Verifying temporal data in geotagged images via sun Azimuth estimation

Sudha, N.; Kakar, Pravin

2012

Kakar, P., & Sudha, N. (2012). Verifying temporal data in geotagged images via sun Azimuth estimation. IEEE Transactions on Information Forensics and Security, 7(3), 1029-1039.

<https://hdl.handle.net/10356/95149>

<https://doi.org/10.1109/TIFS.2012.2188796>

© 2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [DOI: <http://dx.doi.org/10.1109/TIFS.2012.2188796>].

Downloaded on 26 Apr 2023 18:47:34 SGT

Verifying Temporal Data in Geotagged Images via Sun Azimuth Estimation

Pravin Kakar, *Student Member, IEEE* and N. Sudha, *Senior Member, IEEE*

Abstract—Image metadata provides useful information for applications such as image retrieval, content description and geolocation. However, it is relatively easy to tamper with this data using metadata manipulation tools. In this paper, we address the issue of authenticating the time of capture of an image based on its geolocation information. By utilizing a novel two-stage shadow detection process, we are able to estimate the azimuthal direction of the sun in an image and compare it against a calculated theoretical value to establish the authenticity of the time of capture. In case of ambiguous direction estimation from shadows, we improve an existing method using cues from vertical surfaces to resolve the ambiguity. The sky is used as a cue to perform verification in difficult scenarios such as the absence of shadows. Results are provided which show that our technique is able to estimate the sun azimuth with good accuracy. Additional results pertaining to date verification and camera direction estimation are also provided.

Index Terms—Time of capture, EXIF data, Image forensics

I. INTRODUCTION

IMAGE metadata is a useful source of data about the content of the image. It helps provide information for querying image databases, describing the image content and allowing one to establish the place of capture of the image, among myriad other uses. There are various formats in which this metadata may be stored in an image - Exchangeable Image File Format (EXIF), International Press Telecommunications Council's Information Interchange Model (IPTC-IIM) and Adobe Extensible Metadata Platform (XMP) [1]. These formats offer different functionalities for different workflows (image sharing, image copyright management, image processing, etc.) with some overlap. For instance, EXIF provides for accurate GPS coordinates storage for geolocation purposes, while a combination of IPTC-IIM and XMP allows for storing location information in a textual format (country, city, street, etc.).

In recent years, the number of cameras and smartphones that include GPS receivers has been growing. This enables recording the geographical information about the place where the image was captured, known as geotagging. Alternatively, information about mobile network coverage or wireless networks may be used to supplement or substitute for such location information. In certain images with distinctive content, the technique of [2] may be used to estimate the GPS

coordinates in an unsupervised manner. As mentioned earlier, this information is stored in the EXIF data in the image, and so, we only consider this metadata format in the remainder of this paper. While EXIF data can often provide detailed information about the image, it is quite easy to tamper with this data using freely available software packages such as ExifTool¹. Specifically considering geo-temporal data, it may be possible to verify geographical data visually by comparing the image content with data from earth-mapping software. However, visually verifying temporal data is not an easy task due to factors such as unfamiliarity with the location of capture of the image, non-uniform spacing of timezones, and periodic temporal shifts such as those introduced by daylight savings. How then can one authenticate the time of capture of an image?

Our work in this paper presents a solution to this problem. To the best of our knowledge, this is the first ever technique to ascertain temporal information from a single image. In particular, by utilizing the geolocation information (latitude and longitude of the location where the image was captured along with the orientation of the camera obtained from the "GPSLatitude", "GPSLongitude" and "GPSImgDirection" tags respectively in the EXIF data) and the date of capture, we are able to calculate the appropriate position of the sun in the sky relative to the camera. As this position varies with time, we estimate the direction of the sun from the content of the image such as shadows and the sky and verify the time of capture by comparing the closeness of the calculated and estimated directions.

Our major contributions in this paper are as follows:

- A new technique that uses the sun direction estimate and other EXIF data to verify time of capture, date of capture and camera direction of the image.
- A novel two-stage shadow detection technique to estimate the azimuthal direction of the sun.
- Improvement of an existing method to use vertical surfaces in determining the relative sun orientation.

The rest of this paper is organized as follows. In Section II, we discuss the existing related work in image forensics. Section III explains our method for estimating the sun azimuth direction, while Section IV discusses the calculation of the theoretical position of the sun. In Section V, we discuss the handling of cases where the estimated sun direction may be ambiguous or unreliable. We provide results from our technique in Section VI.

Copyright © 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

The authors are with the Institute for Media Innovation, Nanyang Technological University, Singapore - 637553. E-mails: pkakar@pmail.ntu.edu.sg, sudha@ntu.edu.sg

This work was supported by a Ph.D. scholarship awarded to Pravin Kakar by the Institute for Media Innovation.

¹<http://owl.phy.queensu.ca/~phil/exiftool/>

II. RELATED WORK

Traditional approaches to passive image forensics (e.g. [3]–[6]) have focused on detecting inconsistencies in image content in order to detect tampering in images. However, we are concerned with tampering of image metadata, not image content. It is entirely possible that the image content is left untouched, but the metadata is modified, rendering approaches similar to the above ones to declare the image (based on its content) as untampered.

Various approaches have been formulated for the purpose of camera forensics which could ostensibly be used to verify certain fields in the metadata of an image. Certain methods (e.g. [7], [8]) use demosaicking methods and sensor noise characteristics to identify source cameras. The authors of [9] use dust patterns on lenses to identify the lens used for capturing an image. Camera make, model, and lens characteristics are often recorded in image metadata. It may be possible to verify these using the above methods.

Image metadata has also been used directly for image forensics. The utility of EXIF headers in assisting detection of cases of child pornography was discussed in [10]. The authors of [11] use EXIF data along with other cues to identify source cameras and Photoshopped images. A technique presented in [12] treats the invariance of EXIF data in many image processing programs as a source of noise for camera fingerprint matching.

To the best of our knowledge, there has not been much work done in authenticating the temporal data of images, using image metadata or otherwise. An approach is described in [13] to estimate the time of creation of an audio or video recording using the electric network frequency criterion. Similarly, the authors of [14] use time lapse image sequences to model the temporal color changes on outdoor surfaces resulting from variations in sunlight to estimate geo-temporal information. Such approaches obviously cannot be extended to single images. The authors of [15] estimate the age of sensor defects from a database of images taken with the same camera to establish the date of capture of an image on the scale of months. Apart from the requirement of multiple images, the results are too coarse-grained to be useful in determining the time, rather than the date, of capture of an image.

The crux of our technique is the estimation of the azimuthal direction of the illuminating source (the sun, in this case). Within the realm of image forgery detection, approaches in [16], [17] use lighting inconsistencies to estimate the 2-D and 3-D lighting directions, but require human intervention in the form of occluding boundary selection or the creation of 3-D graphics models. Our work is inspired by the technique of [18] which used manual photogrammetric rectification of cast shadows, along with other cues, to resolve the issue of the first person to reach the North Pole. Also, the work of [19] uses shadows along with sky, vertical surface and pedestrian cues to estimate the sun elevation and azimuth angle. This estimate is then used for photorealistic insertion of a 3-D graphics object in the image. However, for verifying time of capture, it is sufficient to only know the sun azimuth angle, which can generally be determined by only using shadows.

We improve the use of the various cues in [19] in our work.

III. SUN AZIMUTH ESTIMATION

The outline of our technique is shown in Fig 1. We first apply various techniques for segmenting the ground plane, detecting ground shadows and rectifying the ground plane. This is followed by measuring the direction of ground shadow lines in order to estimate the sun azimuth direction. Our proposed method improves the reliability of the ground shadow detection process, especially in cases where the cast shadows are thin or weak.

A. Shadow Detection

In order to be able to estimate the sun direction, shadows provide a strong cue. As many structures in the world stand vertically on the ground, the shadows cast by these structures can be used to estimate the direction of the sun [19]. In order to enable this, the ground shadows need to be detected. Some approaches do exist for finding illumination source direction from cast shadows, but suffer from various limitations. The approaches in [20]–[23] require multiple images and/or human intervention in order to determine the direction of the illumination source. The technique in [24] assumes the surface on which the shadow is cast to be Lambertian, which is not a valid assumption for many ground textures like soil and grass [25]. In light of such limitations, we use a two-stage shadow detection process, detecting ground shadows with a technique designed specifically for that purpose [26], and applying a new method developed by us (Section III-A2) in the cases where this technique fails to detect shadows reliably.

1) *Detecting Ground Shadow Lines*: We begin by segmenting the image into ground, vertical objects and sky regions using the technique developed in [27]. Following this segmentation, the technique of [26] is applied to the ground region in order to localize ground shadows.

Long, straight lines in the shadows are detected using the method presented in [28]. For a detected shadow edge consisting of pixels (x_i, y_i) , a connected components algorithm based on the gradient directions of the edge pixels is applied. Components with a significant length ($>5\%$ of image size) are chosen for further analysis. For the matrix D of the line support region,

$$D = \begin{bmatrix} \sum_i \tilde{x}_i^2 & \sum_i \tilde{x}_i \tilde{y}_i \\ \sum_i \tilde{x}_i \tilde{y}_i & \sum_i \tilde{y}_i^2 \end{bmatrix} \quad (1)$$

the eigenvalues λ_1 and λ_2 are calculated. $\tilde{x}_i = x_i - \bar{x}$ where $\bar{x} = \frac{1}{n} \sum_i x_i$. $\tilde{y}_i$ are defined similarly. The ratio $\frac{\lambda_1}{\lambda_2}$ provides the quality of fit of the line. If this ratio is above a threshold proportional to the length of the connected component, the component is declared to be a line.

In order to determine the angle of the shadow lines on the ground, metric rectification is performed using information about the horizon line l_∞ in the image and the focal length f_c of the camera as described in [29]. The latter is obtained from the EXIF data of the image, and the former is chosen to be halfway between the lowest pixel of the sky segment and the highest pixel of the ground segment obtained earlier.

Fig. 1: Flowchart of our technique

This assumption is a useful approximation for most consumer images. As the horizon line is assumed to be horizontal, the equation of l_∞ only has a y coefficient l_y . Solving for the first component of the imaged circular point $I_1 = \alpha - i\beta$ in

$$I_1^2 + 2l_y u_0 I_1 + 2l_y v_0 + l_y^2(u_0^2 + v_0^2 + f_c^2) + 1 = 0 \quad (2)$$

with (u_0, v_0) as the image coordinates of the principal point of the camera, the rectification matrix N is formed as:

$$N = \begin{bmatrix} \frac{1}{\beta} & -\frac{\alpha}{\beta} & 0 \\ 0 & 1 & 0 \\ 0 & l_y & 1 \end{bmatrix} \quad (3)$$

The detected ground shadow lines are warped according to N in order to determine their orientations relative to the camera. The results from various steps of the above process are shown in Fig 2. In all images in this paper, dashed lines indicate the estimated horizon line and solid lines indicate detected shadow lines. All images are taken from the popular photo-sharing website Flickr², unless otherwise stated. Note that only the dominant (most likely) vertical surface segmentation is shown in Fig 2(b). The segmentation used for determining l_∞ consists of all the possible ground and sky pixels, not just the dominant ones.

2) *Proposed Shadow Detection Technique*: In cases where shadows are very thin or weak, the technique of [26] may fail to perform reliably. If some shadows, but no shadow lines, are detected, we apply an intensity-based technique to detect these thin and weak shadows. First, we convert the image to the HSV colorspace and examine the ground region's V-channel for regions containing shadows. We assume that the ground surface has regions consisting of three major types of intensity values - dark regions caused by shadows and occlusions, the majority of regions being the actual ground surface with its

Fig. 2: Various steps in determining shadow line orientation. (a) Original image. (b) Region segmentation. Green = ground, red = vertical surface, blue = sky. (c) Shadow line detection. Dashed line indicates the horizon line. (d) Shadow lines after rectification.

inherent intensity and bright regions such as road markings. Let the set of intensity values of the ground pixels (i_g) be $\mathcal{V}_g = \{i_g | 0 \leq i_g \leq 255\}$. A threshold τ is defined as:

$$\tau = \begin{cases} 30 & \text{if } Mo(\mathcal{V}_g) > 150 \\ 85 & \text{otherwise} \end{cases} \quad (4)$$

where Mo is the mode of the set of values. The above values are determined empirically from a small training set of fifteen images containing shadows. The images in this dataset are

²<http://www.flickr.com>

Fig. 3: Detecting very thin shadows. (a) Original image. (b) Shadow detection with [26]. (c) Shadow detection with our method. (d) Shadow line detection.

Fig. 4: Sun elevation θ_s and azimuth ϕ_s .

different from the ones used for evaluating our technique in Section VI. Denoting the τ th percentile element in $\mathcal{V}_g$ by $i_{g\tau}$,

$$i_g = i_{g\tau} \quad \forall \quad i_g > i_{g\tau} \quad (5)$$

The reasoning for this is as follows. In the process of detecting shadow edges, higher intensity textures, such as road markings, can give rise to false positives. By eliminating higher-intensity pixels (5), we reduce the chance of these false positives occurring. By analyzing $Mo(\mathcal{V}_g)$ (4), we can compensate for the difference in the inherent intensity of the surface across materials. Surfaces made of darker material, with a lower mode, may not have as great a contrast with shadowed regions as brighter surfaces. By raising τ for lower modes, we ensure that no shadows are falsely eliminated.

Segmenting using Otsu's method [30] and performing morphological cleaning, we detect regions where shadows are likely to lie. Applying a Canny edge detector [31] with the shadow regions as a mask, we find the shadow boundaries, and proceed with detecting shadow lines as above. Our method has a higher likelihood of false positives due to analyzing only intensity values. Nevertheless, we have found the detected shadows to outweigh the influence of false positives in practice. An example of a case for using this method is shown in Fig 3.

B. Estimating Sun Azimuth Direction

The position of the sun is expressed in terms of its elevation θ_s and azimuth ϕ_s (Fig 4). The elevation of the sun is the

angle it makes with the horizontal ground plane. The azimuth of the sun is the angle its projection on the ground plane makes with the line in the camera direction. In the rest of this paper, we refer to the sun azimuth direction as the sun direction, for convenience. Mathematically, each shadow line provides an estimate of the sun direction as follows. Consider a ground shadow line $l_i \in \mathcal{G}$, where $\mathcal{G}$ represents the set of all ground shadow lines. Let its orientation on the ground plane after rectification be α_i , and the sun direction be ϕ_s . Then, the angle between the shadow line orientation and the sun direction is

$$\angle(\alpha_i, \phi_s) = \min\{\angle(\alpha_i, \phi_s), \angle(\alpha_i + 180^\circ, \phi_s)\} \quad (6)$$

For a line l_i , one does not generally know which end (if any) of l_i the object casting the shadow line is located. This gives rise to the ambiguity in sun direction with two possible values. The direction estimate of each shadow line l_i , is modeled as the maximum of two normal distributions, denoted by $\mathcal{N}$:

$$P(\phi_s | \alpha_i) = \max_{\phi_s} \{\mathcal{N}(\alpha_i, \sigma_g^2), \mathcal{N}(\alpha_i + 180^\circ, \sigma_g^2)\} \quad (7)$$

where σ_g reflects the uncertainty in the accuracy of the estimate. We improve the direction estimation process by observing that the length of each shadow line plays an important role in determining the reliability of the direction estimated through it. A longer shadow line is less likely to be a false detection resulting from aberrations (potholes on roads, rocky soil, etc.) on the ground, and so, its estimate is more likely to be the correct one. Hence, we sum the preferred sun directions for the shadow lines in a final weighted estimate:

$$P(\phi_s | \mathcal{G}) \propto \sum_{l_i \in \mathcal{G}} |l_i| P(\phi_s | \alpha_i) \quad (8)$$

where $|\cdot|$ denotes length. In comparing the estimated direction with the theoretical value, we ignore the one which cannot possibly lie on the sun path, as discussed below. Failing that, we employ cues from vertical surfaces in the image in order to pick the likelier direction.

IV. THEORETICAL SUN POSITION CALCULATION

For an observer on the earth's surface, the sun appears to be in motion, rising in the east each day, and setting in the west. However, additional factors make this apparent motion of the sun far from uniform, adding to the complexity of estimating time from the sun position or vice versa.

A. Solar Declination

In addition to the east-west motion, the sun also has an apparent annual motion, resulting from an almost constant tilt in the earth's axis of rotation with respect to the plane of its orbit. As the earth orbits the sun, this constant tilt results in the path of the sun in the sky moving along the north-south axis over a year. This annual motion is responsible for phenomena such as the occurrence of seasons, and manifests itself in a phenomenon known as declination, denoted by δ . This is defined as the angular distance of the sun from the equatorial plane. The absolute value of δ is maximum at the solstices in June and December, and zero at the equinoxes in March and September.

Fig. 5: Polar sun plots for (a) Singapore. (b) Barcelona, Spain. Blue curves indicate path of sun on the 21st of each month from June to December. Red curves indicate hourly intervals of equal solar time. Green line indicates north in an overhead view.

B. Equation of Time

In addition to the declination of the sun, a phenomenon known as the equation of time, t_e , causes variations in the time estimated from the position of the sun. This results from two contributing factors - the earth's eccentric orbit around the sun, and the tilt of the earth's axis to its orbital plane. As the earth orbits the sun in an elliptical orbit with an eccentricity of about 0.017, in accordance with Kepler's laws of motion, the speed of the earth varies slightly around its orbit, causing variation in the position of the sun (ignoring the effect of δ) at the same time of the day according to a clock. The earth's axial tilt not only causes the apparent annual motion, it also causes day-to-day shifts in the position of the sun, when projected on to a 2-D surface, to be non-uniform. Tables³ of the values of δ and t_e are published by various astronomical organizations.

C. Calculating Sun Position

The position of the sun depends on all the above factors. Let the apparent path of the sun across the sky lie on the inner surface of a unit sphere, with the observer on the earth's surface at its center. Consider a right-handed Cartesian coordinate system with the positive x -axis lying along north on the ground plane. Then, the position of the sun can be calculated as [32]:

$$x = \sin(\delta) \cos(\theta) - \cos(\lambda) \sin(\theta) \cos(\delta) \quad (9)$$

$$y = \sin(\lambda) \cos(\delta) \quad (10)$$

$$z = \sin(\delta) \sin(\theta) + \cos(\lambda) \cos(\theta) \cos(\delta) \quad (11)$$

where θ is the latitude of the observer's location on the earth's surface, and λ is the sun hour angle, defined as:

$$\lambda = (t_s - 12) \cdot 15^\circ \quad (12)$$

where t_s is the solar time (time according to the sun) in hours for the given location. It is related to clock time or local time t_l as:

$$t_s = t_l - t_o + \frac{\gamma}{15^\circ} + t_e \quad (13)$$

where t_o is the time offset from universal coordinated time (UTC) due to timezones and daylight savings, and γ is the longitude of the observer's location on the earth's surface.

By considering values of λ for which $z > 0$ in (11), the visible path of the sun in the sky for any day and location

on the earth can be determined. Polar sun plots for various locations and days are shown in Fig 5. These plots show the sun path in an overhead view of the location, with the vertical upward direction being the viewing direction of the camera. For Fig 5, this direction is assumed to be north. The sun always lies along the north-south axis at solar noon, providing a useful reference for interpreting these plots.

V. POSTPROCESSING

In this section, we examine the cases that generate undesirable results in our technique and propose solutions for them. Specifically, we discuss the cases where the sun azimuth direction indicates multiple positions on the visible sun path, and cases where shadows not cast by vertical surfaces introduce large errors in our technique.

A. Determining Sun Direction from Vertical Surfaces

As discussed in Section III, shadows provide information about two directions either one of which may correspond to the sun azimuth direction. However, depending on the location and time of capture of the image, only one of these time data may be permissible. We determine the presence of ambiguity by considering the intersection of both directions as determined by the shadow estimation process, with the visible sun path. If there are two intersections, then we resolve the ambiguity by using cues from the vertical surfaces in the image. Although this ambiguity can be resolved by asking the user to choose the correct direction, we avoid human intervention by only utilizing the content of the image. As an example, an image captured in the early morning (Fig 6(a)) gives rise to shadows the direction of which can be caused by multiple positions of the sun on the given path. However, an image captured closer to noon (Fig 6(c)) rarely suffers from such issues.

The technique of [27], used earlier to obtain the ground/vertical surface/sky segmentation, classifies the vertical surface region of the image into 5 subcategories - front-facing, left-facing, right-facing, non-planar solid and porous. Although the first three subcategories are useful for estimating illumination direction, the vertical surfaces are not always assigned correctly to these subcategories (Fig 7). The work of [19] also uses these vertical surface cues for determining the sun direction. However, we observe that the method attempts to compute the sun direction from only 3 coarse, and often unreliable, orientations. Instead, we only utilize the vertical surface cues to resolve any directional ambiguity that may exist. This only requires correctly estimating the sun direction to within $\pm 90^\circ$.

A logistical intensity classifier is used to classify whether the sun is in front of or behind a surface of a given orientation:

$$P(\angle(\beta_i, \phi_s) < 90^\circ | b_i) = \frac{1}{1 + e^{-(x_1 + x_2 b_i)}} \quad (14)$$

In the above equation, b_i is the mean intensity of the vertical surface $w_i \in \mathcal{V}, i \in \{\text{front, left, right}\}$. β_i is the normal surface orientation of w_i , and x_1, x_2 are parameters learned from a database used in [27] with manually labeled sun azimuth directions. Based on the above classification for each

³<http://freepages.pavilion.net/users/aghelyar/sundat.htm>

Fig. 6: Ambiguity in shadow direction. (a) Image containing shadows with ambiguous direction. (b) Polar plot for (a). (c) Image containing shadows with unambiguous direction. (d) Polar plot for (c). Red circle indicates calculated sun position. Solid magenta line indicates correct direction of sun from shadows; dashed magenta line indicates the incorrect one. Dashed blue line indicates the sun direction estimate from vertical surface cues only, where applicable. Black circles indicate hourly intervals along the green sun path. Straight green line indicates north.

Fig. 7: Segmentation of vertical surfaces into left-facing, front-facing and right-facing.

surface orientation, the probability of the sun direction is determined as:

$$P(\phi_s|w_i) = \begin{cases} \mathcal{N}(\beta_i, \sigma_w^2) & \text{if } (14) \geq 0.5 \\ \mathcal{N}(\beta_i + 180^\circ, \sigma_w^2) & \text{if } (14) < 0.5 \end{cases} \quad (15)$$

where σ_w is used to equate the fractional mass of the Gaussian $\pm 90^\circ$ in front of w_i to the value determined from (14). Then, the probabilistic directions obtained from the 3 surface orientations are combined according to:

$$P(\phi_s|\mathcal{V}) \propto \sum_{w_i \in \mathcal{V}} P(\phi_s|w_i) \quad (16)$$

In order to deal with occlusions on the vertical surfaces, the original technique uses a k -means classifier with $k=2$ to pick the brighter intensity. However, while such an intensity selection procedure works well for uniform surfaces, real-world surfaces are far from uniform, containing variations in texture due to windows, different building materials, environmental degradation, etc. Moreover, erroneous segmentation of the horizontal ground as a vertical surface in some images further compounds this problem. Therefore, we mitigate this issue by considering the mean intensity of a surface to be that of the *largest* cluster, rather than the *brightest* cluster. This

Fig. 8: Using sky cue to verify sun position. (a) Image with poor sun direction estimate. (b) Polar plot showing probability of sun location according to sky cue. White indicates maximum probability. (c) Polar plot with the top 10% probability values overlaid with calculated sun position (red), likeliest sun position from sky cue (blue), visible sun path (green) and north indication (straight green line).

ensures that a small bright patch on an otherwise dark vertical surface does not result in wrong classification. As an example, this modification allows us to pick the correct sun direction for the image in Fig 6(a), while the method of [19] fails to do so.

B. Verification of Sun Position using Sky Modeling

In some cases, the estimated sun azimuth direction may differ very widely from the calculated value. In theory, this should be an indicator of inconsistency in the image metadata. However, we acknowledge that a wrong direction estimate can also result from shadows cast by horizontal surfaces, complex-shaped shadows or the complete absence of shadows. In order to deal with this, we turn to the sky in order to obtain information about the position of the sun.

By fitting the sky pixel intensities to models generated from various discretized values of sun elevation and azimuth in [33], the likely position of the sun in the sky can be found. The probability of a particular sun elevation θ_s and azimuth ϕ_s is determined from the sky pixel intensities $s_i \in \mathcal{S}$ as follows:

$$P(\theta_s, \phi_s|\mathcal{S}) \propto \exp \left(\sum_{s_i \in \mathcal{S}} \frac{-(s_i - kg(\theta_s, \phi_s, u_i, v_i, f_c))^2}{2\sigma_s^2} \right) \quad (17)$$

where k is a scale factor, $g(\cdot)$ is the sky model of [33], (u_i, v_i) are the image coordinates of s_i , f_c is the focal length, and σ_s reflects the uncertainty of the model. However, this estimate is generally quite rough, frequently assigning a high probability of the sun being located in a quadrant or more of the sky.

The authors of [19] combine this cue with shadow and vertical surface cues in order to obtain an estimate of the sun position. However, the sky estimate is usually not robust enough to give a reliable estimate independently. In the absence of other cues, instead of determining the sun azimuth direction and thereby verifying the calculated sun position according to the closeness of the estimate as we have done so far, we verify the *likelihood* of the sun being in a particular position in the sky, even though we cannot usually obtain an estimate of the sun position from the sky cue alone.

In order to accomplish this, we perform a simple procedure for problematic images like the one in Fig 8(a). After generating the sky probability map (Fig 8(b)), we consider the points lying above the 90th percentile in the range of probabilities.

Fig. 9: Estimating the sun direction by using different cues independently. (a) Image. (b) Polar sun plot. Red = calculated sun position, blue = estimated sun position from shadows, magenta = estimated sun position from vertical surfaces, amber = estimated sun position from sky. Black circles indicate hourly intervals along the green sun path. Straight green line indicates north.

These points correspond to highly likely positions of the sun, and they do not necessarily form a single cluster (which is another cause for unreliable estimates). If the calculated position of the sun lies among these points (Fig 8(c)), then it is likely that the temporal metadata of the image is reliable. It is interesting to note that we utilize both the elevation and azimuth in the verification from the sky cue, thus preventing the likely regions where the sun may be present from growing too large.

To sum up, shadow cues alone generally provide much more accurate estimates of the time of capture than vertical surfaces or the sky alone. An example of this is shown in Fig 9. We employ the latter two to resolve any ambiguities that may arise in estimating the time of capture of the image from shadow cues alone.

VI. RESULTS

In order to calculate the theoretical position of the sun from the EXIF data of an image, we process it as follows. The latitude θ and longitude γ of the location where the image was captured are obtained from the GPS information in the EXIF data. The date and time of capture are also present in the EXIF data, although the timezone information t_o is often missing. Hence, we combine information from the GeoNames⁴ and tz⁵ databases to find out the timezone and applicability of daylight savings for the date and location of the image. After calculating the position of the sun by applying (9)-(13), we use the camera direction present in the GPS information to rotate the polar plot so that the overhead view has the camera direction pointing up. The camera direction may be recorded with reference to the geographic north pole or the magnetic north pole. If it is recorded with respect to the latter, then we compensate for the temporal variation in the earth's magnetic field (known as magnetic declination) by employing data from the World Magnetic Model⁶.

For convenience, we represent the calculated and estimated sun directions as ϕ_s and $\hat{\phi}_s$ respectively. We also use $\Delta\phi_s = |\phi_s - \hat{\phi}_s|$ to denote the absolute error. Similarly, t_s , $\hat{t}_s$ and Δt_s indicate the above quantities for time of capture.

⁴<http://www.geonames.org/>

⁵<http://cs.ucla.edu/~eggert/tz-link.htm>

⁶<http://www.ngdc.noaa.gov/geomag/>

Fig. 11: Performance of proposed technique with variation in (a) Date of capture. Dotted lines indicate errors with random date of capture. (b) Latitude of capture. (c) Longitude of capture. Red curves indicate mean; blue curves indicate median.

A. Verifying Time of Capture

Results obtained by applying our technique to various images to verify t_s are shown in Fig 10. It can be seen that, knowing the GPS information and date of capture, $\hat{t}_s$ is quite accurate. This is achieved despite simplifying assumptions in the rectification process. In cases where $\hat{\phi}_s$ has two possible values, the ambiguity is resolved by using the vertical surface cue as discussed in Section V-A.

B. Results on a Database

In order to evaluate the performance of our technique, we constructed a database of 100 images from consumer images taken from Flickr. We ensured that the images were outdoor images, contained geolocation information, and had discernible shadows. The images have been captured all over the world, and with various cameras ranging from smartphone cameras to professional DSLRs. We analyzed the performance of our technique by introducing a variety of offsets in the day of capture, and the GPS coordinates of the image. The results of these operations are shown in Fig 11.

1) *Ambiguity in Date of Capture:* We first tested our technique in the presence of ambiguity in the date of capture. The results for this evaluation are shown in Fig 11(a). With

Fig. 10: Verifying time of capture. Top row: Images; Bottom row: Corresponding polar sun plots. Red circle indicates calculated sun position; Blue circle indicates estimated sun position. Black circles indicate hourly intervals along green sun path. Straight green line indicates north.

no ambiguity, our technique has a mean Δt_s of merely 27.7 minutes and a median Δt_s of 27.3 minutes. It can be seen that even with an offset of ± 45 days (a range of approximately 3 months), the median Δt_s of our technique is only about 37 minutes. The large mean Δt_s with positive offsets is due to certain outliers. In these outlying cases, with variation in the visible sun path with the day offset, the multiple possible $\hat{\phi}_s$ are reduced to a single (incorrect) $\hat{\phi}_s$, as the correct intersection of $\hat{\phi}_s$ and the visible sun path is lost. However, the median Δt_s is robust to such outliers and displays the efficacy of our technique. In the case that the date of capture is completely unknown, we tested the performance of our technique by randomly assigning a date of capture for each image in our database. The errors from such random assignment are denoted by the dotted lines in Fig 11(a). As can be seen, the median Δt_s is only about 38 minutes, whereas the mean Δt_s is close to 72 minutes due to similar outliers as above. We observe however that it is often possible to estimate a rough month of capture from the image content (e.g. from seasonal environmental changes) and in the case of smartphone cameras, the date is usually obtained from the mobile network and is quite accurate.

2) *Ambiguity in Latitude of Capture:* Next, our technique is evaluated with variation in the latitude θ of the location of capture. We add or subtract offsets to $|\theta|$. This allows positive offsets to move the location away from the equator, and negative offsets towards it (except for locations close to the equator, where this may be reversed). The results of this analysis are shown in Fig 11(b). Note that Δt_s are lower for negative offsets (moving towards the equator, in general) than positive offsets. This is because Δt_s along the sun path varies less for a given $\Delta\phi_s$ over a significant part of the day as $|\theta|$ moves closer to the equator. Also, the maximum offset of 15° that we tested our technique with corresponds to a distance of about 1,667 km. Our technique has a Δt_s of less than an hour even with this large variation in distance around the location of capture.

3) *Ambiguity in Longitude of Capture:* Finally, we test our technique against variation in longitude γ . Any variations in t_o that may occur due to a change in γ are ignored. It can

TABLE I: Performance of proposed weighted sum of sun direction probabilities. All values are in minutes.

	Weighted Estimates (Proposed)	Equal Estimates [19]
Mean Error	27.7	32.4
Median Error	27.3	30.1

be seen from Fig 11(c) that apart from close to the correct longitude, the mean and median Δt_s increase linearly, due to the γ term in (13).

4) *Effect of Length of Shadow Lines:* We test the performance of our proposed modification in combining the estimated sun azimuth directions obtained from multiple shadow lines (8). We have weighted the probability estimates of each shadow line by its length, instead of considering all estimates equal as in [19]. The results of the experiment are shown in Table I. As can be seen, considering the length of the shadow lines improves the combined $\hat{\phi}_s$.

5) *Sun Orientation Relative to Vertical Surfaces:* We also use our database of images in order to provide evidence for our proposed idea of using the largest cluster rather than the brightest cluster in Section V-A. As can be seen from Table II, the performance with our modification in correctly determining ϕ_s to $\Delta\phi_s = 90^\circ$ is comparable to that of the original technique. In our database, out of 100 images, there are 21 in which ϕ_s gives rise to two possible positions of the sun. When dealing with such images where $\hat{\phi}_s$ is ambiguous, using the largest cluster as proposed by us outperforms using the brightest cluster considerably. Note that such ambiguous cases are exactly where we are using the vertical surface cue, rendering the superior performance of our technique even more relevant.

The reason for the better performance of our technique is as follows. As mentioned earlier, ambiguity in $\hat{\phi}_s$ usually arises in images captured in the early morning or late afternoon. This is also the time when the sun is closer to the ground. This results in the horizontal component of sunlight incident normally on a vertical surface to be greater. As a result, vertical surfaces tend to be more strongly bright or dark as compared to other times of the day. In the case of the original technique, a small bright

TABLE II: Performance of vertical surface cue on database

	Incorrect images	
	Largest cluster (Proposed)	Brightest cluster [19]
All images	41/100	46/100
Ambiguous images	3/21	11/21

Fig. 12: Verifying date of capture. (a) Image to be verified. (b) Polar sun plot for image time of capture and location. Red curve indicates path of sun at image time of capture over the year. Blue circles indicate estimated sun positions on sun path. Green straight line indicates north. (c) Magnified view.

aberration in uniformity on the vertical surface is sufficient to result in a wrong estimate, whereas we mitigate this effect by using the largest cluster.

C. Other Applications

In this section, we discuss other possible applications of our technique. If t_s is known, then the date of capture can be verified. Similarly, if the direction of the camera is unknown, it may be determined from other parameters.

1) *Verifying date of capture:* If t_s and GPS information (θ, γ , and camera direction) for an image are known, then the date when it was captured can be verified, as shown in Fig 12. The red path, known as an analemma, is obtained by calculating the sun position at t_s over all days of the year. The characteristic shape of the path is due to the combined effects of δ and t_e . By finding the intersection of a line oriented at $\hat{\phi}_s$ with the analemma, the dates on which the sun could have been in that position can be determined.

Up to four dates are obtained by the above intersection depending on t_s and the camera direction. Hence, this method is suitable for verifying, not determining, the date of capture. This can find application in legal situations where such information needs to be authenticated. For Fig 12, the estimated dates are May 24 and July 12. The image was captured on July 9 (according to the EXIF data), which is an error of merely 3 days. Note that as the sun path is far less variable in this case as compared to the daily motion, the errors generally tend to be larger than that for time verification. Nevertheless,

Fig. 13: Determining camera direction. (a) Image with wrong camera direction in its metadata. (b) Polar sun plot. (c) View of location in wrong direction. (d) View of location in corrected direction. (e) Sun probability map according to sky cue for (a). (c) and (d) courtesy of Google Street View.

in general, our method offers a reliable way of authenticating the approximate date of capture.

2) *Determining camera direction:* Standalone and smart-phone cameras use digital compasses to determine the camera direction when an image is captured. However, these may not be reliable in the presence of large metallic objects or when the camera orientation is varied. If the time and date of capture of the image, along with its GPS coordinates (θ, γ) are known, then the correct image direction can be determined.

Consider the image in Fig 13(a). The EXIF data purports the camera to have been facing a direction 264° clockwise from north. From Fig 13(b), ϕ_s (from image metadata) is 8.25° from the camera direction, while $\hat{\phi}_s$ (from image content) is 271° from the camera direction, both clockwise. The street view from the reported camera direction is shown in Fig 13(c). Applying a counterclockwise correction of $(271-8.25)^\circ$ to the camera direction as calculated by our technique gives a view as in Fig 13(d) which matches closely to that of the original image. Such camera direction correction has applications in areas such as registration of geotagged images and accident investigations.

It is important to note that an error in the recorded camera direction gives a large error between the estimated and calculated sun directions, similar to the scenario in Section V-B. However, the sky probability map (Fig 13(e)) clearly shows that it is unlikely that the sun could be in the calculated position, thus rendering an error in $\hat{\phi}_s$ far less probable.

TABLE III: Performance of proposed technique in challenging cases. All values are in minutes.

Case	Mean Δt_s	Median Δt_s
Non-vertical shadows	155.6	141.5
Tilted horizon line	126.5	108.7
Tilted ground plane	73.3	55.8
Displaced horizon line	70.2	50.8
Unknown focal length	68.7	50.0

D. Discussion

In this section, we discuss the performance of our technique on certain challenging cases which may sometimes occur.

Fig 14(a) shows that our proposed technique provides an incorrect ϕ_s from shadows when the majority of the cast shadows are from horizontal objects. Note, however, that $\Delta\phi_s$ is about 90° which indicates that our technique works well apart from the assumption of the vertical nature of the shadow-casting objects. We tested the performance of our technique on our database by randomly changing the orientation of each detected shadow line by up to $\pm 90^\circ$. As seen in the first row of Table III, the Δt_s is relatively large as the structural dependence of the shadows is ignored in this case. Minimal user interaction is likely to suffice in mitigating this issue.

Fig 14(b) presents a case where a horizontal horizon line is a poor approximation of the actual horizon line. Unsurprisingly, $\Delta\phi_s$ is relatively large. For testing on our database, we rotated each image by a random value of up to $\pm 30^\circ$. As expected, the values of Δt_s in the second row of Table III are quite large at close to two hours. We will investigate the handling of such cases, by possibly analyzing the dominant orientations of edges in the image, in our future work.

Figs 14(c) and (d) respectively show cases where the ground surface is not horizontal and where the estimated horizon line is displaced quite far away (>100 pixels) from the actual horizon line. Nevertheless, our technique is able to deal with such problems fairly well, and Δt_s is reasonably low. In our database, we simulated a tilted ground plane by rotating the cast ground shadows by a random value of up to $\pm 30^\circ$. For checking the performance with a displaced horizon line, we introduced a random error of up to ± 200 pixels in the position of the estimated horizon line. As seen in the third and fourth rows of Table III, the performance is reasonably good with a median Δt_s of less than an hour.

Finally, Fig 14(e) presents an example of the performance of our technique in case of missing f_c information. For the camera used (Sony DSC-HX5V), the range of possible f_c values is 4.25 mm - 42.5 mm. For a non-telephoto image such as this one, the f_c is likely to be at the lower end of this range. The ϕ_s for the two limits of the range, as well as for the actual f_c (5.09 mm) show that our technique performs quite well even with some ambiguity about the actual f_c value. Note that such ambiguity is usually not an issue for smartphone cameras which have fixed focal length lenses. In our database, we randomly changed the f_c values of images taken with variable focal length lenses by up to $10\times$ the f_c stated in the image metadata. Although this can result in a f_c greater than possible with the lens used, our technique performs quite well as seen in the last row of Table III. The median Δt_s of

our technique is only 50 minutes in this case.

VII. CONCLUSION

Image metadata in the form of EXIF data contains GPS information about the location where the image was captured in the case of geotagged images. We use this information in combination with information about the sun azimuthal direction determined from shadows in order to verify the time/date of capture of an image in our work. Also, we are able to determine the camera direction from the geo-temporal information present in the image metadata.

We have detected ground shadows using an existing technique, and applied a novel technique in cases where the former technique fails. We have estimated the direction of the sun based on straight ground shadow lines with longer lines providing stronger estimates. This estimated direction is then compared against a theoretical value calculated taking into account solar declination, equation of time and timezone offset effects.

By comparing the theoretical and estimated directions, we are able to verify the time of capture of the image quite accurately. In case of ambiguity in the estimated azimuthal direction or in the absence of shadows, we have exploited secondary cues from vertical surfaces and the sky to perform verification.

We have tested our technique on a database of images, and shown that it performs effectively even with ambiguity in the date of capture and the geolocation information of the image. Additionally, we have provided results showing the applicability of our method in verifying the date of capture and determining the camera direction using results from astronomy.

REFERENCES

- [1] Metadata Working Group, "Guidelines for Handling Image Metadata v2.0," November 2010. [Online]. Available: <http://www.metadataworkinggroup.org/specs/>
- [2] J. Hays and A. Efros, "IM2GPS: Estimating geographic information from a single image," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2008, pp. 1-8.
- [3] J. Fridrich, D. Soukal, and J. Lukáš, "Detection of copy-move forgery in digital images," in *Proc. Digital Forensic Research Workshop*, 2003, pp. 90-105.
- [4] M. Johnson and H. Farid, "Exposing Digital Forgeries through Chromatic Aberration," in *Proc. ACM Workshop on Multimedia and Security*, 2006, pp. 48-55.
- [5] M. Kirchner and J. Fridrich, "On detection of median filtering in digital images," in *Media Forensics and Security II*, vol. 7541. SPIE, 2010, pp. 754 110-1-754 110-12.
- [6] T.-T. Ng, S.-F. Chang, C.-Y. Lin, and Q. Sun, "Passive-blind image forensics," in *Multimedia Security Technologies for Digital Rights*. Elsevier, 2006, pp. 383-412.
- [7] S. Bayram, H. T. Sencar, and N. Memon, "Improvements on Source Camera Model Identification Based on CFA Interpolation," in *Int. Conf. on Digital Forensics*, 2006.
- [8] J. Lukas, J. Fridrich, and M. Goljan, "Digital Camera Identification from Sensor Noise," *IEEE Trans. Inf. Forensics Security*, vol. 1, no. 2, pp. 205-214, 2006.
- [9] A. E. Dirik, H. T. Sencar, and N. Memon, "Digital Single Lens Reflex Camera Identification From Traces of Sensor Dust," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 539-552, 2008.
- [10] P. Alvarez, "Using Extended File Information (EXIF) File Headers in Digital Evidence Analysis," *Int. J. of Digital Evidence*, vol. 2, no. 3, pp. 1-5, 2004.

Fig. 14: Challenging cases for our proposed technique. Top row: Images. Bottom row: Corresponding polar sun plots. (a) Shadows cast by horizontal surfaces. (b) Tilted horizon line. (c) Tilted ground surface. (d) Poor estimation of horizon line. (e) Missing focal length information.

- [11] E. Kee, M. K. Johnson, and H. Farid, "Digital Image Authentication from JPEG Headers," *IEEE Trans. Inf. Forensics Security*, vol. 6, no. 3, pp. 1066–1075, 2011.
- [12] B. Mahdian, S. Saic, and R. Nedbal, "JPEG Quantization Tables Forensics: A Statistical Approach," in *Int. Workshop on Computational Forensics*, 2010, pp. 150–159.
- [13] C. Grigoros, "Applications of ENF Criterion in Forensic Audio, Video, Computer and Telecommunication Analysis," *Forensic Science International*, vol. 167, no. 2-3, pp. 136–145, 2007.
- [14] K. Sunkavalli, F. Romeiro, W. Matusik, T. Zickler, and H. Pfister, "What do color changes reveal about an outdoor scene?" in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [15] M. Goljan and J. Fridrich, "Determining Approximate Age of Digital Images Using Sensor Defects," in *SPIE Conf. on Media Watermarking, Security, and Forensics*, 2011.
- [16] M. K. Johnson and H. Farid, "Exposing Digital Forgeries in Complex Lighting Environments," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 2, pp. 450–461, 2007.
- [17] E. Kee and H. Farid, "Exposing Digital Forgeries from 3-D Lighting Environments," in *IEEE Int. Workshop on Information Forensics and Security*, 2010, pp. 1–6.
- [18] T. D. Davies, "Photographs prove Peary was close to the Pole," *National Geographic Mag.*, vol. 1, pp. 52–53, 1990.
- [19] J.-F. Lalonde, "Understanding and Recreating Visual Appearance Under Natural Illumination," Ph.D. dissertation, Carnegie Mellon University, January 2011.
- [20] I. Sato, Y. Sato, and K. Ikeuchi, "Illumination distribution from shadows," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 1999, pp. 306–312.
- [21] I. Junejo and H. Foroosh, "Estimating Geo-temporal Location of Stationary Cameras Using Shadow Trajectories," in *Proc European Conf. Computer Vision*, 2008, pp. 318–331.
- [22] L. Wu, X. Cao, and H. Foroosh, "Camera calibration and geo-location estimation from two shadow trajectories," *Computer Vision and Image Understanding*, vol. 114, no. 8, pp. 915–927, 2010.
- [23] T. Wu and C. Tang, "A Bayesian Approach for Shadow Extraction from a Single Image," in *Proc. IEEE Int. Conf. on Computer Vision*, 2005, pp. 480–487.
- [24] T. Kim and K. Hong, "A practical single image based approach for estimating illumination distribution from shadows," in *Proc. IEEE Int. Conf. Computer Vision*, 2005, pp. 266–271.
- [25] I. Mariotto and V. Gutschick, "Non-Lambertian Corrected Albedo and Vegetation Index for Estimating Land Evapotranspiration in a Heterogeneous Semi-Arid Landscape," *Remote Sensing*, vol. 2, no. 4, pp. 926–938, 2010.
- [26] J. Lalonde, A. Efros, and S. Narasimhan, "Detecting ground shadows in outdoor consumer photographs," *Proc. European Conf. Computer Vision*, pp. 322–335, 2010.
- [27] D. Hoiem, A. Efros, and M. Hebert, "Recovering surface layout from an image," *Int. J. of Computer Vision*, vol. 75, no. 1, pp. 151–172, 2007.
- [28] J. Kosecka and W. Zhang, "Video compass," in *Proc. European Conf. Computer Vision*, 2002, pp. 476–490.
- [29] D. Liebowitz, A. Criminisi, and A. Zisserman, "Creating architectural models from images," in *Computer Graphics Forum*, vol. 18, no. 3, 1999, pp. 39–50.
- [30] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. 9, no. 1, pp. 62–66, 1979.
- [31] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, no. 6, pp. 679–698, 1986.
- [32] D. Savoie, *Sundials: Design, Construction, and Use*. Praxis, 2009, ISBN 0387098011.
- [33] R. Perez, R. Seals, and J. Michalsky, "All-weather model for sky luminance distribution—preliminary configuration and validation," *Solar energy*, vol. 50, no. 3, pp. 235–245, 1993.

media signal processing.

Pravin Kakar (S'10) obtained the B.Tech. degree in Electronics Engineering from the Veermata Jijabai Technological Institute (VJTI), Mumbai, India in 2009. He is currently pursuing the Ph.D. degree in Computer Engineering at the Institute for Media Innovation, Nanyang Technological University (NTU), Singapore. He is also working as a Project Officer at gameLAB in the School of Computer Engineering at NTU.

His research interests include multimedia forensics, computer vision, image processing and multi-

and neural networks.

N. Sudha (M'99-SM'07) received her B.E degree from Madurai Kamaraj University, India in 1992, M.S from Indian Institute of Technology (IIT) Madras in 1997 and Ph.D from IIT Guwahati in 2001. She has been working as an Assistant Professor in the School of Computer Engineering, Nanyang Technological University, Singapore since 2006. Previously she was employed at IIT Madras and IIT Guwahati.

Her research areas include computer vision and image processing, biometrics, embedded systems