light extinction, the poorer the visibility.

Light extinction is typically measured using a transmissometer [2, 11, 4]. This device consists of a light source (transmitter) and light detector (receiver), generally sepa-

rated by a distance of several kilometers, and assesses visibility impairment by measuring the amount of light lost over this known distance. The transmitter emits a uniform light beam of known constant intensity. The receiver separates this light from ambient light, computes the amount of light lost, and reports the extinction coefficient b_{ext} which is commonly measured in units of inverse megameters ($1 \text{ Mm}^{-1} = 1.0 \times 10^{-6} \text{ m}^{-1}$).

Our goal is to measure b_{ext} using a camera instead of a transmissometer. We do this by noting that b_{ext} is inversely related to *atmospheric transmission* t through the exponential equation [26]

$$t = \exp^{-b_{ext}r} \quad (1)$$

where r is the distance of the scene. Further, atmospheric transmission t can be related to the observed image I through [28, 6, 16, 18, 8]

$$I = Jt + A(1 - t) \quad (2)$$

where J is the scene radiance and A is the ambient (atmospheric) light. The first term on the right side this equation is inversely related to the amount of light radiating from the scene that is scattered out of the visual pathway and thus increases with improved transmission. The second term is the amount of ambient light typically from the sun that is scattered into the visual pathway and thus decreases with improved transmission. In the extremes, the perceived image can either be just the scene radiance, i.e., no atmospheric interference, or just the scattered ambient light.

Intuitively, reduced visibility results in an image with less detail especially in the distance. This reduced acuity results from two sources: the objects and their backgrounds become more similar due to increased attenuation and scattering; and the atmosphere acts as a low-pass filter [10], suppressing the higher-frequency image components or details. We use the term local contrast to refer to image acuity and define it as the magnitude of difference in image intensity over a short spatial distance $C = |\nabla I|$ where the gradient is with respect to the two-dimensional image space. The same spatial difference can be computed on the right side of equation 2 to get

$$|\nabla I| = |\nabla(Jt + A(1 - t))| \quad (3)$$

$$= |\nabla Jt| \quad (4)$$

$$= t|\nabla J|. \quad (5)$$

Line 4 results from the assumption that the ambient light A is locally constant and line 5 results from the positivity of transmission t and the assumption that it too is locally constant. The quantity $|\nabla J|$ is the “true” contrast of the scene when imaged under perfect transmission. This equation shows transmission has the intuitive interpretation as the ratio of the observed contrast to the true contrast.

We use Sobel filters to estimate the image gradient at each pixel. To compensate for slight camera movement and other sources of image noise, we compute image contrast C as the average of the gradient magnitude over an image region Ω :

$$C = \frac{1}{|\Omega|} \sum_{\Omega} |\nabla I|. \quad (6)$$

Finally, putting it all together, we can relate the quantity we are trying to estimate, the coefficient of extinction b_{ext} , to what we measure from the image, contrast C (or, more

ID	Dim	Description
B64	432	Blocks 64x64 px
B128	108	Blocks 128x128 px
B192	48	Blocks 192 x 192 px
B256	24	Blocks 256 x 256 px
R64	24	Rows 64 px tall
R128	12	Rows 128 px tall
R192	8	Rows 192 px tall
All	224	All of the above, except B64

Table 1: The image partitionings considered and the dimensionality (Dim) of the resulting feature vectors (one contrast value per region). See the end of section 4 for more details on the partitionings which are identified by ID.

precisely, the log of the contrast) through the linear relation:

$$b_{ext} = \frac{\ln C}{r} - \frac{\ln |\nabla J|}{r}. \quad (7)$$

Here, $1/r$ is the scaling factor and $\ln |\nabla J|/r$ is the offset which we treat as unknowns in our predictive models below.

4. DATASET AND EVALUATION

We evaluate our prediction models using an extensive set of images and ground truth extinction readings from the Arizona Department of Environmental Quality which manages the PhoenixVis.net visibility web cameras website [2]. This website contains live images from six visibility cameras of scenic urban and rural vistas in the Phoenix, Arizona region. Our dataset consists of the following acquired over 2006:

- Digital images of South Mountain (SOMT) captured every 15 minutes.
- The extinction coefficient b_{ext} measured every hour using a transmissometer. This serves as the ground truth.

The images measure 1536×1152 pixels, are in the RGB colorspace, and (unfortunately) have been JPEG compressed at an unknown quality level. We only consider images taken at the top of each hour, since this is when the transmissometer readings are made, and during daylight hours, approximately 10 am to 4 pm. This results in a dataset of 6,897 images from the SOMT camera which are “labelled” with ground truth extinction b_{ext} values.

We evaluate the accuracy of our model using the coefficient of determination R^2 between the predicted and ground truth values. Let $b_{ext}'^i$ and b_{ext}^i be the predicted and true extinction coefficients for image i then

$$R^2 = 1 - \frac{\sum_{i=1}^n (b_{ext}'^i - b_{ext}^i)^2}{\sum_{i=1}^n (b_{ext}'^i - \bar{b}_{ext})^2} \quad (8)$$

where n is the number of images in the evaluation set and \bar{b}_{ext} is the mean of the true values. R^2 has a maximum value of 1 with higher values indicating a more accurate model.

In order to provide an intuitive feel for the predictions, we also compute the mean absolute error (MAE) between the

predicted and true values:

$$MAE = \frac{1}{n} \sum_{i=1}^n |b_{ext}'^i - b_{ext}^i| \quad (9)$$

In previous work [7], we demonstrated that local contrast computed using Sobel filters outperforms both contrast computed in the frequency domain and an image haze model based on the so-called dark channel prior [8] for estimating light extinction. We thus restrict ourselves here to local contrast based on Sobel filtering. The best results on the SOMT dataset were an R^2 value of 0.646 and an MAE value of 12.5 Mm^{-1} . Those results were for using a single image region; i.e., a univariate linear regression. In this paper, we extend that work to consider multiple image regions using regression trees as well as multivariate linear regression. The hypothesis and motivation is that no single image region will be effective for estimating light extinction under different visibility regimes. Far regions are not useful when visibility is relatively poor since they are not observable at all and close regions are not useful when visibility is relatively good since there is not enough intervening atmosphere to reduce visual acuity by a measurable amount.

We experiment with a number of image partitionings. First, we partition the images into square blocks of different sizes including: a 24×18 grid of 64×64 pixel blocks (B64); a 12×9 grid of 128×128 pixel blocks (B128); an 8×6 grid of 192×192 pixel blocks (B192); and a 6×4 grid of 256×256 pixel blocks (B256). We also consider regions corresponding to rows of these blocks. This results in three additional partitionings: a 24×1 grid of 64×1536 pixel regions (R64); a 12×1 grid of 128×1536 pixel regions (R128); and an 8×1 grid of 192×1536 pixel regions (R192). The motivation here is that each such row represents a large image region which is situated at approximately the same distance. These extended horizontal regions might thus be more effective than smaller blocks for estimating visibility. Yet, we still have a set at regions at various distances. Table 1 summarizes the different partitionings and also indicates the dimensionality of the feature vectors that result since we extract one contrast value per region using equation 6. The terms in parenthesis above are the IDs assigned to the different partitionings.

5. METHOD: REGRESSION TREES

Transmission is an inverse exponential function of both distance and the extinction coefficient. For areas with a low transmission value, there will be a minimal change in contrast for small shifts in the extinction coefficient. However, these same areas can become very informative as the visibility improves. Regression trees generate a decision-making process by which multiple linear fits can be considered within the same model. As the conditions of the image change, different image regions may be used to generate the output. By using a regression tree, the system can be trained to observe closer landmarks if distant points become overly attenuated.

We extend our earlier work by allowing for the use of all regions in the image and optimizing their contributions to the overall estimation. This problem is handled by segmenting the image into regions and using an M5’ regression tree [22, 29] to estimate visibility. This approach results in a discontinuous but piece-wise linear predictive model.

Regression trees are a type of decision tree learning method.




Figure 2: A mock-up example of a regression tree for a 2-feature visibility problem. Here, B_{ext} is expressed in terms of X_1 and X_2 . The number of training points that are used to learn each of the constituent parts of the model is indicated by the count N .

They are very similar to classification trees, but the primary goal is to output a real number instead of a class. The performance at each iteration of tree growth is measured by the information gain. This provides a quantitative value for the quality of each data split, and is calculated as the standard deviation of the examples assigned to the resulting leaf nodes. Intuitively, this means that we try to select splits that group similar regimes together. Regression trees select inequalities that split data in a way that maximizes this information gain. Each leaf node is then represented by a numerical value or a simple linear model to produce the output. An example tree using two features is shown in figure 2.

M5’ regression generates a decision tree using standard deviation to calculate the information gain. The feature with the largest information gain as measured in reduction of standard deviation, represents the best way to split the current examples into two distinct groups. Standard deviation reduction (SDR) is calculated as

$$SDR = sd(T) - \sum_i \frac{|T_i|}{|T|} \times sd(T_i) \quad (10)$$

where $sd(T)$ is the standard deviation of the parent node, $sd(T_i)$ is the standard deviation of each child node, $|T|$ is the number of data points in the parent node, and $|T_i|$ the number of data points in the child nodes. This splits the training data in a way that minimizes the standard deviation of each subset proportional to the size of those sets. Because variables contain continuous values, we perform the fit based on inequalities (e.g. $X_2 > 4.5$). The function stops splitting when the SDR becomes sufficiently low or the number of examples in the set becomes too small to continue splitting. The algorithm then begins pruning the tree for reduced error. This is necessary because the building process relies solely on SDR and does not test fits as it progresses. If the

Partitioning	Regression Trees		Multi. Linear Reg.	
	R^2	MAE	R^2	MAE
All	0.780	8.90	0.845	7.95
B64	0.776	8.88	0.830	8.21
B128	0.762	9.22	0.813	8.94
B192	0.735	9.82	0.830	8.21
R64	0.733	9.72	0.748	10.5
R128	0.722	10.3	0.729	11.1
R192	0.698	10.5	0.719	11.1

Table 2: Results of the regression tree and multivariate linear regression predictive models for different image partitionings. See table 1 for details on the partitionings. MAE is in units of Mm^{-1} .

summed error of child nodes is greater than a parent node, the parent node will become a leaf node instead. Each leaf node consists of a set of data points which can be characterized by a single value or a linear model. We use the M5PrimeLab [1] implementation of M5’ regression trees.

We evaluate the regression trees using five-fold cross validation. The 6,897 labelled images are partitioned randomly into five approximately equal sized sets. A regression tree is then learned using four of these partitions—the training set—and then evaluated on the fifth partition—the test set. This is then repeated four times using different training/test sets and the final results are averaged. This is performed for each of the image (spatial) partitionings described in section 4.

Table 2 shows the results for the different image partitionings. It is immediately apparent that using regression trees to incorporate multiple regions significantly improves upon our previous approach [7] which used only a single region and resulted in an R^2 value of 0.646 and an MAE value of $12.5 Mm^{-1}$. We will discuss these results further below.

6. METHOD: MULTIVARIATE LINEAR REGRESSION

In addition to the regression tree approach above, we also consider incorporating multiple image regions through multivariate linear regression. Given n image regions, we have n instances of equation 7. Assuming a constant value of the light extinction b_{ext} for the scene (i.e., the atmosphere is approximately uniform), we can sum these instances to get

$$nb_{ext} = \sum_{i=1}^n \frac{\ln C_i}{r_i} - \frac{\ln |\nabla J_i|}{r_i} \quad (11)$$

where C_i , r_i , and J_i are the measured contrast, distance, and true contrast of region i respectively. We can rewrite this as

$$b_{ext} = \sum_{i=1}^n \frac{\ln C_i}{nr_i} - K \quad (12)$$

where the individual offsets have been grouped into a single offset K . We can consider this as a single multivariate linear regression model in which the (logs of the) contrasts of the regions C_i are the variables, the $1/nr_i$ are the scaling factors, and K is the offset. Given a training set, we can then use least-squares to estimate the scaling factors and offset.

We evaluated this model using the same five-fold cross validation and image (spatial) partitionings as for the re-

gression tree model. The results for the multivariate linear regression model are summarized in table 2. This model is shown to outperform the regression tree model.

7. DISCUSSION

We make the following observations based on the results in table 2. First, not surprisingly, models which incorporate multiple image regions are more effective than those with only a single region. The best results for the regression tree and multivariate linear regression models are $R^2 = 0.780$ (MAE=8.90 Mm^{-1}) and $R^2 = 0.845$ (MAE=7.95 Mm^{-1}) respectively while the best result using a single region in our previous work was $R^2 = 0.646$ (MAE=12.5 Mm^{-1}). This represents improvements of 20.7% and 30.8% in terms of R^2 . This improvement is likely in large part due to the fact that scene regions at different distances are effective for estimating light extinction under different visibility regimes. Far regions are not useful when visibility is relatively poor since they are not observable at all and close regions are not useful when visibility is relatively good since there is not enough intervening atmosphere to reduce visual acuity by a measurable amount.

We also observe that models incorporating a larger number of smaller regions are better than those incorporating a smaller number of larger regions. This is true for both the regression tree and multivariate linear regression models. For the most part, more smaller blocks performed better than fewer larger blocks and more small rows performed better than fewer larger rows. Using the smaller blocks also performed better than the larger rows. This finding is further evidence that it is useful to have regions at varying distances so as to better model different visibility regimes.

The multivariate linear regression model is shown to outperform the regression tree model for all image partitionings. While it is difficult to make any general conclusions based on this result since these are two very different approaches, it is likely that the regression trees are over-fitting the training data. Many of the leaf nodes of the trees contain very few samples—some as few as four examples—which makes it difficult to fit even a simple linear univariate model. It is possible that more aggressive pruning of the tree could improve this. This is potentially a topic for further investigation.

Finally, it is instructive to examine how the prediction error varies with visibility regime. Figure 3 plots the error of the multivariate linear regression model for different ground truth values of b_{ext} . These are for the “All” image partitioning. It is clear that the model performs better for improved visibility regimes. This decrease in performance in poor visibility is most likely due to two factors: 1) the dataset contains fewer samples in this regime; and 2) the inverse exponential relationship between extinction and transmission results in reduced sensitivity of the model to changes in image contrast in this regime.

8. CONCLUSION

We proposed and evaluated two predictive models that incorporate multiple image regions to estimate light extinction. Our results show that models which incorporate multiple regions are more effective than those with a single region. They also show that using a larger number of smaller regions is better than smaller number of larger regions. These findings are informative for the goal of using the growing number

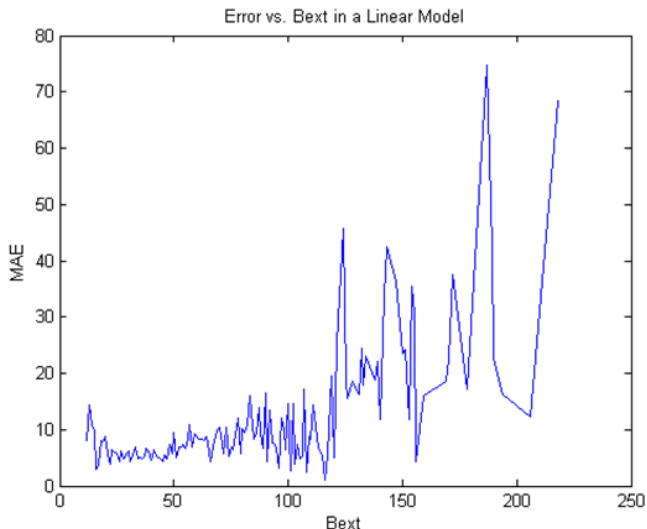


Figure 3: By plotting MAE (the mean absolute error) versus b_{ext} (the coefficient of extinction), we note that the model is more accurate for the improved visibility regime. These results are for the multivariate linear regression model using the “All” image partitioning.

of cameras in our ecosystem to estimate light extinction and possibly monitor atmospheric pollution.

While the predictive models learned here are specific to the evaluation dataset in terms of the regression coefficients, we expect the general findings—multiple regions at varying distances are better than a single region and multivariate linear regression performing better than regression trees—to be applicable to other cameras and scenes. The long-term challenging goal of this work is to learn predictive models for visibility camera systems for which we do not have access to ground truth data for calibration such as from a transmissometer. Such systems would, of course, be limited to providing relative and not absolute estimates of visibility. This would, however, allow the predictive models to be deployed to the growing number of cameras already present in our ecosystem such as web, surveillance, and traffic cameras.

9. ACKNOWLEDGEMENTS

This work was partially funded by the Center for Information Technology Research in the Interest of Society (CITRIS). We would like to thank the Arizona Department of Environmental Quality and Air Resource Specialists, Inc., for providing the images and transmissometer data.

10. REFERENCES

- [1] M5primelab: M5’ regression tree and model tree toolbox for matlab/octave. <http://www.cs.rtu.lv/jekabsons/>.
- [2] Phoenix visibility web camera website managed by the Arizona Department of Environmental Quality (ADEQ). <http://www.phoenixvis.net/>.
- [3] D. Baumer, S. Versick, and B. Vogel. Determination of the visibility using a digital panorama camera. *Atmospheric Environment*, 42(11):2593–2602, 2008.

- [4] J. Betts. The instrumental assessment of visual range. *Proceedings of the IEEE*, 59(9):1370–1371, September 1971.
- [5] F. Caimi, D. Kocak, and J. Justak. Remote visibility measurement technique using object plane data from digital image sensors. In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*, pages 3288–3291, 2004.
- [6] R. Fattal. Single image dehazing. In *ACM SIGGRAPH*, pages 72:1–72:9, 2008.
- [7] N. Graves and S. Newsam. Using visibility cameras to estimate atmospheric light extinction. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 577–584, 2011.
- [8] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2341–2353, 2009.
- [9] K. Kim and Y. Kim. Perceived visibility measurement using the HSI color different method. *Journal of the Korean Physical Society*, 46(5):1243–1250, 2005.
- [10] V. Krishnakumar and P. Venkatakrishnan. Determination of the atmospheric point spread function by a parameter search. *Astronomy & Astrophysics Supplement Series*, 126:177–181, November 1997.
- [11] P. Lee, T. Hoffer, D. Schorran, E. Ellis, and J. Moyer. Laser transmissometer—a description. *Science of The Total Environment*, 23:321–335, 1982.
- [12] C.-H. Luo, C.-Y. Wen, C.-S. Yuan, J.-J. Liaw, C.-C. Lo, and S.-H. Chiu. Investigation of urban atmospheric visibility by high-frequency extraction: Model development and field test. *Atmospheric Environment*, 39(14):2545–2552, 2005.
- [13] J. V. Molenar, D. S. Cismoski, F. Schreiner, and W. C. Malm. Analysis of digital images from Grand Canyon, Great Smoky Mountains, and Fort Collins, Colorado. In *Regional and Global Perspectives on Haze: Causes, Consequences and Controversies Visibility Specialty Conference*, 2004.
- [14] E. Namer and Y. Y. Schechner. Advanced visibility improvement based on polarization filtered images. In *Proc. SPIE 5888: Polarization Science and Remote Sensing II*, pages 36–45, 2005.
- [15] E. Namer, S. Schwartz, and Y. Y. Schechner. Skyless polarimetric calibration and visibility enhancement. *Optics Express*, 17(2):472–493, 2009.
- [16] S. Narasimhan and S. Nayar. Chromatic framework for vision in bad weather. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 598–605, 2000.
- [17] S. Narasimhan and S. Nayar. Removing weather effects from monochrome images. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 186–193, 2001.
- [18] S. Narasimhan and S. Nayar. Vision and the atmosphere. *International Journal on Computer Vision*, 48(3):233–254, July 2002.
- [19] S. Narasimhan and S. Nayar. Contrast restoration of weather degraded images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(6):713–724, June 2003.
- [20] S. Narasimhan and S. Nayar. Interactive (de)weathering of an image using physical models. In *ICCV Workshop on Color and Photometric Methods in Computer Vision*, 2003.
- [21] S. Narasimhan and S. Nayar. Shedding light on the weather. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 665–672, 2003.
- [22] J. Quinlan. Learning with continuous classes. In *Australian Joint Conference on Artificial Intelligence*, pages 343–348, 1992.
- [23] D. S. Raina, N. J. Parks, W.-W. Li, R. W. Gray, and S. L. Dattner. Innovative monitoring of visibility using digital imaging technology in an arid urban environment. In *Regional and Global Perspectives on Haze: Causes, Consequences and Controversies Visibility Specialty Conference*, 2004.
- [24] Y. Schechner, S. Narasimhan, and S. Nayar. Instant dehazing of images using polarization. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 325–332, 2001.
- [25] Y. Schechner, S. Narasimhan, and S. Nayar. Polarization-based vision through haze. *Applied Optics, Special issue*, 42(3):511–525, January 2003.
- [26] J. Seinfeld and S. Pandis. *Atmospheric Chemistry and Physics: From Air Pollution to Climate Change*. Wiley, 2006.
- [27] S. Schwartz, E. Namer, and Y. Y. Schechner. Blind haze separation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1984–1991, 2006.
- [28] R. Tan. Visibility in bad weather from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [29] Y. Wang and I. Witten. Induction of model trees for predicting continuous classes. In *European Conference on Machine Learning*, pages 128–137, 1997.