

Mining Photo-sharing Websites to Study Ecological Phenomena

Haipeng Zhang
School of Informatics &
Computing
Indiana University
Bloomington, IN
zhanhaip@indiana.edu

Mohammed Korayem
School of Informatics &
Computing
Indiana University
Bloomington, IN
mkorayem@indiana.edu

David J. Crandall
School of Informatics &
Computing
Indiana University
Bloomington, IN
djcran@indiana.edu

Gretchen LeBuhn
Department of Biology
University of California
San Francisco, CA
lebuhn@sfsu.edu

ABSTRACT

The popularity of social media websites like Flickr and Twitter has created enormous collections of user-generated content online. Latent in these content collections are observations of the world: each photo is a visual snapshot of what the world looked like at a particular point in time and space, for example, while each tweet is a textual expression of the state of a person and his or her environment. Aggregating these observations across millions of social sharing users could lead to new techniques for large-scale monitoring of the state of the world and how it is changing over time. In this paper we step towards that goal, showing that by analyzing the tags and image features of geo-tagged, time-stamped photos we can measure and quantify the occurrence of ecological phenomena including ground snow cover, snow fall and vegetation density. We compare several techniques for dealing with the large degree of noise in the dataset, and show how machine learning can be used to reduce errors caused by misleading tags and ambiguous visual content. We evaluate the accuracy of these techniques by comparing to ground truth data collected both by surface stations and by Earth-observing satellites. Besides the immediate application to ecology, our study gives insight into how to accurately crowd-source other types of information from large, noisy social sharing datasets.

Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications—*Data Mining, Image Databases, Spatial Databases and GIS*; I.4.8 [Image Processing and Computer Vision]: Scene Analysis

General Terms

Measurement, Theory

Keywords

Data mining, social media, photo collections, crowd-sourcing, ecology

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.
WWW 2012, April 16–20, 2012, Lyon, France.
ACM 978-1-4503-1229-5/12/04.

1. INTRODUCTION

The popularity of social networking websites has grown dramatically over the last few years, creating enormous collections of user-generated content online. Photo-sharing sites have become particularly popular: Flickr and Facebook alone have amassed an estimated 100 billion images, with over 100 million new images uploaded every day [18]. People use these sites to share photos with family and friends, but in the process they are creating immense public archives of information about the world: each photo is a record of what the world looked like at a particular point in time and space. When combined together, the billions of photos on these sites combined with metadata including timestamps, geo-tags, and captions are a rich untapped source of information about the state of the world and how it is changing over time.

Recent work has studied how to mine passively-collected data from social networking and microblogging websites to make estimates and predictions about world events, including tracking the spread of disease [11], monitoring for fires and emergencies [9], predicting product adoption rates and election outcomes [16], and estimating aggregate public mood [5, 22]. In most of these studies, however, there is either little ground truth available to judge the quality of the estimates and predictions, or the available ground truth is an indirect proxy (e.g. since no aggregate public mood data exists, [22] evaluates against opinion polls, while [5] compares to stock market indices). While these studies have demonstrated promising results, it is not yet clear when crowd-sourcing data from social media sites can yield reliable estimates, or how to deal with the substantial noise and bias in these datasets. Moreover, these studies have largely focused on textual content and have not taken advantage of the vast amount of visual content online.

In this paper, we study the particular problem of estimating geotemporal distributions of ecological phenomena using geo-tagged, time-stamped photos from Flickr. Our motivations to study this particular problem are three-fold. First, biological and ecological phenomena frequently appear in images, both because photographers take photos of them purposely (e.g. close-ups of plants and animals) or incidentally (a bird in the background of a family portrait, or the snow in the action shot of children sledding). Second, for the two phenomena we study here, snowfall and vegetation cover, large-scale (albeit imperfect) ground truth is available in the form of observations from satellites and ground-based weather stations. Thus we can explicitly evaluate the accuracy of various techniques




Figure 1: Comparing MODIS satellite snow coverage data for North America on Dec 21, 2009 with estimates produced by analyzing Flickr tags (best viewed on screen in color). *Left:* Original MODIS snow data, where white corresponds with water, black is missing data because of cloud cover, grey indicates snow cover, and purple indicates no significant snow cover. *Middle:* Satellite data coarsened into 1 degree bins, where green indicates snow cover, blue indicates no snow, and grey indicates missing data. *Right:* Estimates produced by the Flickr photo analysis proposed in this paper, where green indicates high probability of snow cover, and grey and black indicate low-confidence areas (with few photos or ambiguous evidence).

for extracting semantic information from large-scale social media collections.

Third, while ground truth is available for these particular phenomena, for other important ecological phenomena (like the geo-temporal distribution of plants and animals) no such data is available, and social media could help fill this need. In fact, perhaps no community is in greater need of real-time, global-scale information on the state of the world than the scientists who study climate change. Recent work shows that global climate change is impacting a variety of flora and fauna at local, regional and continental scales: for example, species of high-elevation and cold-weather mammals have moved northward, some species of butterflies have become extinct, waterfowl are losing coastal wetland habitats as oceans rise, and certain fish populations are rapidly declining [23]. However monitoring these changes is surprisingly difficult: plot-based studies involving direct observation of small patches of land yield high-quality data but are costly and possible only at very small scales, while aerial surveillance gives data over large land areas but cloud cover, forests, atmospheric conditions and mountain shadows can interfere with the observations, and only certain types of ecological information can be collected from the air. To understand how biological phenomena are responding to both landscape changes and global climate change, ecologists need an efficient system for ground-based data collection to give detailed observations across the planet. A new approach for creating ground-level, continental-scale datasets is to use passive data-mining of the huge number of visual observations produced by millions of users worldwide, in the form of digital images uploaded to photo-sharing websites.

Challenges. There are two key challenges to unlocking the ecological information latent in these photo datasets. The first is how to recognize ecological phenomena appearing in photos and how to map these observations to specific places and times. Fortunately, modern photo-sharing sites collect a rich variety of non-visual information about photos, including metadata recorded by the digital camera — exposure settings and timestamps, for example — as well as information generated during social sharing — text tags, comments, and ratings, for example. Many sites also record the geographic coordinates of where on Earth a photo was taken, as reported either by a GPS-enabled camera or smartphone, or input manually by the user. Thus online photos include the ingredients necessary to produce geo-temporal data about the world, including information about content (images, tags and comments), and when (timestamp) and where (geotag) each photo was taken.

The second challenge is how to deal with the biases and noise inherent in online data. People do not photograph the Earth evenly,

so there are disproportionate concentrations of activity near cities and tourist attractions. Photo metadata is often noisy or inaccurate; for example, users forget to set the clock on their camera, GPS units fail to find fixes, and users carelessly tag photos. Even photos without such errors might be misleading: the tag “snow” on an image might refer to a snow lily or a snowy owl, while snow appearing in an image might be artificial (as in an indoor zoo exhibit).

This paper. In this paper we study how to mine data from photo-sharing websites to produce crowd-sourced observations of ecological phenomena. As a first step towards the longer-term goal of mining for many types of phenomena, here we study two in particular: ground snow cover and vegetation cover (“green-up”) data. Both are critical features for ecologists monitoring the earth’s ecosystems. Importantly for our study, these two phenomena have accurate fine-grained ground truth available at a continental scale in the form of observations from aerial instruments like NASA’s Terra earth-observing satellites [12, 19] or networks of ground-based observing stations run by the U.S. National Weather Service. This data allows us to evaluate the performance of our crowd-sourced data mining techniques at a very large scale, including thousands of days of data across an entire continent. Using a dataset of nearly 150 million geo-tagged Flickr photos, we study whether this data can potentially be a reliable resource for scientific research. An example comparing ground truth snow cover data with the estimates produced by our Flickr analysis on one particular day (December 21, 2009) is shown in Figure 1. Note that the Flickr analysis is sparse in places with few photographs, while the satellite data is missing in areas with cloud cover, but they agree well in areas where both observations are present. This (and the much more extensive experimental results presented later in the paper) suggests that Flickr analysis may produce useful observations either on its own or as a complement other observational sources.

To summarize, the main contributions of this paper include:

- introducing the novel idea of mining photo-sharing sites for geo-temporal information about ecological phenomena,
- introducing several techniques for deriving crowd-sourced observations from noisy, biased data using both visual and textual tag analysis, and
- evaluating the ability of these techniques to accurately measure these phenomena, using dense large-scale ground truth.

2. RELATED WORK

A variety of recent work has studied how to apply computational techniques to analyze online social datasets in order to aid research

in other disciplines [20]. Much of this work has studied questions in sociology and human interaction, such as how friendships form [8], how information flows through social networks [21], how people move through space [6], and how people influence their peers [4]. The goal of these projects is not to measure data about the physical world itself, but instead to discover interesting properties of human behavior using social networking sites as a convenient data source.

Crowd-sourced observational data. Other studies have shown the power of social networking sites as a source of observational data about the world itself. Bollen *et al* [5] use data from Twitter to try to measure the aggregated emotional state of humanity, computing mood across six dimensions according to a standard psychological test. Intriguingly, they find that these changing mood states correlate well with the Dow Jones Industrial Average, allowing stock market moves to be predicted up to 3 days in advance. However their test dataset is relatively small, consisting of only three weeks of trading data. Like us, Jin *et al* [16] use Flickr as a source of data for prediction, but they estimate the adoption rate of consumer photos by monitoring the frequency of tag use over time. They find that the volume of Flickr tags is correlated with sales of two products, Macs and iPods. They also estimate geo-temporal distributions of these sales over time but do not compare to ground truth, so it is unclear how accurate these estimates are. In contrast, we evaluate our techniques against a large ground truth dataset, where the task is to accurately predict the distribution of a phenomenon (e.g. snow) across an entire continent each day for several years.

Crowd-sourced geo-temporal data. Other work has used online data to predict geo-temporal distributions, but again in domains other than ecology. Perhaps the most striking is the work of Ginsberg *et al* [11], who show that by monitoring the geospatial distribution of search engine queries related to flu symptoms, the spread of the H1N1 flu can be estimated several days before the official statistics produced by traditional means. DeLongueville *et al* [9] study tweets related to a major fire in France, but their analysis is at a very small scale (a few dozen tweets) and their focus is more on human reactions to the fire as opposed to using these tweets to estimate the fire’s position and severity. In perhaps the most related existing work to ours, Singh *et al* [24] create geospatial heat maps (dubbed “social pixels”) of various tags, including snow and greenery, but their focus is on developing a formal database-style algebra for describing queries on these systems and for creating visualizations. They do not consider how to produce accurate predictions from these visualizations, nor do they compare to any ground truth.

Citizen science. While some volunteer-based biology efforts like the Lost Ladybug Project [3] and the Great Sunflower Project [2] use social networking sites to organize and recruit volunteer observers, we are not aware of any work that has attempted to passively mine ecological data from social media sites. The visual data in online social networking sites provide a unique resource for tracking biological phenomena: because they are images, this data can be verified in ways that simple text cannot. In addition, the rapidly expanding quantity of online images with geo-spatial and temporal metadata creates a fine-scale record of what is happening across the globe. However, to unlock the latent information in these vast photo collections, we need mining and recognition tools that can efficiently process large numbers of images, and robust statistical models that can handle incomplete and incorrect observations.

3. OUR APPROACH

We use a sample of nearly 150 million geo-tagged, timestamped Flickr photos as our source of user-contributed observational data

about the world. We collected this data using the public Flickr API, by repeatedly searching for photos within random time periods and geo-spatial regions, until the entire globe and all days between January 1, 2007 and December 31, 2010 had been covered. We applied filters to remove blatantly inaccurate metadata, in particular removing photos with geotag precision less than about city-scale (as reported by Flickr), and photos whose upload timestamp is the same as the EXIF camera timestamp (which usually means that the camera timestamp was missing).

For ground truth we use large-scale data originating from two independent sources: ground-based weather stations, and aerial observations from satellites. For the ground-based observations, we use publicly-available daily snowfall and snow depth observations from the U.S. National Oceanic and Atmospheric Administration (NOAA) Global Climate Observing System Surface Network (GSN) [1]. This data provides highly accurate daily data, but only at sites that have surface observing stations. For denser, more global coverage, we also use data from the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument aboard NASA’s Terra satellite. The satellite is in a polar orbit so that it scans the entire surface of the earth every day. The MODIS instrument measures spectral emissions at various wavelengths, and then post-processing uses these measurements to estimate ground cover. In this paper we use two datasets: the daily snow cover maps [12] and the two-week vegetation averages [19]. Both of these sets of data including an estimate of the percentage of snow or vegetation ground cover at each point on earth, along with a quality score indicating the confidence in the estimate. Low confidence is caused primarily by cloud cover (which changes the spectral emissions and prevents accurate ground cover from being estimated), but also by technical problems with the satellite. As an example, Figure 1 shows raw satellite snow data from one particular day.

3.1 Estimation techniques

Our goal is to estimate the presence or absence of a given ecological phenomenon (like a species of plant or flower, or a meteorological feature like snow) on a given day and at a given place, using only the geo-tagged, time-stamped photos from Flickr. One way of viewing this problem is that every time a user takes a photo of a phenomenon of interest, they are casting a “vote” that the phenomenon actually occurred in a given geospatial region. We could simply look for tags indicating the presence of a feature – i.e. count the number of photos with the tag “snow” – but sources of noise and bias make this task challenging, including:

- *Sparse sampling:* The geospatial distribution of photos is highly non-uniform. A lack of photos of a phenomenon in a region does not necessarily mean that it was not there.
- *Observer bias:* Social media users are younger and wealthier than average, and most live in North America and Europe.
- *Incorrect, incomplete and misleading tags:* Photographers may use incorrect or ambiguous tags — e.g. the tag “snow” may refer to a snowy owl or interference on a TV screen.
- *Measurement errors:* Geo-tags and timestamps are often incorrect (e.g. because people forget to set their camera clocks).

A statistical test. We introduce a simple probabilistic model and use it to derive a statistical test that can deal with some such sources of noise and bias. The test could be used for estimating the presence of any phenomenon of interest; without loss of generality we use the particular case of snow here, for ease of explanation. Any given photo either contains evidence of snow (event s) or does not contain evidence of snow (event \bar{s}). We assume that a given photo taken at a time and place with snow has a fixed probability $P(s|snow)$

of containing evidence of snow; this probability is less than 1.0 because many photos are taken indoors, and outdoor photos might be composed in such a way that no snow is visible. We also assume that photos taken at a time and place without snow have some non-zero probability $P(s|\overline{snow})$ of containing evidence of snow; this incorporates various scenarios including incorrect timestamps or geo-tags and misleading visual evidence (e.g. man-made snow).

Let m be the number of snow photos (event s), and n be the number of non-snow photos (event \bar{s}) taken at a place and time of interest. Assuming that each photo is captured independently, we can use Bayes' Law to derive the probability that a given place has snow given its number of snow and non-snow photos,

$$\begin{aligned} P(snow|s^m, \bar{s}^n) &= \frac{P(s^m, \bar{s}^n|snow)P(snow)}{P(s^m, \bar{s}^n)} \\ &= \frac{\binom{m+n}{m} p^m (1-p)^n P(snow)}{P(s^m, \bar{s}^n)}, \end{aligned}$$

where we write s^m, \bar{s}^n to denote m occurrences of event s and n occurrences of event \bar{s} , and where $p = P(s|snow)$ and $P(snow)$ is the prior probability of snow. A similar derivation gives the posterior probability that the bin does not contain snow,

$$P(\overline{snow}|s^m, \bar{s}^n) = \frac{\binom{m+n}{m} q^m (1-q)^n P(\overline{snow})}{P(s^m, \bar{s}^n)},$$

where $q = P(s|\overline{snow})$. Taking the ratio between these two posterior probabilities yields a likelihood ratio,

$$\frac{P(snow|s^m, \bar{s}^n)}{P(\overline{snow}|s^m, \bar{s}^n)} = \frac{P(snow)}{P(\overline{snow})} \left(\frac{p}{q}\right)^m \left(\frac{1-p}{1-q}\right)^n. \quad (1)$$

This ratio can be thought of as a measure of the confidence that a given time and place actually had snow, given photos from Flickr.

A simple way of classifying a photo into a positive event s or a negative event \bar{s} is to use text tags. We identify a set \mathcal{S} of tags related to a phenomenon of interest. Any photo tagged with at least one tag in \mathcal{S} is declared to be a positive event s , and otherwise it is considered a negative event \bar{s} . For the snow detection task, we use the set $\mathcal{S}=\{\text{snow, snowy, snowing, snowstorm}\}$, which we selected by hand.

The above derivation assumes that photos are taken independently of one another, which is generally not true in reality. One particular source of dependency is that photos from the same user are highly correlated with one another. To mitigate this problem, instead of counting m and n as numbers of *photos*, we instead let m be the number of *photographers* having at least one photo with evidence of snow, while n is the numbers of photographers who did not upload any photos with evidence of snow.

The probability parameters in the likelihood ratio of equation (1) can be directly estimated from training data and ground truth. For example, for the snow cover results presented in Section 4, the learned parameters are: $p = p(s|snow) = 17.12\%$, $q = p(s|\overline{snow}) = 0.14\%$. In other words, almost 1 of 5 people at a snowy place take a photo containing snow, whereas about 1 in 700 people take a photo containing evidence of snow at a non-snowy place.

Figure 1 shows a visualization of the likelihood ratio values for the U.S. on one particular day using this simple technique with $\mathcal{S}=\{\text{snow, snowy, snowing, snowstorm}\}$. High likelihood ratio values are plotted in green, indicating a high confidence of snow in a geospatial bin, while low values are shown in blue and indicate high confidence of no snow. Black areas indicate a likelihood ratio near 1, showing little confidence either way, and grey areas lack data entirely (having no Flickr photos in that bin on that day).

3.2 Learning features automatically

The confidence score in the last section has a number of limitations, including requiring that a set of tags related to the phenomenon of interest be selected by hand. Moreover, it makes no attempt to incorporate visual evidence or negative textual evidence — e.g., that a photo tagged “snowy owl” probably contains a bird and no actual snow. We use machine learning techniques to address these weaknesses, both to automatically identify specific tags and tag combinations that are correlated with the presence of a phenomenon of interest, and to incorporate visual evidence into the prediction techniques.

Learning tags. We consider two learning paradigms. The first is to produce a single exemplar for each bin in time and space consisting of the set of all tags used by all users. For each of these exemplars, the NASA and/or NOAA ground truth data gives a label (snow or non-snow). We then use standard machine learning algorithms like Support Vector Machines and decision trees to identify the most discriminative tags and tag combinations. In the second paradigm, our goal instead is to classify individual *photos* as containing snow or not, and then use these classifier outputs to compute the number of positive and non-positive photos in each bin (i.e., to compute m and n in the likelihood ratio described in the last section).

Learning visual features. We also wish to incorporate visual evidence from the photos themselves. There is decades of work in the computer vision community on object and scene classification (see [27] for a recent survey), although most of that work has not considered the large, noisy photo collections we work with here. We tried a number of approaches, and found that a classifier using a simplified version of GIST augmented with color features [14,28] gave a good trade-off between accuracy and tractability.

Given an image I , we partition the image into a 4×4 grid of 16 equally-sized rectangular regions. In each region we compute the average pixel values in each of the red, green, and blue color planes, and then convert this color triple from sRGB space to the CIELAB color space [15]. CIELAB has a number of advantages, including separating greyscale intensity from the color channels and having greater perceptual uniformity (so that Euclidean distances between two CIELAB color triples are approximately proportional to the human perception of difference between the colors). For each region R we also compute the total gradient energy $E(R)$ within the grayscale plane I_g of the image,

$$\begin{aligned} E(R) &= \sum_{(x,y) \in R} \|\nabla I_g(x,y)\| \\ &= \sum_{(x,y) \in R} \sqrt{I_x(x,y)^2 + I_y(x,y)^2}, \end{aligned}$$


where $I_x(x,y)$ and $I_y(x,y)$ are the partial derivatives in the x and y directions evaluated at point (x,y) , approximated as,

$$\begin{aligned} I_x(x,y) &= I_g(x+1,y) - I_g(x-1,y), \\ I_y(x,y) &= I_g(x,y+1) - I_g(x,y-1). \end{aligned}$$

For each image we concatenate the gradient energy in each of the 16 bins, followed by the 48 color features (average L, a, and b values for each of the 16 bins), to produce a 64-dimensional feature vector. We then learn a Support Vector Machine (SVM) classifier from a labeled training image set.

4. EXPERIMENTS AND RESULTS

We now turn to presenting experimental results for estimating the geo-temporal distributions of two ecological phenomena: snow



	NYC	Chicago	Boston	Philadelphia
Mean active Flickr users / day	65.6	94.9	59.7	43.7
Approx. city area (km^2)	3,712	11,584	11,456	9,472
User density (avg users/unit area)	112.4	52.5	33.5	29.6
Mean daily snow (inches)	0.28	0.82	0.70	0.35
Snow days (snow>0 inches)	185	418	373	280
Number of obs. stations	14	20	41	26

Figure 2: *Top:* New York City geospatial bounding box used to select Flickr photos, and locations of NOAA observation stations. *Bottom:* Statistics about spatial area, photo density, and ground truth for each of the 4 cities.

and vegetation cover. In addition to the likelihood ratio-based score described in Section 4 and machine learning approaches, we also compare to two simpler techniques: *voting*, in which we simply count the number of users that use one of a set S of tags related to the phenomenon of interest at a given time and place, and *percentage*, in which we calculate the ratio of users that use one of the tags in S over the total number of users who took a photo in that place on that day.

4.1 Snow prediction in cities

We first test how well the Flickr data can predict snowfall at a local level, and in particular for cities in which high-quality surface-based snowfall observations exist and for which photo density is high. We choose 4 U.S. metropolitan areas, New York City, Boston, Chicago and Philadelphia, and try to predict both daily snow presence as well as the quantity of snowfall. For each city, we define a corresponding geospatial bounding box and select the NOAA ground observation stations in that area. For example, Figure 2 shows the the stations and the bounding box for New York City. We calculate the ground truth daily snow quantity for a city as the average of the valid snowfall values from its stations. We call any day with a non-zero snowfall or snowcover to be a snow day, and any other day to be a non-snow day. Figure 2 also presents some basic statistics for these 4 cities. All of our experiments involve 4 years (1461 days) of data from January 2007 through December 2010; we reserve the first two years for training and validation, and the second two years for testing.

Daily snow classification for 4 cities. Figure 3(a) presents ROC curves for this daily snow versus non-snow classification task on New York City. The figure compares the likelihood ratio confidence score from equation (1) to the baseline approaches (voting and percentage), using the tag set $S=\{\text{snow, snowy, snowing, snow-storm}\}$. The area under the ROC curve (AUC) statistics are 0.929, 0.905, and 0.903 for confidence, percentage, and voting, respectively, and the improvement of the confidence method is statistically significant with $p = 0.0713$ according to the statistical test of [29]. The confidence method also outperforms other methods for the other three cities (not shown due to space constraints). ROC curves for all 4 cities using the likelihood scores are shown in Figure 3(b). Chicago has the best performance and Philadelphia has the worst; a possible explanation is that Chicago has the most active Flickr users per day (94.9) while Philadelphia has the least (43.7).

These methods based on presence or absence of tags are simple

and very fast, but they have a number of disadvantages, including that the tag set must be manually chosen and that negative correlations between tags and phenomena are not considered. We thus tried training a classifier to learn these relationships automatically. For each day in each city, we produce a single binary feature vector indicating whether or not a given tag was used on that day. We also tried a feature selection step by computing information gain and rejecting features below a threshold, as well as adding the likelihood score from equation (1) as an additional feature. For all experiments we used feature vectors from 2007 and 2008 for training and tested on data from 2009 and 2010, and used a LibLinear classifier with L2-regularized logistic regression [10]. Table 1 presents the results, showing that information gain (IG) and confidence scores (Conf) improve the results for all cities, and that the classifier built with both IG and Conf generally outperforms other classifiers, except for Boston. Figure 3(c) shows ROC curves from different classifiers for NYC and Figure 3(d) compares ROC curves for the 4 cities using the classifier using both feature selection and confidence. Note that the machine learning-based techniques substantially outperform the simple likelihood ratio approach (compare Figures 3(b) and (d)).

Predicting snow quantities. In addition to predicting simple presence or absence of a phenomenon, it may be possible to predict the degree or quantity of that phenomenon. Here we try one particular approach, using our observation that the numerical likelihood score of equation (1) is somewhat correlated with depth of snow ($R^2=0.2972$) — i.e., that people take more photos of more severe storms (see Figure 4). Because snow cover is temporally correlated, we fit a multiple linear regression model in which the confidence scores of the last several days are incorporated. The prediction on day t is then given by,

$$\begin{cases} \sum_{i=0}^T \alpha_i \log(\text{conf}_{t-i}) + \beta & \text{if } \text{conf}_t \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

where conf_t represents the likelihood ratio from equation (1) on day t , T is the size of the temporal window, and the α and β pa-

Table 1: Daily snow classification results for a 2 year period (2009–2010) for four major metropolitan areas.

Features	Accuracy	Precision	Recall	F-Measure	Baseline
NYC					
Tags	0.859	0.851	0.859	0.805	0.85
Tags+Conf.	0.926	0.927	0.926	0.917	0.85
Tags+IG	0.91	0.906	0.91	0.898	0.85
Tags+IG+Conf.	0.93	0.93	0.93	0.923	0.85
Boston					
Tags	0.899	0.897	0.899	0.894	0.756
Tags+Conf.	0.93	0.929	0.93	0.929	0.756
Tags+IG	0.91	0.911	0.91	0.91	0.756
Tags+IG+Conf.	0.923	0.923	0.923	0.923	0.756
Chicago					
Tags	0.937	0.938	0.937	0.935	0.728
Tags+Conf.	0.949	0.952	0.949	0.948	0.728
Tags+IG	0.938	0.938	0.938	0.938	0.728
Tags+IG+Conf.	0.953	0.954	0.953	0.953	0.728
Philadelphia					
Tags	0.849	0.851	0.849	0.815	0.805
Tags+Conf.	0.912	0.917	0.912	0.903	0.805
Tags+IG	0.903	0.899	0.903	0.897	0.805
Tags+IG+Conf.	0.927	0.926	0.927	0.924	0.805




Figure 3: ROC curves for binary snow predictions: (a) ROC curves for New York City, comparing likelihood ratio confidence score to voting and percentage approaches, (b) ROC curves for 4 cities using the likelihood scores, (c) ROC curves from SVM classifiers with different features for New York City, and (d) ROC curves for 4 cities using the logistic regression (LibLinear) classifier with tags, information gain and confidence features. (Best viewed in color.)




Figure 4: Time series of actual daily snow (top) and score estimated from Flickr (bottom) for New York City, 2007–2010.

rameters are learned from the training data. We found that increasing T generally improves performance on the 4 cities, but that no additional improvement occurred with $T > 3$. We can measure the error of our predictions with the root-mean-squared error between the time series of our predictions and the actual snow data (following [16]). We achieve an RMS error of between about 1 and 1.5 inches across the 4 cities; Philadelphia has the largest error (1.44), followed by Boston (1.26), New York (1.15), and Chicago (1.06). As an example, Figure 5 presents a visual comparison of the prediction time series versus the actual snow time series for Chicago.

An alternative way of evaluating the snow quantity estimates is to view it as a multi-way classification task. We follow an existing snowfall impact scale [25] and quantize daily snow quantity into 7 buckets: no snow, 0-1 inches, 1-4 inches, 4-10 inches, 10-20 inches, 20-30 inches, or more than 30 inches. We then build a classifier to predict the snow ranges for the four cities using the numbers of snow and non-snow users. We include the numbers of users from the previous three days as extra features. We use a Naive Bayesian classifier [17], which performed best on this task. These multi-way classification results are better than a majority class baseline, with 7-way correct classification rates at 87.5% for Philadelphia, 87.9% for New York, 84.0% for Boston, and 83.7% for Chicago (versus baselines of 80.5%, 85.1%, 75.6%, and 72.9%, respectively).

4.2 Continental-scale snow prediction

Predicting snow for individual cities is of limited practical use because accurate meteorological data already exists for these highly




Figure 5: Comparing time series of actual daily snowfall (in mm) for Chicago with estimates using Flickr, for Jan 2009–Dec 2010 and $T = 3$. Red dots show predictions, and vertical bars show actual values.

populated areas. In this section we ask whether phenomena can be monitored at a continental scale, a task for which existing data sources are less complete and accurate. We use the photo data and ground truth described in Section 4, although for the experiments presented in this paper we restrict our dataset to North America (which we defined to be a rectangular region spanning from 10 degrees north, -130 degrees west to 70 degrees north, -50 degrees west). (We did this because Flickr is a dominant photo-sharing site in North America, while other regions have other popular sites — e.g. Fotolog in Latin America and Renren in China.)

The spatial resolution of the NASA satellite ground truth datasets is 0.05 degrees latitude by 0.05 degrees longitude, or about $5 \times 5 km^2$ at the equator. (Note that the surface area of these bins is non-uniform because lines of longitude get closer together near the poles.) However, because the number of photos uploaded to Flickr on any particular day and at any given spatial location is relatively low, and because of imprecision in Flickr geo-tags, we produce estimates at a coarser resolution of 1 degree square, or roughly $100 \times 100 km^2$. To make the NASA maps comparable, we downsample them to this same resolution by averaging the high confidence observations within the coarser bin. We then threshold the confidence and snow cover percentages to annotate each bin with one of three ground truth labels:

- Snow bin, if confidence is above 90 and coverage above 80,
- Non-snow bin, if confidence is above 90 and coverage is 0,
- Unknown bin, otherwise.

Our goal is to predict whether or not each geospatial bin had snow-cover on each day, given the photos from Flickr.

Retrieving snow or non-snow bins. In many real applications, ecologists would be satisfied in finding bins for which the phenomenon is present, rather than actually classifying all bins. It is thus useful to view this problem as a retrieval task, in which the

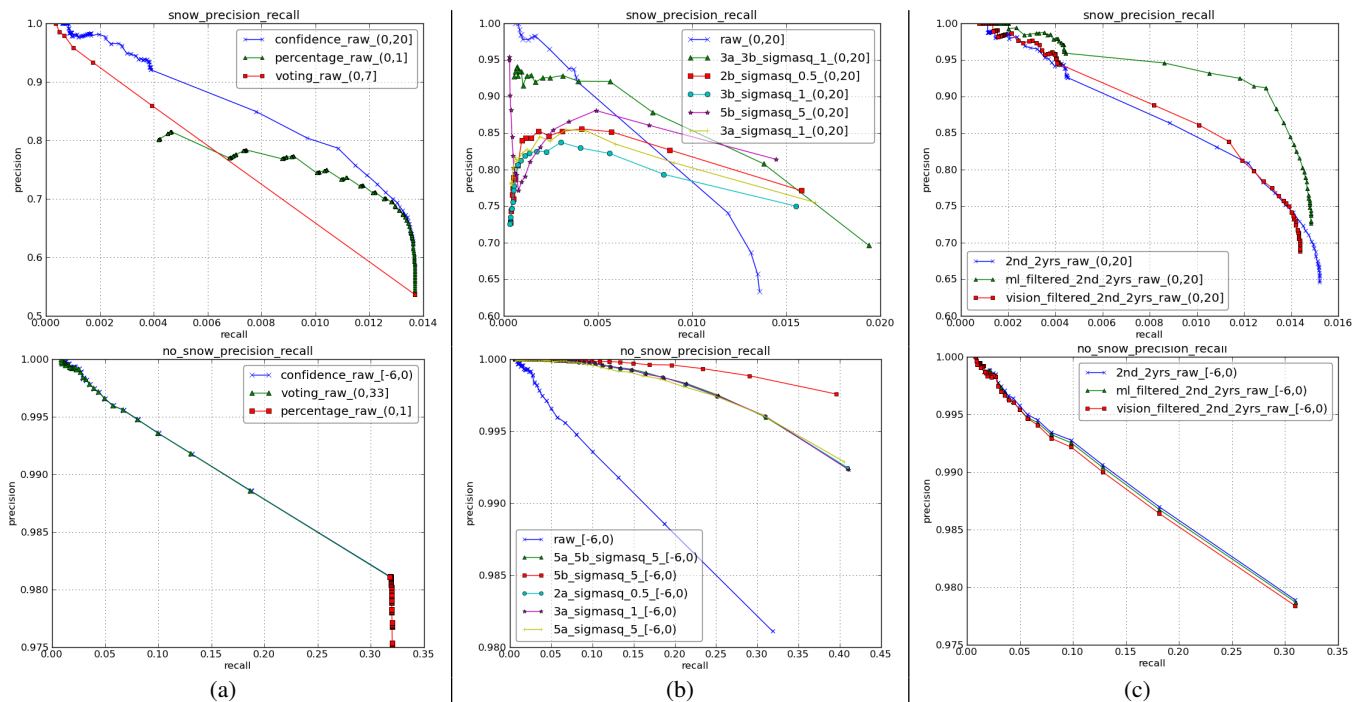


Figure 6: Precision and recall curves for retrieving snow (top) and non-snow (bottom) instances, where an instance is a single geospatial bin on a single day, using different techniques: (a) comparing the voting, percentage, and statistical confidence estimation techniques, (b) comparing different temporal smoothing strategies, (c) using classifiers to reject falsely-tagged snow images using visual and textual features.

goal is to identify bins likely to contain the phenomenon, or likely not to contain it. We thus turn to evaluating the performance of our estimation techniques using precision-recall curves, where

$$\text{precision} = \frac{|R \cap G|}{|R|} \quad \text{recall} = \frac{|R \cap G|}{|G|},$$

where R is the set of retrieved bins and G is the set of correct bins according to the ground truth. Precision-recall curves are also easier to interpret in situations where the classification baselines are so high, as in our case.

Figure 6(a) shows precision-recall curves for retrieving bins and days containing snow (top) and those not containing snow (bottom). In total, these curves involve classifying about 7 million exemplars (each of which is a single geospatial bin on a single day), of which 11.0% have ground truth. 82.2% of the bins with ground truth are no-snow bins, while snow bins account for 17.8%. We observe that the confidence method performs significantly better than the other two methods for retrieving snow bins, achieving about 98% precision at 0.2% recall, and about 80% precision at 1% recall. For retrieving non-snow bins the three techniques are almost the same, and all three perform better than the random baseline.

While the precisions in these curves are high, the recall values are alarming low. The main reason for this is that large areas of North America, particularly most of Canada and Alaska, have sparse populations resulting in a very limited number of photos uploaded in these areas. We showed in the last section that accurate snow estimates can be inferred for highly populated cities; the low recalls here are because of low photographic density in much of the continent. Restricting to specific subsets significantly increases the density of observations: for example, the average number of photos per bin over our four years of data is nearly ten times larger for the northeast US compared to all of North America (70,398 vs 8,134). The performance is significantly better in these more densely populated areas; for example, in the Northeast US the precision is 96.3% at a recall of 19.5% for snow retrieval, and 99.9% precision at 9.1%

recall for non-snow retrieval. Moreover, recall would naturally improve as our dataset grows; our sample of 150 million images is less than 3% of the photos on Flickr, and thus the recall would improve significantly if we had access to the entire dataset.

Temporal smoothing. For many phenomena (including snow), the existence of an event on one day is strongly correlated with its existence on the next day. Thus one way of addressing the sparsity of Flickr photos in some locations is to propagate evidence forward and backward in time. To do this, we apply a Gaussian filter on the Flickr confidence values for each bin in an attempt to achieve better recalls. We vary the degree of smoothing by using Gaussians with different variance values. We tried smoothing with many different parameters, including smoothing both forward and backwards in time, or in only one direction. Figure 6(b) shows curves for several of the best combinations that we found, including the raw confidence score (blue X's), 3 days before and after with variance 1.0 (brown triangles), 2 days before with variance 0.5 (red squares), 3 days before with variance 1.0 (blue circles), 5 days before with variance 5.0 (purple stars), and 3 days after with variance 1.0 (yellow +s). We find that temporal smoothing three days before and after with variance 1.0 significantly improves performance for both snow and non-snow retrieval, increasing snow retrieval precision by about 7 percentage points at 1% recall.

Voting. Voting performs worse than the statistical confidence given by the Bayesian likelihood ratio, but it is an interesting technique to study in more detail because of its simplicity. Voting simply counts the number of users who have annotated at least one photo in a given bin and day with a snow-related tag. Figure 7 plots precision versus the number of votes for snow retrieval. The shape of these curve illustrates why crowd-sourced observations of the world can be reliable, if enough people are involved: as the number of votes for snow increases, it becomes progressively less likely that these independent observations are coincidental, and more likely that




Figure 7: Precision vs number of votes for snow predictions using the voting method.

Table 2: Taxonomy of manually-labeled false-positive photos (which have at least one snow-related tag despite being taken at a snowless time and place according to the ground truth).

Class	Description	# of photos
little or distant	photos with trace amount of snow or snow in the distance	585 (33.0%)
man made	photos with snow made by humans (e.g. at a ski slope)	152 (8.6%)
no snow	photos without visible snow	737 (41.5%)
snow	photos with significant snow	279 (15.7%)
not sure	other photos	21 (1.2%)

they are caused by the presence or absence of an actual phenomenon. It is interesting to notice that when there are 7 or more snow voters, snow prediction precision becomes 100%, while the same is true for non-snow prediction when the number of non-snow voters reaches 33 if there are no snow voters in the bin.

Case study of false positives. To understand the failure modes of estimating attributes about the world from Flickr photos, we performed a case study of false positives — bins and days in which our Flickr mining predicted the presence of snow, but the NASA ground truth indicated that there was no snow cover. In particular, we studied snow false positives at the operating point at which the likelihood ratio method gives a precision of 74.1% and a recall of 1.2% (i.e. when the threshold is 4). At this operating point, 34,323 total predictions are made (each corresponding to a single geospatial bin on a single day), 2,208 of which have valid ground truth. Of these 2,208 bins, 1,636 (74.1%) are correctly classified, while the 572 false positive bins have a total of 1,855 photos tagged with one of the snow terms (despite the fact that they were taken at places and times in which the NASA satellite did not record snow). We manually examined these 1,855 false positive photos and classified them into 5 different classes according to their visual content, as shown in Table 2. Nearly 60% of these photos do actually appear to contain some snow; of these, 33% either show trace amount of snow or snow in the distance (usually on a distant mountain peak), and 8.6% have man-made snow that would not show up on the NASA maps (like in a zoo or ski slope), while only about 16% include a significant amount of natural snow. About 40% of the photos tagged with a snow-related term do not appear to contain any snow at all; these are caused by mis-tagged images or snow-related tags that are used to describe something else (like the interference on a TV screen). Figure 8 shows some sample false positives from each class.

For images that seem to contain natural snow, there are several possible explanations for why the ground truth does not indicate




Figure 8: Sample photos that were not taken at a place and time with snow according to the ground truth, but that were uploaded with a snow-related tag: (a) photo with trace amounts of snow, (b) photo with distant snow, (c) photo with man made snow, and (d) photo with no snow (but with a “snowy egret”).

snow cover at that time and place. One is that the satellite passes over at an unknown time of day, so it is possible that snowfall occurred after the satellite’s observation was taken. Another cause are photos with incorrect time stamps or geo-locations; we assume that such errors occur frequently, although it is hard to quantify the frequency just by looking at the photos. Other photos clearly contain snow, but the amount is so little that it might not be visible from the satellite (e.g. Figure 8(a)), or the snow is so far in the distance that it is in a different geospatial bin (e.g. Figure 8(b)).

There are some cases where the Flickr evidence for snow is overwhelming, but the NASA ground truth does not indicate snow. This could be caused by the timing issue described above, or by satellite resolution and confidence issues. For example, on February 21, 2008, 5 Flickr users reported snowfall in New York. This bin is marked as a no-snow bin in the ground truth because the vast majority of it has zero snow coverage according to the satellite, but there is a small area within the bin that has low confidence (due to cloud cover) and probably corresponds to a snow squall.

Machine learning for tag selection. Many of the above error modes can be addressed by training classifiers on textual tag and visual images features. As discussed in Section 3.1, we are interested in two learning paradigms: the first is to learn combinations of tags that classify geospatial bins well according to the NASA ground truth, while the second task is to reduce false positives by rejecting photos that are tagged with a snow term but do not actually contain snow.

In the first task, we want to learn to classify whether a given bin contains snow on a given day, based on a binary feature vector encoding the set of tags used by all users in that bin on that day. We tried four different classifiers to address this problem: REPTree, a fast decision tree learner which builds a decision tree using information gain and variance and prunes it using reduced-error pruning [13], Support Vector Machines (SVMs) [7], Discriminative Multinomial Naive Bayes (DMNB) [26] and LibLinear classifier with L2-regularized logistic regression [10]. To reduce the large number of features (a total of 404324 tags), we compute information gain and keep all features (13442 tags) with information gain greater than zero. Figure 9 presents ROC curves for this task, showing that the learned classifier outperforms the likelihood ratio from equation (1), and that feature selection with information gain




Figure 9: ROC curves for classifying whether a geo-bin has snow on a given day, comparing the LibLinear classifier with various tag features to the confidence method using hand-selected tags.




Figure 10: ROC curve for classifying whether photos contain snow, using decision trees with various features: AllTags includes all tags, IntersTags excludes tags corresponding to specific geographic areas, and AllTags+Time and IntersTags+Time include the month of the year as an additional feature.

and using the confidence ratio as an additional feature all improve performance.

Next we try the second learning paradigm, in which our goal is to examine photos that have a snow-related tag, and use the other tags as well as visual features to decide whether or not they actually contain snow. For example, the classifier might learn that a photo with “snowy” should be discarded if it also contains the tag “egret,” since that photo is likely of a bird and not of actual snow. For training these classifiers, we had a human judge evaluate 1,855 images and to annotate them as to whether or not they actually contain evidence of snow.

We used decision trees for this task because it is easy to understand and interpret what features the classifier is using. In initial experimentation, we found that many of the most discriminative features were place names, like “sandiego” or “canada.” These geographic tags are understandably strongly correlated with snowfall, but we would like our classifier to base its decisions on the content of an image (because, for example, climate change might cause snowfall in San Diego some day, and we would like our classifier to be able to detect this). To avoid selecting these tags, we first divide North America into four regions (northeast, northwest, southeast, southwest) and get the intersection of the sets of tags used in these four regions. We then use only this set of intersected tags (“IntersTags”) for building the decision tree. Besides tags, we also tried including the photo’s timestamp month as an additional feature.

ROC curves are presented in Figure 10. We see that the time feature helps to improve the results, as does using all tags instead of just the spatially-intersected ones. The baseline (majority class) is 86.3%. It is interesting to examine the top few levels of the trained decision tree, to get a sense for which tags are most discrimina-

tive. The top decision node is “summer:” if this tag is present, then the photo is classified as not snow. If summer is not present, then the next few layers look at tags like “mountain,” “clouds,” “ski,” “geese,” and “egret.”

Machine learning to suppress false positives. Finally, we consider using the photo classifier as a filter while computing the likelihood ratios of Section 3.1, in order to reject photos that are marked with a snow tag but do not contain snow, using both visual and textual features. For the textual features, we use the decision tree classifier just described. For visual features, we trained an SVM using the GIST-like visual features described in Section 3.1, on the same hand-labeled dataset of about 2,000 images explained above. As with all other experiments, the training and testing sets were kept separate by training on data from 2007-2008 and testing on data from 2009-2010. For the photos in these latter two years, we use our decision tree to try to filter out false positives (photos tagged “snow” but not containing snow), and then re-compute the likelihood ratio confidence score. We find that using a classifier to reject false positives based on tags increased precision by nearly 10 percentage points, as shown in Figure 6(f): at 1% recall, precision increased from about 84% to about 93% for snow retrieval. For the visual features, we find a significant but more modest improvement, from about 84% to 86% at this level of recall.

4.3 Estimating vegetation cover

Another important measure of the ecological state of the planet is vegetation cover. We perform greenery versus no greenery predictions similarly to snow and no snow predictions using the Flickr confidence threshold method discussed in Section 3.1. As with snow, the ground truth is obtained from down-sampling and thresholding the NASA MODIS greenery data which has the same resolution as the snow cover data with similar coverage and quality (confidence) values. The Flickr greenery confidence values of bins are obtained in a similar way as with snow, except that we use a different set of target tags, including “tree,” “trees,” “leaf,” “leaves”, and “grass.” One important difference between the NASA greenery and snow datasets is that the greenery data is an average of daily observations spanning 16 days. Thus our goal is to predict the geospatial distribution of greenery for each 16-day period of the year.

We require a bin to have no less than 50% greenery coverage and above middle quality to be considered as a ground truth bin. We report experiments using two different definitions of non-green bins: those having less than 1% coverage, and those having less than 5% coverage. For the 50% and 1% threshold combination, 25.6% of the bins with ground truth are greenery bins, while for the 50% and 5% threshold combination, 15.8% of the bin with ground truth are greenery bins. As shown in Figure 11, both curves outperform a random baseline for greenery prediction, but the estimates are not as accurate as those observed during in snow predictions. There seem to be several reasons for this drop in performance. One is that the boundary between greenery and no greenery seems more vague than the snow/no-snow boundary. Moreover, the greenery ground truth data has a much coarser temporal resolution (16 days). Finally, it’s less clear which tags should be used to estimate greenery; using color analysis of the visual content of images may be a better approach, which we leave for future work. We also tried a learned classifier to predict greenery/non-greenery bins based on the set of tags used by all users in each bin and on each day. We used the LibLinear classifier [10] because it performed well in case of snow classification. Figure 12 presents the ROC curve for this classification task, showing an equal-error rate of about 91.6%.




Figure 11: Greenery precision-recall curve using two different ground truth thresholds.




Figure 12: ROC curve for classifying greenery of bins, using tag features and LibLinear classifier.

5. CONCLUSION AND FUTURE WORK

In this paper, we propose using the massive collections of user-generated photos uploaded to social sharing websites as a source of observational evidence about the world, and in particular as a way of estimating the presence of ecological phenomena. As a first step towards this long-term goal, we used a collection of 150 million geo-tagged, timestamped photos from Flickr to estimate snow cover and greenery, and compared these estimates to fine-grained ground truth collected by earth-observing satellites and ground stations. We compared several techniques for performing the estimation from noisy, biased data, including simple voting mechanisms and a Bayesian likelihood ratio. We also tested several possible improvements to these basic methods, including using temporal smoothing and machine learning to improve the accuracy of estimates. We found that while the recall is relatively low due to the sparsity of photos on any given day, the precision can be quite high, suggesting that mining from photo sharing websites could be a reliable source of observational data for ecological and other scientific research. In future work, we plan to study additional features including using more sophisticated computer vision techniques to analyze visual content. Also we plan to study a variety of other ecological phenomena, including those for which high quality ground truth is not available, such as migration patterns of wildlife and the distributions of blooming flowers.

6. ACKNOWLEDGMENTS

We thank Professor Michael Trosset for discussions on the linear regression models. This work was supported in part by the Lilly Endowment and by the Data to Insight Center at Indiana University, and used computational resources funded in part by NSF grant EIA-0202048 and by IBM.

7. REFERENCES

- [1] <http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/>.
- [2] The great sunflower project. <http://www.greatsunflower.org>.
- [3] Lost ladybug project. <http://www.lostladybug.org>.
- [4] A. Anagnostopoulos, R. Kumar, and M. Mahdian. Influence and correlation in social networks. In *KDD*, 2008.
- [5] J. Bollen, H. Mao, and X.-J. Zeng. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8, 2011.
- [6] D. Brockmann, L. Hufnagel, and T. Geisel. The scaling laws of human travel. *Nature*, 439:462–465, 2006.
- [7] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Trans. on Intelligent Systems and Technology*, 2:1–27, 2011.
- [8] D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri. Feedback effects between similarity and social influence in online communities. In *KDD*, 2008.
- [9] B. De Longueville, R. S. Smith, and G. Luraschi. "OMG, from here, i can see the flames!". In *Proc. Intl. Workshop on Location Based Social Networks*, pages 73–80, 2009.
- [10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: a library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [11] J. Ginsberg, M. Mohebbi, R. Patel, L. Brammer, M. Smolinski, and L. Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457:1012–1014, 2009.
- [12] D. K. Hall, G. A. Riggs, and V. V. Salomonson. MODIS/Terra Snow Cover Daily L3 Global 0.05Deg CMG V004. Boulder, CO, USA: National Snow and Ice Data Center, 2011, updated daily.
- [13] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
- [14] J. Hays and A. A. Efros. IM2GPS: Estimating geographic information from a single image. In *CVPR*, 2008.
- [15] R. Hunt. *Measuring Color*. London: Fountain Press, 1998.
- [16] X. Jin, A. Gallagher, L. Cao, J. Luo, and J. Han. The wisdom of social multimedia: Using Flickr for prediction and forecast. In *ACM Multimedia*, 2010.
- [17] G. John and P. Langley. Estimating continuous distributions in bayesian classifiers. In *Proc. Uncertainty in AI*, 1995.
- [18] K. Kremerskothen. <http://blog.flickr.net/en/2011/08/04/6000000000/>
- [19] Land Processes Distributed Active Archive Center. MODIS/Terra Vegetation Indices 16-Day L3 Global 0.05Deg CMG V005. Sioux Falls, SD: U.S. Geological Survey., 2011.
- [20] D. Lazer et al. Life in the network: the coming age of computational social science. *Science*, 323(5915):721–723, 2009.
- [21] D. Liben-Nowell and J. Kleinberg. Tracing information flow on a global scale using internet chain-letter data. *PNAS*, 12(105):4633–4638, 2008.
- [22] B. O'Connor, R. Balasubramanian, B. Routledge, and N. Smith. From tweets to polls: Linking text sentiment to public opinion time series. In *ICWSM*, 2010.
- [23] M. L. Parry, O. F. Canziani, J. P. Palutikof, P. J. van der Linden, and C. E. Hanson. *IPCC, 2007: Climate Change 2007: Impacts, Adaptation, and Vulnerability*. Cambridge University Press, 2007.
- [24] V. K. Singh, M. Gao, and R. Jain. Social pixels: genesis and evaluation. In *ACM Multimedia*, 2010.
- [25] M. Squires and J. Lawrimore. Development of an operational northeast snowfall impact scale. Technical report, NOAA National Climatic Data Center, 2006.
- [26] J. Su, H. Zhang, C. Ling, and S. Matwin. Discriminative parameter learning for bayesian networks. In *ICML*, 2008.
- [27] R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer, 2010.
- [28] A. Torralba. Contextual priming for object detection. *International Journal of Computer Vision*, 53(2):161–191, 2003.
- [29] I. Vergara, T. Norambuena, E. Ferrada, A. Slater, and F. Melo. StAR: a simple tool for the statistical comparison of ROC curves. *BMC Bioinformatics*, 9(1):265, 2008.