# The Global Network of Outdoor Webcams: Properties and Applications

Nathan Jacobs[1,4], Walker Burgin[1], Nick Fridrich[1], Austin Abrams[1], Kylia Miskell[1], Bobby H. Braswell[2], Andrew D. Richardson[3], Robert Pless[1]
[1]Department of Computer Science and Engineering, Washington University, St. Louis, MO, 63119
[2]Atmospheric and Environmental Research, Inc., Lexington, MA, 02421
[3]Dept. of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA, 02138
[4]jacobsn@cse.wustl.edu

## ABSTRACT

There are thousands of outdoor webcams which offer live images freely over the Internet. We report on methods for discovering and organizing this already existing and massively distributed global sensor, and argue that it provides an interesting alternative to satellite imagery for global-scale remote sensing applications. In particular, we characterize the live imaging capabilities that are freely available as of the summer of 2009 in terms of the spatial distribution of the cameras, their update rate, and characteristics of the scene in view. We offer algorithms that exploit the fact that webcams are typically static to simplify the tasks of inferring relevant environmental and weather variables directly from image data. Finally, we show that organizing and exploiting the large, ad-hoc, set of cameras attached to the web can dramatically increase the data available for studying particular problems in phenology.

## Categories and Subject Descriptors

I.4.9 [**Computing Methodologies**]: IMAGE PROCESSING AND COMPUTER VISION—*Applications*

## General Terms

Algorithms, Measurement, Performance

## Keywords

webcam, phenology, camera network, global, calibration, metadata

## 1. INTRODUCTION

Many thousands of outdoor cameras are currently connected to the Internet—they are placed by governments, companies, conservation societies, national parks, universities, and private citizens. Individually, these cameras observe scenes in order to show the current traffic and weather conditions, to advertise the beauty of a particular beach or mountain, or to give a view of animal or plant life at a particular location. Collectively, however, this set of cameras has an untapped potential to monitor global trends—changes in weather, snow-cover, vegetation, and traffic density are all observable in some of these cameras, and together they give a relatively dense sampling over the US, Europe and parts of Asia. Figure 1 shows locations of the more than 16000 webcams we have discovered that are currently giving live images.

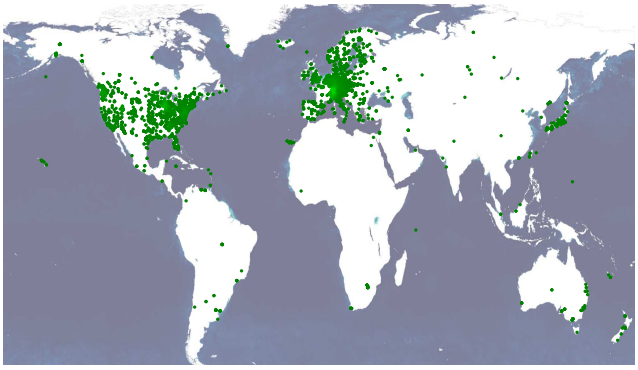

**Figure 1: Tens of thousands of outdoor webcams, cameras attached to the Internet, are currently delivering live images from across the planet (camera locations shown above). We argue that these images represent an underutilized resource for monitoring the natural world. In this work, we describe our efforts to discover, organize, and use these cameras to measure several environmental properties.**

While cameras are used as important sensors for a wide variety of problems including measuring plant growth, surveying animal populations, monitoring surf conditions, and security, often there are limitations due to the cost of installing and maintaining these cameras, especially in biology and sociological research. Because the cameras discovered within the global camera network are already installed and being maintained, the marginal, additional cost of including images from those cameras in a novel study lies in (a) discovering cameras, (b) archiving camera images, (c) selecting images for particular tasks, and (d) developing algorithms specific to a task.

The contributions of this paper are tools to address each of


*ACM GIS '09, November 4-6, 2009* Seattle, WA, USA

these problems, either by solving them, or by offering ways to simplify them on a task-specific basis. Together, they minimize the time/cost to make use of this webcam network. First, we offer a characterization of the set of live images currently available through a web URL—both our methods of discovering the list of available webcams and a statistical sampling of this set of images to determine where those cameras are located and properties of the scene in view. Second, we describe a set of tools that we have created to visualize archives of images from 1000 of these webcams that we have been capturing over the last three years. Third, we use this data archive to demonstrate algorithms that automatically infer weather conditions strictly from image data.

We will conclude by considering one problem in greater depth. Phenology is the study of periodic plant and animal life cycle events. For plant studies, ground cameras are an appealing alternative to satellite-based approaches (e.g. MODIS platform [35]) because satellite imagery is corrupted by aerosols and water vapor, the effect of changes in viewing angle on the retrieved signal, and the relatively coarse spatial (approx. 500 meters) and temporal resolution of satellite data products. This problem is important enough that explicit camera monitoring networks have been deployed on a small scale [26]. In this paper we demonstrate that with relative simple image processing, many cameras already online (e.g. traffic cameras and campus "quad-cams"), already support this analysis. This demonstrates one example where properly organizing the global webcam network supports large scale environmental monitoring studies with limited additional cost.

Our main argument is that the global webcam network is a dramatically under-utilized resource. In Section 2 we describe our work in discovering, understanding, and organizing this resource. Section 3 demonstrates methods to visualize long-term time-lapse data, and highlight the benefits of camera geo-location in scene understanding. Finally, Section 4 highlights that the global webcam network is already able to augment or replace dedicated camera networks for many important environmental monitoring tasks over large spatial scales.

## 1.1 Related Work

To our knowledge, ours is the first work that attempts to systematically catalog, characterize, and use the set of all publicly available webcams. There have been limited attempts to automatically calibrate static webcams, and several projects which create dedicated camera networks to monitor particular environmental features.

### *Static cameras.*

Within the Computer Vision community, two data-sets have grounded most work on understanding images from static cameras. First, the "Weather and Illumination Database" (WILD), captured images from a single, well calibrated and controlled camera. This dataset includes simultaneous weather measurements and was captured over the course of a year in a dense and urban part of New York City [20]. The "Archive of Many Outdoor Scenes" (AMOS) extended this to hundreds of webcam images in a broader collection of settings, and has gathering images since March 2006 [11].

Many algorithms have been developed to infer scene and camera information using long sequences of images from a fixed view. These new tools have the potential to overcome the challenges of interpreting images from webcams by providing accurate sensor calibration. Examples include a methods for clustering pixels based on the surface orientation [16], for factoring a scene into components based on illumination properties [28], for obtaining the camera orientation and location [13, 12, 29, 17], and for automatically estimating the time-varying camera response function [15].

### *Camera Networks for Environmental Monitoring.*

There is a long history of using camera networks to monitor environmental changes and social behaviors. Notable examples which use large dedicated camera networks include the Argus Imaging System [5] with 30 locations and 120 camera which explicitly focuses on coastal monitoring. Cameras within the Argus network, and similar cameras set up on an ad-hoc basis for individual experiments have been used to quantify density of use of beach space [18], the use of beaches as a function of weather [14], and trends both in beach usage and beach erosion [9]. Another large, dedicated camera network is the Haze Cam Pollution Visibility Camera Network [4]. In this case, the cameras are placed near measurement systems for air pollution and other meteorological data, but the images are primarily used to provide the public a visual awareness of the effects of air pollution. To our knowledge, these cameras have not been systematically used to provide additional quantitative measurements to augment the explicit pollution or meteorological data, but recent work has validated that similar cameras have high correlation with explicit measurements of atmospheric visibility, based both on ground [34], and satellite measurements [33, 32].

Additional work has focused on phenology and recent studies have shown that phenology is a robust indicator of climate change effects on natural systems; for example, earlier budburst and flowering by plants have been documented in response to recent warming trends. Improved monitoring of vegetation phenology is viewed as an important, yet simple, means of documenting biological responses to a changing world [23]. New and inexpensive monitoring technologies are resulting in a dramatic shift in the way that phenological data are now being collected [19], and already several networks based around color digital imagery (e.g. "PhenoCam" [24] and the "Phenological Eyes Network" [25]) have been established to monitor phenology at a regional scale. Previous studies have provided solid evidence that both qualitative and quantitative information about seasonal changes in the condition and state of vegetation canopies can be extracted from webcam images [27, 26]. Bradley et al. demonstrate [1] that some of the necessary processing can be performed using a web-based interface.

Recent work by Graham et al. demonstrates, similar to the example application in Section 4.2, that a large collection of webcams can be used to estimate temporal properties of leaf growth on trees. Their work supports our claim of the importance of careful sensor calibration to handle, for example, automatic color balance compensation.

### *Discovering webcams.*

Significant efforts to collect and annotate large numbers of webcams has been undertaken. Notably, Google Maps now includes a "webcam layer" that organizes live webcam feeds from approximately 6700 webcams. Other large collections of webcam URLs [31, 30, 22] have been created and

many of these cameras are geo-located. However, these collections are not as spatially dense as the cameras we have discovered (see Section 2), and to our knowledge are not yet being used to explicitly infer geographic information or for environmental monitoring.

## 2. CAMERA DISCOVERY AND CHARACTERIZATION

Our overall goal is to create a comprehensive list of URLs that point to live images captured by a webcam and then use images from them to measure environmental properties.

Our strategy for finding URLs involves merging lists from webcam aggregators and explicit searches for cameras in, for example, national parks, state departments of transportation, and similar searches targeting each country around the world. Many of the cameras we have discovered come from web sites that contain lists of large numbers of webcams, either cameras that they explicitly own (e.g. the Weatherbug camera network [31]), or cameras that individuals register to be part of a collective (e.g. the Weather Underground webcam registry [30]). Additionally, we use a collection of Google searches for unusual terms, such as "inurl:axis-cgi/jpg", that are primarily used in default webpages generated by webcams.

To date, we have found 16112 webcam URLs (including the URLs from the AMOS dataset [11]) that give different live images, and these URLs correspond to images with different sizes, refresh rates, and scenes. The remainder of this section provides an overview of the properties of these webcams. We consider both low-level properties of the images, such as file size, and high-level properties of the scene, such as whether or not a mountain is visible.

### 2.1 Webcam Image Properties

We begin by describing low-level properties of individual webcam image files and the sequence of images generated by a webcam. Figure 2 shows the distribution of file sizes and image dimensions that reflects the fact that most webcams provide small, highly-compressed images. In order to understand the distribution of temporal refresh rates, we estimate the refresh rate for a set of 500 randomly selected cameras using a standard method [3]. The distribution in Figure 2(c) reflects the fact that many webcams are configured to capture new images every 5 minutes.

To begin to characterize statistics of the scenes viewed by this set of cameras, we manually estimated the minimum and maximum distance of objects in the scene from the camera for a randomly chosen subset of 300 cameras. We grouped our estimates into the following intervals: 1–10 meters, 10–100 meters, 100–1000 meters, and greater than 1000 meters. Most cameras have objects both near and far; this is highlighted in Figure 2(d) where the cumulative distribution functions for min- and max-depth show that 80% of the scenes have an object within 10 meters of the camera, and only 20% of the scenes do not contain an object more than 100 meters away.

Additionally, we manually labeled the scenes imaged by all the cameras in the AMOS dataset and 300 randomly sampled new cameras. We tagged each scene based on several characteristics: if it was outdoors, if it contained a road, trees, buildings, or substantial sky, or water (where we define 'substantial' to mean 'at least a fifth of the picture').

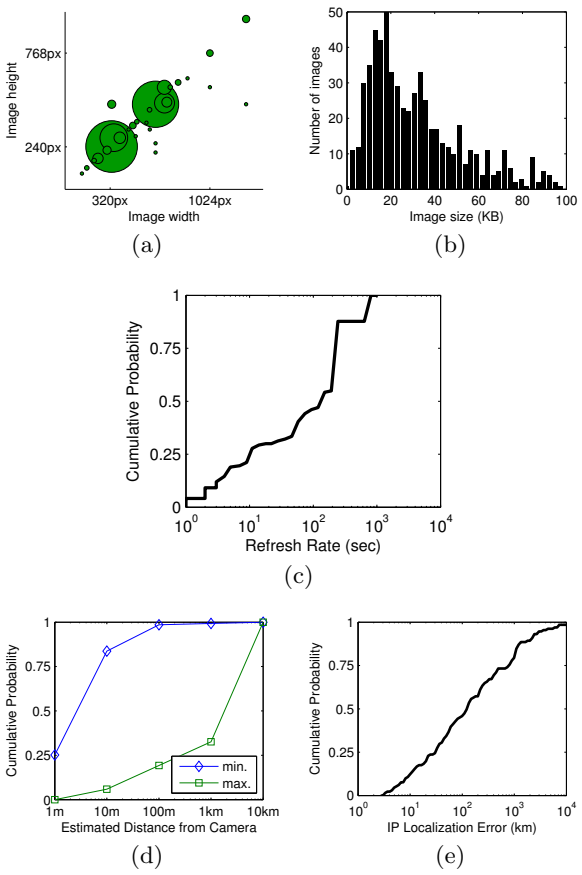


**Figure 2: (a) The distribution of image sizes measured in pixels. Each circle is centered at an image size with area proportional to the number of cameras. (b) The distribution of image sizes in kilobytes. (c) The cumulative distribution of refresh rates of webcams. Note the large number of cameras that refresh every 5 minutes. (d) The cumulative density function of the minimum depth of an object in the scene (blue line, near the top) and the maximum depth of an object in the scene (green line, near the bottom). Most cameras see objects both in the near field (10m or less) and far field (at least 1 km). (e) The cumulative distribution of localization errors of IP-addressed based localization. The median error is 111km.**

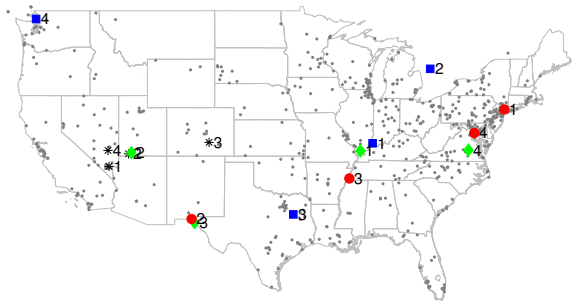

(a)

(b) contains a mountain (black-star)

(c) contains natural water and a building (blue-square)

(d) contains a building but no trees (red-circle)

(e) contains trees but no building (green-diamond)

| Tags | Percentage |
|---|---|
| Outside | 95 |
| Road | 50 |
| Trees | 67 |
| Buildings | 60 |
| Sky | 68 |
| Water | 22 |

(f)

**Figure 3: (a) A subset of webcams in AMOS dataset, and (b) the percentage of those cameras that have particular tags (c–e) Four random sample images from subsets defined by the presence (absence) of a manual metadata tag. The locations of each of these images is presented on the map in part (a) of this figure. (f) Global statistics of the labels.**

Figure 3 shows specific examples of this labeling and gives some global statistics. This type of manual labeling is especially helpful for selecting a subset of cameras for a particular measurement task (see Section 3).

## 2.2 Camera Locations

To estimate the spatial distribution of webcams, shown in Figure 1, we used the IPInfoDB geo-location database [8] to obtain an approximate camera location based only on the IP-address of the camera for a randomly selected subset of 200 camera. We then manually localized these cameras resulting in a set of 138 cameras for which we were confident that our location estimate was within 100 meters of the true location. Comparison of the IP-based estimates with manually generated ground-truth camera locations, see Figure 2(e), shows that there is significant error in the IP-address based location estimates. In fact, half of the cameras have a localization error greater than 111km. This motivates the use of techniques that use images to estimate camera location [13, 29]. These methods have shown significantly lower error rates and work by reasoning about image variations caused by the motion of the sun and correlations in local weather patterns.

Obtaining accurate location estimates for the cameras is critical for our goal of measuring environmental properties. Therefore, we use a combination of automatic [13] and manual techniques (using webpage metadata and visual landmarks) to determine locations for the cameras we have discovered. The next section describes an additional benefit of having accurate geo-location information.

## 2.3 Automatic Scene Labeling

One advantage of having an accurate location estimate for a camera is that it facilitates integration with existing GIS databases. This enables, for example, the rich annotations present in these databases to be transferred to the cameras. These annotations can determine, for example, if a camera is in a dense urban area or farm land, or if it is more likely viewing a road or a river.

In addition to labeling cameras, individual images can be automatically labeled. To demonstrate this we spatially and temporally registered sequences of webcam images to historical weather readings from the National Climatic Data Center [21]. Figure 4 shows the type of filtering possible when local weather readings are registered to webcam images. A side benefit of this labeling is that it provides an interesting additional form of context for Computer Vision algorithms, this is an area we are leaving for future work.

We continue to work to expand our list of live webcam URLs and integrate our images with a variety of sources of annotation. This is done in the context of our overall goal to use webcams to measure specific environmental properties. For a given measurement task the first step in using the global webcam network is likely selecting a set of cameras that view suitable scenes or images that have a certain property (e.g. with low haze or with no wind). The next section describes our work that eases this preliminary step.

# 3. BROWSING WEBCAMS

In working with a large archive of images from many webcams, we find that visualization tools are critical for debugging, updating, and maintaining the capture system, as well as for finding relevant images for particular tasks. Currently,

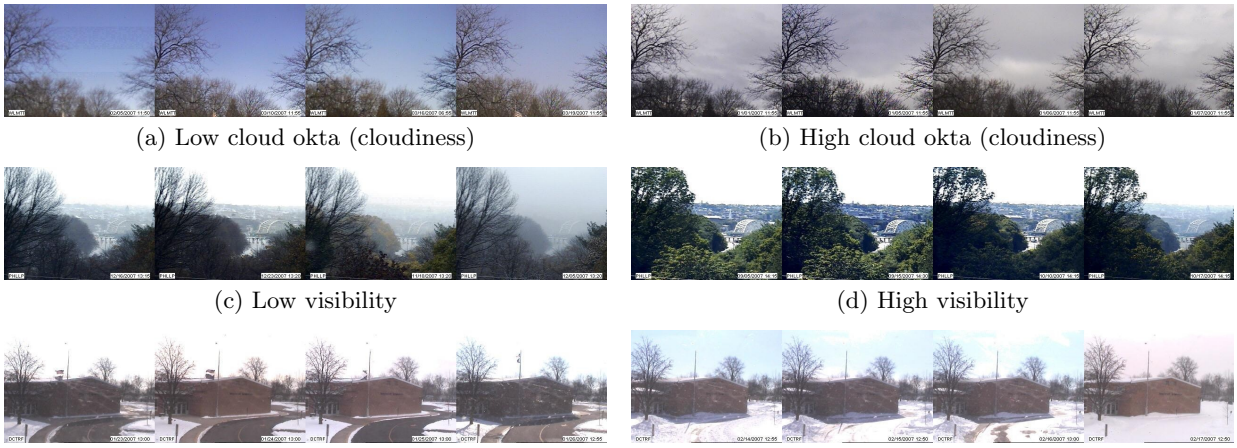

(a) Low cloud okta (cloudiness)

(b) High cloud okta (cloudiness)

(c) Low visibility

(d) High visibility

(e) Low snow depth

(f) High snow depth

**Figure 4: Automatic image labels can be created by spatially and temporally registering webcam images to weather reports. Above are montages of webcam images that correspond to extreme weather readings for a variety of weather properties.**

we use a two-layer web interface; first, a page shows the current image of every camera in the data set, or a subset based on keyword/tag filtering as shown in Figure 3. This filtering and browsing interface is important for determining how many cameras may support a particular task, and for determining whether cameras are currently broken or not delivering images (which is common since we have no control over the cameras, and they are supported and maintained with varying degrees of attention). Second, each camera has a dedicated page which includes all known meta-information (e.g. tags, geo-location, and, if known, calibration and geo-orientation), as well as an interface that supports searching for a specific image from a camera.

Searching for a specific image from a camera is done using a summary of the image appearance over the course of each year. The first instantiation of this yearly summary is an image indexed by time of year (on the x-axis) and time-of-day (on the y-axis). At each pixel, this image shows the mean color of the entire image captured at that time on that day. Figure 5(a) shows two example cameras and the annual summaries of those cameras for 2008. This interface makes it very easy to see when day (gray), night (black), and missing images (dark red) occur at a camera. For example, the left image shows the nighttime growing shorter during the middle of the summer. The right side of Figure 5(a) shows the unusual circumstance of a camera for which the night-time seems to drift throughout the year; this camera is mounted on the bridge of a Princess Cruise Lines ship which circumnavigated the globe in 2008.

The web interface allows the user to click a pixel on this summary image, and shows the image taken on that day, closest to the selected time. This gives an intuitive way to view, for example, a large set of images taken near dawn, by selectively clicking along the day-night interface. Additionally, keyboard interfaces allow moving to the image taken at the same time the previous day, or moving forward and backward within a day, to give time-lapse movies at different time resolutions.

However, the two summary visualizations shown immediately below the images in Figure 5(a) are less informative than one would like. Since many cameras perform both contrast equalization and color balancing in order to give reasonable pictures at all times of the day, this summary image often shows little more than a clear indication of when daytime is, and other changes such as shifts in camera viewpoint or changes in scene color may not be visible.

A more abstract but informative visualization can be achieved by performing principle component analysis (PCA) on the set of images, which we compute incrementally using Brand's [2] algorithm. The set of images from a camera $\{I_1, I_2, \ldots, I_k\}$ is approximated as the linear combination of the mean image $\mu$ and of three basis images $\{b_1, b_2, b_3\}$ so that for each image $i$, $I_i \approx \mu + \alpha_{i,1} b_1 + \alpha_{i,2} b_2 + \alpha_{i,1} b_3$. The vector of coefficients $(\alpha_{i,1}, \alpha_{i,2}, \alpha_{i,3})$ gives a compact description of image $i$ relative to the overall variation in images seen at that camera. We normalize these coefficients and use them to define the RGB channel of a (false-color) summary image.

This summary visualization make it simple to find if cameras have captured data at relevant times of year; if they have moved, and allow rapid navigation through the large image dataset. The bottom of each part of Figure 5 shows this visualization for three different cameras. In particular, this highlights the consistency of the daily variations of the desert scene, the inconsistency throughout the year of the view from the cruise ship, and the slight change in view point of the view of the golden rotunda.

# 4. MEASURING ENVIRONMENTAL PROPERTIES

Local environmental properties often directly affect the images we collect from the webcams; whether it is cloudy or sunny is visible by the presence of shadows; wind speed and direction is visible in smoke, flags, or close up views of trees; particulate density is reflected in haziness and the color spectrum during sunset. We explore techniques to automatically extract such environmental properties from long sequence of webcam images. This allows the webcams *already* installed across the earth to act as generic sensors to improve our understanding of local weather patterns and variations.

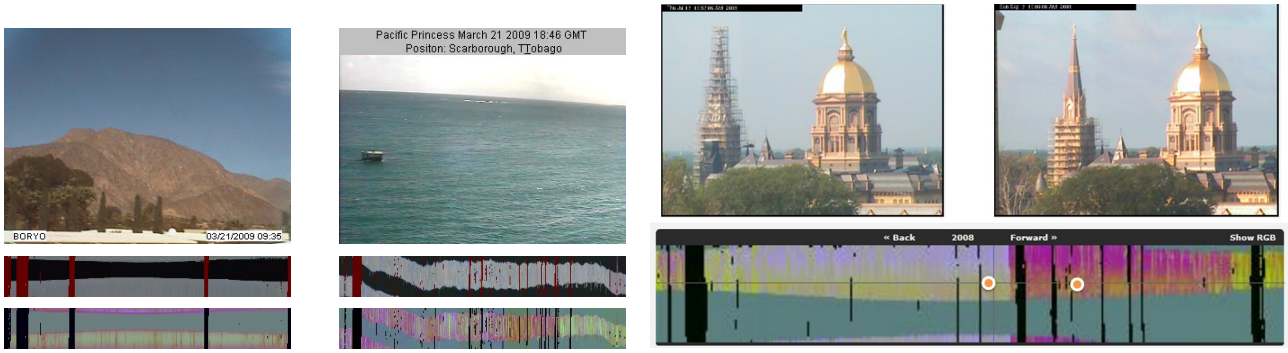


(a) (b)

**Figure 5: (a) Images from two cameras in our database with the corresponding RGB- and PCA-based annual summary images (x-axis is day-of-year and y-axis is time-of-day). The right camera is on a cruise ship that circumnavigated the globe during 2008; this causes nighttime to "wrap" during the year. (b) An example where the PCA-based summary image highlights a small change in the camera viewpoint; the dots on the summary image correspond to the images (above) which show a small viewpoint shift. The time when the shift occurs corresponds to the summary image changing from yellow/blue to purple.**

We first demonstrate that a generic supervised learning technique can automatically learn to estimate the relationship between a time-lapse of images and a time-varying weather signal (in this case, wind velocity) [10]. The supervised setting, while limited to situations in which a collocated sensor is available, demonstrates that extracting a variety of environmental properties is possible. A side benefit is that the models trained in this fashion often show interesting relationships to the calibration of the camera (i.e. in this case we find a relationship between a model for predicting wind velocity and the the geo-orientation of the camera).

Second, we consider another application more in depth, and show that minimal additional human intervention provides robust tools to quantify the timing and rate of the "spring onset" of leaf growth on trees. In many regions of the world, the timing of spring onset has advanced at between 2 and 5 days per decade over the last 30 years [23], and the length of the growing season is an important factor controlling primary productivity and hence carbon sequestration. Our analysis here expands on the ongoing efforts of the PhenoCam project [24, 26], in which a smaller number of dedicated, high-resolution (1296 x 960 pixel) cameras were deployed specifically for this purpose at forest research sites in the Northeastern U.S. While there are additional challenges in working with a much larger set of cameras for which the camera settings and internal processing algorithms are unknown, results presented here show that the spring greenup signal is visible in many cameras not dedicated to this monitoring task.

## 4.1 Supervised Weather Signal Estimation

We consider a time series of images $I_1, I_2, \ldots I_n$ captured from a camera with a known geographic location, and a synchronized time series of wind velocity estimates $Y_1, Y_2, \ldots Y_n$ captured from a nearby weather station. Canonical correlation analysis [7] (CCA) is a tool for finding correlations between a pair of synchronized multi-variate time signals. Applied to this problem, it finds a projection vector $A$ and a projection vector $B$ that maximizes the correlation between the scalar values $AI_t$ and $BY_t$, over all time steps $t$. Then, given a new image $I_{t+1}$, we can predict the projected version of the wind velocity signal as: $BY_{t+1} \approx AI_{t+1}$. We find that both the $A$ and the $B$ matrices tell us interesting features about the scene in view.

Figure 6 shows results for one camera in the AMOS dataset. The image projection $A$ can be represented as an image, and clearly highlights that the orientation of a flag within the scene is highly correlated with the wind speed. The plot shows the first dimension that CCA predicts from both the webcam images and the weather data for our test data. In Figure 6(d) we show the relationship of the CCA projection vector and the geographic structure of the scene. We find that the wind velocity projection vector $B$ is perpendicular to the viewing direction of the camera.

The algorithm described above demonstrates that webcams can be used to extract environmental properties. However, the method is limited because it requires a collocated sensor to train the model. Generalizing the method to work on cameras without collocated training data is an important problem to making the global webcam imaging network useful for monitoring the environment. In the next section we show an example of a signal estimator that does not require a collocated sensor.

## 4.2 Spring Leaf Growth

This section describes efforts to estimate the timing of spring leaf development from webcam images. Importantly, this method does not require a co-located sensor or ground observations of vegetation phenology, and human input is minimal. We use the simple "relative greenness" signal [27] and show that it can be extended to many of the cameras in the AMOS dataset. The relative greenness, $g/(r+g+b)$, is defined as the average of the green color channel divided by the sum of all color channels.

We begin by selecting a set of cameras with a significant number of trees in the field of view. For each camera we extract a set of images (at most one for each day) captured around noon for the first 275 days of 2008. We manually draw a polygon around the trees (since the cameras are static, or registered post-capture, only one polygon must

(a)

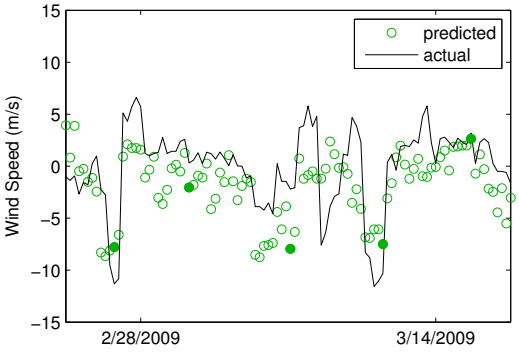


(b)

(c)


(d)

**Figure 6: An example of predicting wind speed from webcam images. (a) An example image and the CCA projection used to linearly predict the wind speed from a webcam image. (b) Predicted wind speed values and corresponding ground truth. (c) A montage in which each image corresponds to a filled marker in the plot above. (d) An image from Google Maps of the area surrounding the camera. The camera FOV was manually estimated by visually aligning scene elements with the satellite view. The dashed line (red) is the CCA projection axis defined as $B$ in Section 4.1. This image confirms that, as one would expect, our method is best able to predict wind direction when the wind approximately perpendicular to the principal axis of the camera.**

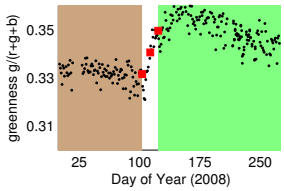


(a)

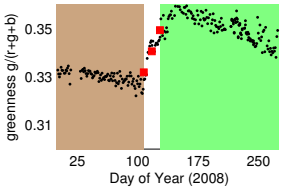


(b)

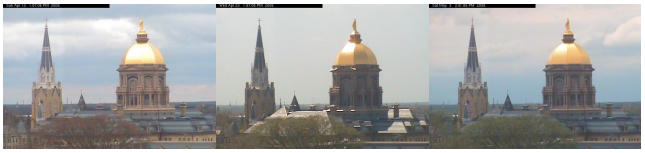

(c)

**Figure 7: Estimating spring leaf growth using an outdoor webcam. (a) The raw value of greenness over time (black dots) and a partitioning of the year based on the presence/absence of leaves. (b) Results for the same set of images with color correction, based on an ad-hoc color-standard, applied. The color correction reduces the local variance of the greenness score but, in this case, does not significantly impact the estimated onset of spring leaf growth. (c) Three images, each corresponding to a square (red) marker in (b), to verify the model fit.**

be drawn). We then compute the average greenness value of the tree region for each image. In order to characterize the timing of spring leaf growth, we fit a 4-parameter sigmoid model [27],

$$g(t) = a + \frac{b}{1 + \exp(c - dt)} \tag{1}$$

where $t$ is the day of year, to the greenness signal. Note that $c/d$ corresponds to the day-of-the-year of the vertical midpoint of the model.

Some cameras in the dataset automatically adjust the color balance to respond to changing illumination conditions (due, for example to clouds, solar elevation, and aerosols). This causes problems because the colors measured by the camera vary even when the underlying color of the scene does not change. To compensate for this automatic color balancing we use scene elements such as buildings or street signs (whose true color we assume to be constant over time) as an ad-hoc color standard. We then solve for the linear color axis scaling which maintains the color of the color standard, and apply this scaling to the entire image to create a color balanced image.

Figure 7 shows the raw and color-corrected greenness signals and the estimated sigmoidal model for a single camera. In addition the figure contains a montage of three images for manual inspection. The images in the montage are selected by first determining the vertical mid-point, $\hat{t}$, of the sigmoid function. The images selected for the montage are the images closest to $\hat{t} - 10$ days, $\hat{t}$, and $\hat{t} + 10$ days. More results, as well as a color-coded map, are shown in Figure 8. The map shows, as expected, a slight linear correlation between latitude and the "spring onset" [6]. The highlighted montages show that the estimated dates are accurate.

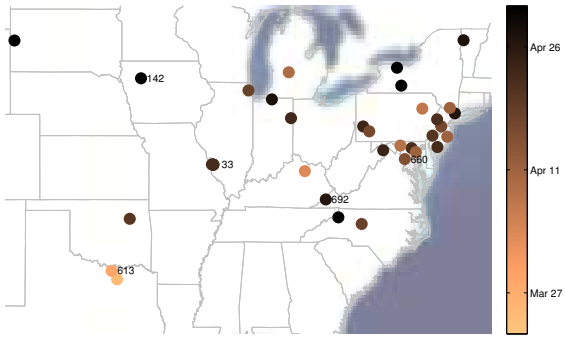


(a)

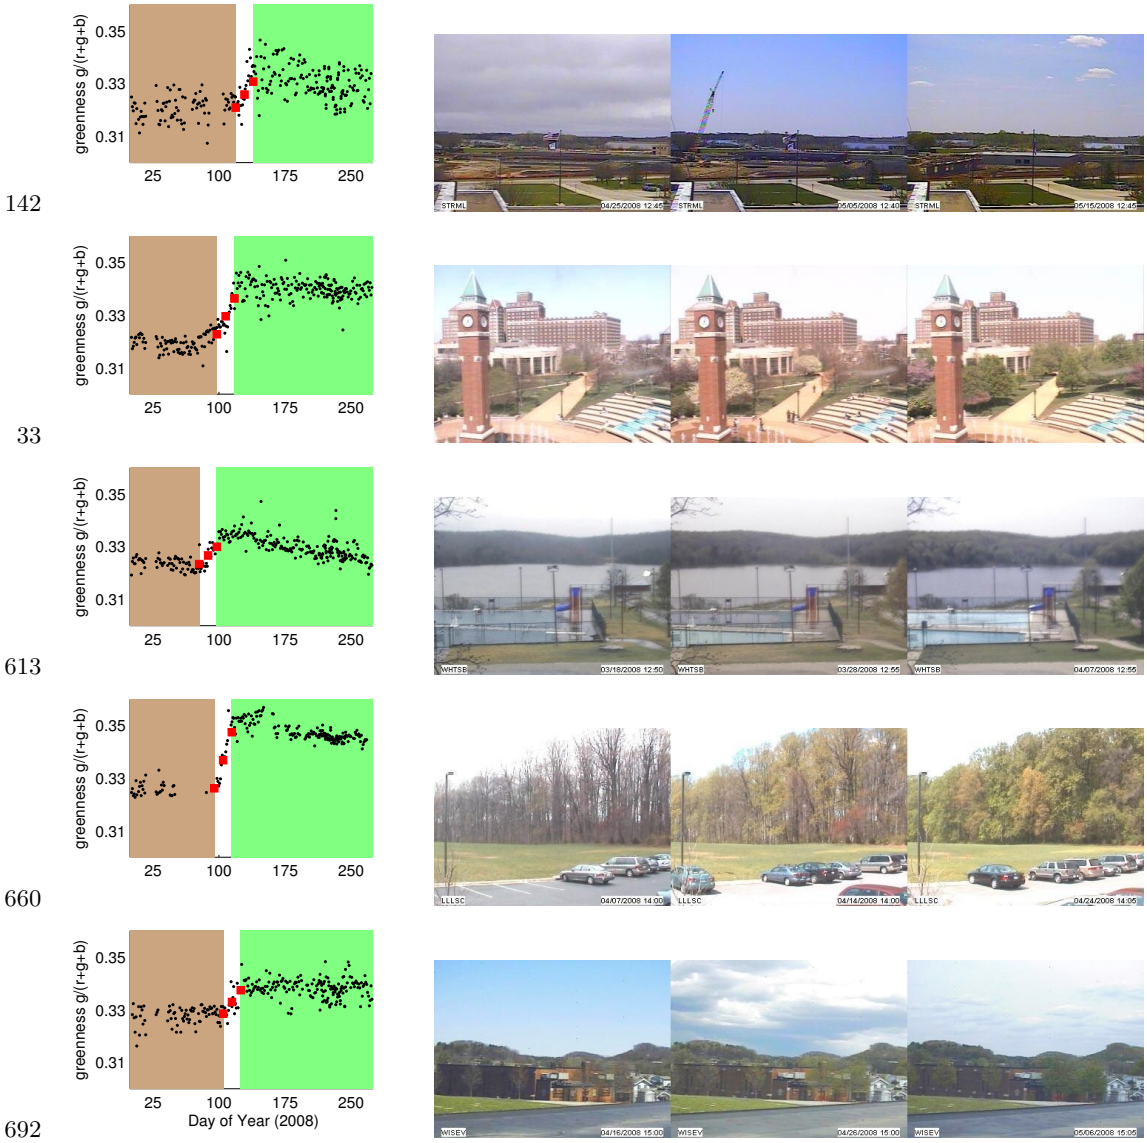


(b)

**Figure 8: Determining the onset of spring leaf growth using webcams. (a) A scatter plot of the locations of webcams used in the experiment (points colors correspond to the midpoint of spring leaf growth and are determined by a sigmoidal model of greenness). (b) (left column) The greenness signal of the webcams and the corresponding leaf-growth transitions determined by the sigmoidal model. (right column) The images that correspond to the square (red) markers in the plot in the left column. Careful inspection reveals that our model correctly finds the transition between no-leaves and leaves.**

In some webcam images temporal variations in the average greenness signal are caused by multiple species of trees. This same problem occurs in satellite imagery but, unlike when using satellite imagery, webcams allow us to address the problem. In fact it is possible to factor the average greenness signal into components due to multiple tree species.

Our approach is to first fit a mixture-of-sigmoids model,

$$g(t) = a + \frac{b}{1+\exp(c-dt)} + \frac{e}{1+\exp(f-gt)},$$

to the greenness signal (we use Levenberg-Marquardt to fit the model). Figure 9 shows the result of fitting this model to the average greenness signal from a camera that views multiple tree species. The time-series shows that the new function is a more accurate model of the data (i.e. the extra sigmoid allows the model to fit the small rise that occurs roughly 20 days before the main rise).

The coefficients of mixture-of-sigmoids model helps to segment the image into regions that correspond to the individual mixture components. To obtain the segmentation, shown in Figure 9, we first fit two single-sigmoid models, one for each component in the mixture model, separately to each pixels greenness signal. Each new model has the same form as (1) except two parameters, $c$ and $d$, are held fixed to the values from the corresponding mixture component (these correspond to the horizontal shift and stretch of the sigmoid). For each pixel, the model with the lowest mean-squared error is chosen as the correct model and the pixel is labeled accordingly. This segmentation approximately breaks the scene into the two types of trees in the field of view.

These results offer exciting possibilities for low-cost automated monitoring of vegetation phenology around the world. There are numerous potential applications of the resulting data streams [19], including real-time phenological forecasting to improve natural resource management (particularly agriculture and forestry) and human health (e.g. the dispersal of allergenic pollen) as well as validation and improvement of algorithms for extracting phenological information from satellite remote sensing data.

# 5. CONCLUSION

The global network of outdoor webcams represents an underutilized resource for measuring the natural world. We conjecture that this resource has been ignored because of the significant challenges in finding, organizing, archiving images from, and calibrating a large number of webcams. This paper outlines our work to overcome these challenges and demonstrates several applications that use the images to measure environmental properties.

In addition to these direct benefits, there are outstanding opportunities for outreach to the general public, e.g. by linking webcam-based monitoring with educational programs to inform the public about the effects of climate change on our natural environment.

## Acknowledgement

We gratefully acknowledge the support of NSF CAREER grant (IIS-0546383) which partially supported this work. ADR acknowledges support from the Northeastern States Research Cooperative.

# 6. REFERENCES

[1] E. Bradley, D. Roberts, and C. Still. Design of an image analysis website for phenological and meteorological monitoring. *Environmental Modelling and Software*, 2009.

[2] M. Brand. Incremental singular value decomposition of uncertain data with missing values. In *Proceedings European Conference on Computer Vision*, 2002.

[3] J. Cho and H. Garcia-Molina. Estimating frequency of change. *ACM Trans. Internet Technology*, 3(3):256–290, 2003.

[4] Haze Cam Pollution Visibility Camera Network. http://www.hazecam.net/.

[5] R. Holman, J. Stanley, and T. Ozkan-Haller. Applying video sensor networks to nearshore environment monitoring. *IEEE Pervasive Computing*, 2(4):14–21, 2003.

[6] A. D. Hopkins. Periodical events and natural law as guides to agricultural research and practice. *Monthly Weather Review. Supplement No. 9*, pages 1–42, 1918.

[7] H. Hotelling. Relations between two sets of variates. *Biometrika*, 28:321–377, 1936.

[8] IPInfoDB. http://www.ipinfodb.com/.

[9] G. J., A. García-Olivares, E. Ojeda, A. Osorio, O. Chic, and R. Gonzalez. Long-Term Quantification of Beach Users Using Video Monitoring. *Journal of Coastal Research*, 24(6):1612–1619, 2008.

[10] N. Jacobs, W. Burgin, R. Speyer, D. Ross, and R. Pless. Adventures in archiving and using three years of webcam images. In *Proceedings IEEE CVPR Workshop on Internet Vision*, 2009.

[11] N. Jacobs, N. Roman, and R. Pless. Consistent temporal variations in many outdoor scenes. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

[12] N. Jacobs, N. Roman, and R. Pless. Toward fully automatic geo-location and geo-orientation of static outdoor cameras. In *Proceedings IEEE Workshop on Applications of Computer Vision*, 2008.

[13] N. Jacobs, S. Satkin, N. Roman, R. Speyer, and R. Pless. Geolocating static cameras. In *Proceedings IEEE International Conference on Computer Vision*, 2007.

[14] M. Kammler and G. Schernewski. Spatial and temporal analysis of beach tourism using webcam and aerial photographs. In G. Schernewski and N. Löser, editors, *Managing the Baltic Sea: Coastline Reports*. 2004.

[15] S. J. Kim, J.-M. Frahm, and M. Pollefeys. Radiometric calibration with illumination change for outdoor scene analysis. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 2008.

[16] S. J. Koppal and S. G. Narasimhan. Clustering appearance for scene analysis. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 2006.

[17] J.-F. Lalonde, S. G. Narasimhan, and A. A. Efros. What does the sky tell us about the camera? In *Proceedings European Conference on Computer Vision*, 2008.

[18] A. Moreno. The role of weather in beach recreation—a case study using webcam images. In A. Matzarakis,


(a) (b)

**Figure 9: Using webcam images for phenological monitoring has advantages in operating at a much finer spatial and temporal resolution than satellite imagery. Here we show that the higher spatial resolution enables distinguishing between trees with different spring leaf growth rates. (a) The value of the greenness signal (black dots) for a camera viewing a clock-tower in a plaza. The thin (red) curve is the value of a single-sigmoid model fit to the data. The thick (green) curve is the value of a mixture-of-sigmoids model fit to the data. (b) A montage of images captured by the camera. The first image (left) has a color overlay that corresponds to which component of the mixture-of-sigmoids model best fits the time-series of the underlying pixel. The other images provide evidence that the segmentation is meaningful (the purple regions grow leaves earlier than the red regions).**

C. R. de Freitas, and D. Scott, editors, *Developments in Tourism Climatology*. International Society of Biometeorology: Commission on Climate, Tourism and Recreation, 2007.

[19] J. T. Morisette, A. D. Richardson, A. K. Knapp, J. I. Fisher, E. A. Graham, J. Abatzoglou, B. E. Wilson, D. D. Breshears, G. M. Henebry, J. M. Hanes, and L. Liang. Tracking the rhythm of the seasons in the face of global change: phenological research in the 21st century. *Frontiers in Ecology and the Environment*, 7(5):253–260, 2009.

[20] S. G. Narasimhan, C. Wang, and S. K. Nayar. All the images of an outdoor scene. In *Proceedings European Conference on Computer Vision*, 2002.

[21] National Climatic Data Center. http://www.ncdc.noaa.gov/.

[22] Opentopia Webcams. http://www.opentopia.com/hiddencam.php.

[23] M. Parry, O. Canziani, J. Palutikof, P. van der Linden, and C. Hanson, editors. *Climate Change 2007: Impacts, Adaptation and Vulnerability*. Cambridge University Press, 2007.

[24] PhenoCam. http://phenocam.unh.edu/.

[25] Phenological Eyes Network. http://pen.agbi.tsukuba.ac.jp/.

[26] A. D. Richardson, B. H. Braswell, D. Y. Hollinger, J. P. Jenkins, and S. V. Ollinger. Near-surface remote sensing of spatial and temporal variation in canopy phenology. *Ecological Applications*, 19(6):1417–1428, 2009.

[27] A. D. Richardson, J. P. Jenkins, B. H. Braswell, D. Y. Hollinger, S. V. Ollinger, and M.-L. Smith. Use of digital webcam images to track spring green-up in a deciduous broadleaf forest. *Oecologia*, 152(2), 2007.

[28] K. Sunkavalli, W. Matusik, H. Pfister, and S. Rusinkiewicz. Factored time-lapse video. *ACM Transactions on Graphics*, 26(3), 2007.

[29] K. Sunkavalli, F. Romeiro, W. Matusik, T. Zickler, and H. Pfister. What do color changes reveal about an outdoor scene? In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 2008.

[30] Weather Underground. http://www.wunderground.com/.

[31] Weatherbug Inc. http://weatherbugmedia.com/.

[32] C. J. Wong, M. Z. MatJafri, K. Abdullah, and H. S. Lim. Temporal and spatial air quality monitoring using internet surveillance camera and alos satellite image. In *Proceedings IEEE Aerospace Conference*, 2009.

[33] C. J. Wong, M. Z. MatJafri, K. Abdullah, H. S. Lim, and K. L. Low. Temporal air quality monitoring using internet surveillance camera and alos satellite image. In *Proceedings IEEE Geoscience and Remote Sensing Symposium*, 2007.

[34] L. Xie, A. Chiu, and S. Newsam. Estimating atmospheric visibility using general-purpose cameras. In *Proceedings International Symposium on Advances in Visual Computing*, pages 356–367, 2008.

[35] X. Zhang, M. A. Friedl, C. B. Schaaf, A. H. Strahler, J. C. F. Hodges, F. Gao, B. C. Reed, and A. Huete. Monitoring vegetation phenology using modis. *Remote Sensing of Environment*, 84:471–475, 2003.