Images Don't Forget: Online Photogrammetry to Find Lost Graves

Abby Stylianou, Joseph D. O'Sullivan, Austin Abrams and Robert Pless
Department of Computer Science
Washington University
Saint Louis, Missouri
Email: astylianou@wustl.edu, osullj@wustl.edu, abramsa@cse.wustl.edu, pless@wustl.edu

***Abstract*—The vast amount of public photographic data posted and shared on Facebook, Instragram and other forms of social media offers an unprecedented visual archive of the world. This archive captures events ranging from birthdays, trips, and graduations to lethal conflicts and human rights violations. Because this data is public, it has led to a new genre of journalism, one led by citizens finding, analyzing, and synthesizing data into stories that describe important events. To support this, we have built a set of browser-based tools for the calibration and validation of online images.**

This paper presents these tools in the context of their use in finding two separate lost burial locations. Often, these locations would have been marked with a headstone or tomb, but for the very poor, the forgotten, or the victims of extremist violence buried in unmarked graves, the geometric cues present in a photograph may contain the most reliable information about the burial location. The tools described in this paper allow individuals without any significant geometry background to utilize those cues to locate these lost graves, or any other outdoor image with sufficient correspondences to the physical world. We highlight the difficulties that arise due to geometric inconsistencies between corresponding points, especially when significant changes have occurred in the physical world since the photo was taken, and visualization features on our browser-based tools that help users address this.

I. Introduction

Images are a special type of data; they may carry within themselves strong cues about the identity of the people or objects in the scene, and the geolocation of those objects or the camera. These cues become easier to explore and more likely to be useful when there are large numbers of other images of the same objects or places. Fortunately, people, companies and governments share an immense set of imagery through, for example, Flickr and Facebook and Instragram and Google Street view. These provide the reference imagery to make it possible to understand events with a precision and geometric accuracy never before possible.

The constraints relating images to the calibration of the camera that took the picture and the real-world coordinates of points that is sees have long been studied within the fields of Computer Vision and Photogrammetry. But exploiting those constraints is typically limited to those with a working knowledge of linear algebra and special training in imaging geometry — even though the basic question "Where could the camera be to make the image look like this?" is quite accessible.

Fig. 1: Geocalibration.org, shown here, was used to determine a grave location for the family in this photograph. Users can calibrate any online image by selecting a set of points in the image (A), then the location of those points in the real world using a Google Maps interface (B). (C) The program solves for the camera that best projects those real world points onto the image points, and provides the user with a visualization of that camera's frustum (it's location and field of view) which updates as the user modifies the corresponding points. (D) Given that camera, the user can also view exactly what location each real world point projects to on the image, in the form of reprojection lines. A perfect ground truth camera's reprojection lines would lie exactly on top of the image points. (E) The user is also provided all of the camera's parameters in text format.

Therefore, we have worked to develop a set of web-based tools that allow anyone to geolocate imagery, and to solve for camera locations and focal lengths. We hope this democratizes the ability to formally reason about images across a large range of applications, from individuals who have a picture important to them and want to verify exactly where it was taken, to citizen journalists who seek to determine or validate the locations where pictures are taken in conflicts of global importance [1], [2]. This paper describes work that is publicly shared through a website: http://Geocalibration.org, describing trade-offs that we made in what type of camera model to use, how to design a user interface, and how to share results.

Our specific contributions are:

- an introduction to publicly available online photogrammetry tools that allow untrained individuals to calibrate images,
- a characterization of the sensitivity of geolocation results as a function of poor correspondences between an image and the physical world, and the features built into these online tools to circumvent this,
- an interface for viewing networks of calibrated images,
- two separate case studies using these tools to locate lost burial locations given varying degrees of confidence in point correspondences between a photo and the physical world.

II. Related Work

The problem of localizing a camera relative to a collection of points in a scene is not a new one in photogrammetry [11], [5], [4]. While significant progress has been made recently in the Computer Vision community to automatically recognize scenes and estimate their location, methods that determine accurate and precise camera locations, as opposed to more general localizations [12], [13], require either large networks of images where camera locations can be determined through feature matching between the non-geolocated images or images with known geolocation [9].

Other work in camera calibration for single images uses checkerboards with known 3D coordinates [11] or 3D geometry from Google Earth [3]. In cases such as those detailed in Section V, where the images needing to be calibrated are up to 30 years old, and where there is no Google Earth geometry for the scene, determining precise 3D locations for each correspondence can be difficult.

This paper builds on previous work in simplifying the camera calibration process into two dimensions [10].

III. Camera Calibration

Camera calibration is a method to optimize for the position of a camera given some information about the projective geometry in the scene. Given a set of real-world coordinates (X, Y, Z) and the corresponding locations (x, y) on the image, we consider the camera calibration to recover the 3D position, orientation, and focal length, or zoom level, of the camera most consistent with that set of correspondences. Here, we discuss camera calibration with known 3D world-coordinates, and then discuss a special case of camera calibration that only requires 2D world-coordinates (i.e. latitude and longitude).

A. Notation

We assume that we have many correspondences $(X_i, Y_i, Z_i) \rightarrow (x_i, y_i)$. The goal is to recover the real world 3D location of the camera (L_X, L_Y, L_Z) most consistent with those correspondences. In the process, we will also optimize over the 3D orientation—expressed as a (pan, tilt, roll) = (p, t, r) triple in radians—and the focal length, f, of the camera[1].

Calibration In 3D

Given the location, orientation, and focal length of the camera, a 3D point (X, Y, Z) projects into the image at (x, y) as determined by the following linear equation (see [7] for more details):

$$\begin{bmatrix} x\omega \\ y\omega \\ \omega \end{bmatrix} = \begin{bmatrix} f & 0 & \frac{w}{2} \\ 0 & f & \frac{h}{2} \\ 0 & 0 & 1 \end{bmatrix} \left[\begin{array}{c|c} R(p,t,r) & \begin{matrix} L_X \\ L_Y \\ L_Z \end{matrix} \end{array} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \tag{1}$$

where $R(p, t, r)$ is a 3×3 rotation matrix equivalent to panning, tilting, and rolling by the specified amounts, w and h are the width and height of the image (in pixels), and ω is a scaling factor. To recover the final (x, y) coordinate, we divide the left hand side by the third element:

$$\text{proj}\left(\begin{bmatrix} x\omega \\ y\omega \\ \omega \end{bmatrix}\right) = \begin{bmatrix} x \\ y \end{bmatrix} \tag{2}$$

To simplify notation, we rewrite the right hand side of Equation (1) as a 3×4 matrix M_{3D} that depends on the unknown parameters $L, p, t, r,$ and f. Camera calibration performs an optimization over these unknowns so that this *reprojection error* over all n correspondences is as small as possible:

$$\underset{L,p,t,r,f}{\text{argmin}} \sum_{i=1}^{n} \left\| \text{proj}\left(M_{3D}(L,p,t,r,f) \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix} \right) - \begin{bmatrix} x_i \\ y_i \end{bmatrix} \right\|^2 \tag{3}$$

Calibration in 2D

Determining the world-coordinates (X, Y, Z) for a given (x, y) point in an image can be challenging, whether due to scene changes like plant growth over time that actually make it impossible to measure the 3D location of an image point in the real world, or a lack of 3D geometry available from sources like Google Earth. If we only know the latitude and longitude of each correspondence, however, this provides the coordinates in a 2D space (i.e. the X and Y, but not Z). Assuming the camera has no roll, a similar derivation shows the projective equations still depend on the 2D location of the camera:

$$\begin{bmatrix} x\omega \\ \omega \end{bmatrix} = \begin{bmatrix} f & \frac{w}{2} \\ 0 & 1 \end{bmatrix} \left[\begin{array}{c|c} R(p) & \begin{matrix} L_X \\ L_Y \end{matrix} \end{array} \right] \begin{bmatrix} X \\ Y \\ 1 \end{bmatrix} \tag{4}$$

Notice that now, the projection only depends on the x coordinate in the image, and that only the pan angle p can be recovered. This is equivalent to calibrating a camera from a 1-dimensional image.

[1]The focal length, expressed in pixels, is a measure of the zoom level of the camera. A small focal length (say 50 pixels in a 640×480 image) corresponds to a very wide-angle field of view, while a larger focal length (say 2000 pixels) would be used in cameras more zoomed in, like a telephoto lens.

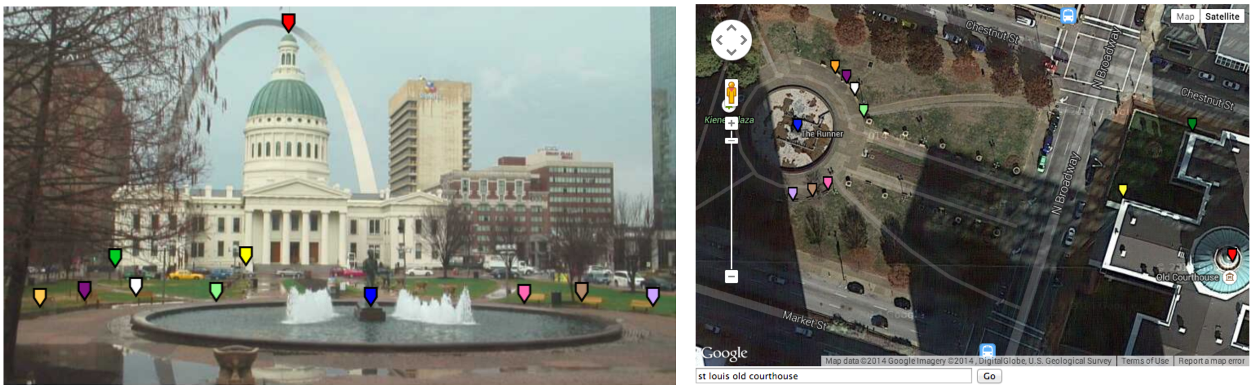

Fig. 2: (left) Correspondences in image. (right) Correspondences on map.

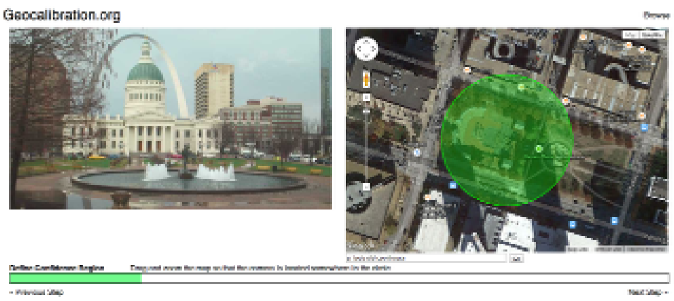

Fig. 3: Confidence region.

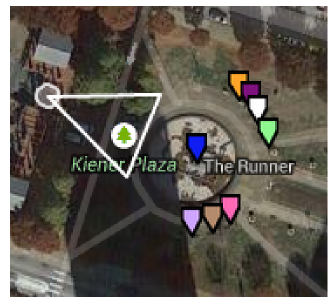

Fig. 4: Optimal camera frustum.

This simplified projection offers an optimization that returns L_X and L_Y, the geo-location of the camera on a ground plane:

$$\underset{L,p,f}{\text{argmin}} \sum_{i=1}^{n} \left|\left| \text{proj}\left(M_{2D}(L,p,f) \begin{bmatrix} X_i \\ Y_i \\ 1 \end{bmatrix} \right) - x_i \right|\right|^2 \quad (5)$$

where M_{2D} is the 2×3 matrix from Equation (4), and here $\text{proj}([\substack{x \\ y}]) = \frac{x}{y}$.

Therefore, given a set of image location (x, y) and their corresponding geolocations (X, Y), we can optimize the above objective to find the most consistent camera location L, camera pan angle p, and focal length f.

B. Optimization

Instead of directly using degrees latitude and longitude as (X, Y), we convert our coordinate space into a Cartesian coordinate frame with axes corresponding to meters East and North of some arbitrary origin point. Although this does not model the curvature of the Earth, we are only working with points in an approximately kilometer square region, so these errors are insignificant. After calibrating in this coordinate frame, we convert L back into a geographic coordinate frame.

Given a pan angle and focal length, the optimization becomes a system of linear equations over L. To speed up the optimization, we only optimize over p and f, and at each step, solve for the best L with linear least squares and report the reprojection error with respect to that L.

We evaluate the objective function from Equation (5) at a grid of (p, f) guesses, and optimize from the single parameter choice with the smallest error. We sample p at 20 degree intervals from 0 to 360 degrees, and the horizontal field of view, θ, at 15 degree intervals from 5 to 140 degrees. For a camera modeled as a pin hole, the θ is easily converted to f. This grid search gives us a good initial point, and from here, we use the Nelder-Mead simplex method [8] to find the best answer close to the initialization.

IV. GEOCALIBRATION.ORG

This tool is actively deployed at http://Geocalibration.org, a publicly accessible, browser-based tool where any individual can perform this calibration without any background in camera calibration or even geometry. A first generation version of this tool, called Project Live 2D, was developed in order to locate the lost grave of a crime victim [10] . We will briefly describe that work in Section V, but primarily in the context of the improvements made to Geocalibration.org as a result of the lessons learned in developing and using that initial tool.

The first step in using Geocalibration.org is to select an image that needs to be calibrated. The tool accepts any valid image URL, and then asks the user to define a region in which they believe the camera is location, as seen in Figure 3. This tool requires that a user have some general idea of the camera's location, as they have to be able to locate corresponding points between the image and the Google Maps aerial views.

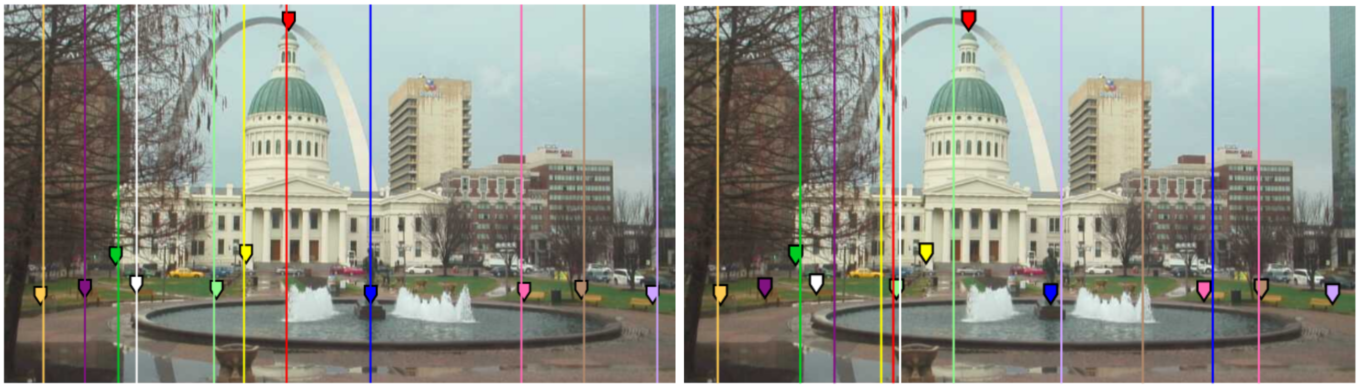

Fig. 5: (left) Good reprojection lines. (right) Bad reprojection lines.

Fig. 6: (left) (a) The global network of calibrated images as of publication. (right) Two cameras with overlapping fields of view.

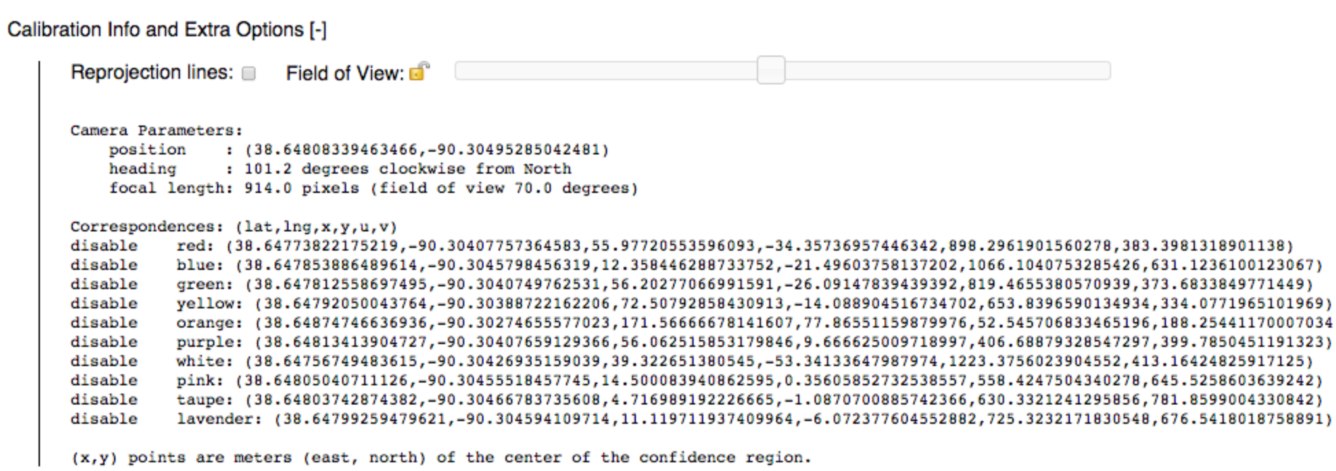

Fig. 7: Calibration output from Geocalibration.org.

Once the user has set this region, they can select up to 16 locations points in the image to then locate on Google Maps. This can be seen on the left side of Figure 2. Having a user place more than the minimum of four points minimizes the effect of errors from individual correspondences. The user then selects the corresponding real world locations, which can be seen on the right side of Figure 2. Once the user is satisfied with their correspondences and clicks **Next**, the program performs the calibration optimization detailed in Section III-A, the result of which is visualized on the map with the camera frustum in white, as seen in Figure 4. This shows the user the location and field of view for a camera that projects the points in the real world, (X, Y), onto the image with the smallest total distance between those projected points and the corresponding points in the image, (x, y), or the reprojection error. As the user adds more correspondences or adjusts existing ones, the optimization runs in real time, adjusting the frustum as the user moves the points. This provides a user with constant feedback and an intuition of the impact specific points have on the calibration.

If the overall reprojection error is high, then the calibration is not good, despite the optimization reaching a minimum. To help users visualize whether a calibration is good or not, the user has the ability to turn on "Reprojection Lines" as an overlay on the image. These reprojection lines show where the real world point projects into the image given the optimal camera parameters. We draw a line rather than a point, as we only know the x-coordinate that the point projects onto. This is a result of the geometric simplification of our calibration into 2d as described in Section III-A. Given a perfect, ground truth calibration, and points selected in the image at exactly the correct location, these reprojection lines would fall exactly on top of their matching colored points in the image. A poorly calibrated camera would have lines far away from their matching points, or lines on the wrong sides of nearby points. This reprojection error visualization is shown in Figure 5.

This visualization is turned on in a section on the page titled "Calibration Info and Extra Options." In this section, the user is provided the text results of the calibration, as well as the input information for each correspondence. In addition those items, there are two additional tools available to the user in this section. The first is a slider bar that allows the user to modify that camera's focal length by hand, and to lock it to a particular value. This is important if the user knows the focal length already, either from the image's exif data or from another non-geographic calibration technique (i.e., vanishing lines [6]). The second is an "enable/disable" button beside each correspondence's set of information. This allows the user to remove points from the calibration to see the impact that particular point had, without permanently deleting it. This can help to highlight points that are incorrectly positioned. These features can be seen in Figure 7.

Once a geometrically valid calibration has been found (i.e., no points reprojecting onto the wrong sides of each other, no extreme focal lengths, sufficiently low reprojection error), the correspondences and camera parameters are saved to a database. Every saved calibration is viewable on a global map on Geocalibration.org's **Browse** page, as shown in Figure 6(a). Not only does this page allow users to quickly return to previous calibrations, it also provides each camera the context of other, nearby calibrated cameras. Occasionally these cameras even have overlapping fields of view, as seen in Figure 6(b).

Fig. 8: An image from the 1983 burial of a Jane Doe, and a comparable view from 2013. Note the amount of dense growth in the background of the scene, as well as the large tree in the foreground that was not present in 1983. Significant changes over the last 30 years made selecting correspondences between the original images and the current scene difficult.

V. CASE STUDIES

Case Study 1: Lost Jane Doe Grave Location

In 1983, a young girl was found decapitated in the basement of an apartment building in St. Louis, Missouri. After the police failed to make progress in finding either her head or her killer, they buried her in a pauper's burial at Washington Park Cemetery. The funeral was photographed by St. Louis Post-Dispatch photographer Ed Sedej. In the years following the burial, the cemetery fell into disrepair and become completely overgrown. One of the original photographs from 1983, as well as a photograph from approximately the same location in 2013 are shown in Figure 8.

In 2011, the St. Louis Police Department (STLPD) attempted to exhume the girl's remains in order to perform modern forensic analysis to determine her identity or location of origin. Multiple efforts to exhume her remains failed, prompting a ban on further exhumation attempts until further evidence of her burial location was provided. We used a first generation version of the tool described in this paper to successfully identify the location of her remains from a single photograph of the funeral. This primary challenge in locating that burial site was working with the significant physical change that had taken place at the cemetery over thirty years

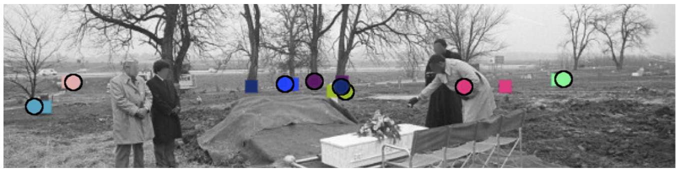

Fig. 9: A prior attempt at visualizing reprojection error in a calibration.

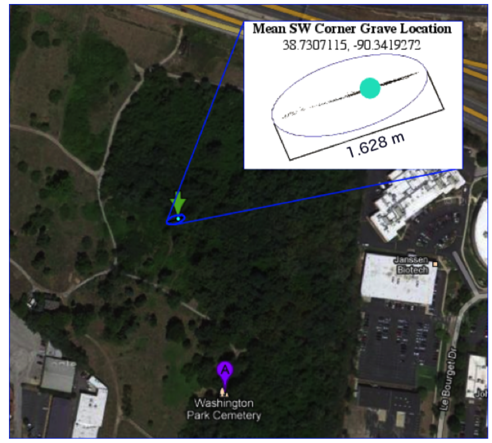

Fig. 10: The correct location of the lost Jane Doe grave is shown with present day aerial imagery of Washington Park Cemetery. Compared to the sparse trees in the original photo, the cemetery is now heavily forested. Imagery ©2013 DigitalGlobe, U.S. Geological Survey, USDA Farm Service Agency, Map data ©2013 Google.

and the difficulties that created in selecting valid corresponding points. This work is described at length in [10].

Through that case, we learned several noteworthy lessons that have informed to development of Geocalibration.org and led to the development of several of the features detailed in Section IV. Detailed below are the specific reasons for the addition of these features.

Starting Location (Figure 3) Our first attempts at geolocating this lost grave used cell phones to collect GPS coordinates at the real world locations for our correspondences. Attempting to use these inaccurate locations resulted in a camera in completely the wrong location. While we did not know the precise camera calibration, we knew the general vacinity that the camera would have had to be located and any time spent on calibrations outside of that region were wasted time. While we knew this starting location down to about a 10 meter across region, Geocalibration.org allows a starting region up to about a mile across. In cases where that camera location isn't known even at that level, Geocalibration.org is not the correct tool for performing camera calibration, as the user will likely be unable to find corresponding points between the image and the real world.

Additional Points While it is only geometrically necessary to select four correspondences in order to solve for the camera's location, focal length and heading, it is important to have as many points as possible across as much of the width of image as possible, as well as points at significantly different depths. Additional points help to account for error in measuring a correspondence's real world location or selecting its correct pixel location in the image.

The tool that we used to locate this lost grave initially used only a few corresponding points, at least one of which was we had measured to be in completely the wrong location. This had a larger impact on our calibration than when we updated the tool to use more points. Geocalibration.org now allows up to sixteen corresponding points. More points than that have diminishing returns on the improvement to the calibration and become infeasible to manage in the user interface.

Enable and Disable Correspondences (Figure 7) In adding additional points, we occasionally tried adding points whose locations we were less confident about, in order to observe the impact that correspondence had on the calibration (i.e., a tree in 2013 that looked similar to a tree in the original photograph but that might not actually be the correct tree). The original interface had no simple way to include or exclude points from the calibration, without deleting the points from the set of correspondences. To minimize frustration, especially in scenes that have repeated items, like power lines or telephone poles that may be difficult to uniquely identify correctly in aerial imagery, allowing a user to quickly label which points to include in the calibration makes it much easier to find a consistent set of matches.

Reprojection Lines (Figure 5) The initial tool used in locating this lost grave did have a first attempt at visualizing the reprojection error. This is shown in Figure 9, where the projected location is shown in the image as a square whose color matches its corresponding image location. This, however, gave the impression that we knew the y-coordinate of the reprojected location, when in fact we only know the x-coordinate, as explained in Section III-A. The small icons also occasionally occluded each other, making it difficult to analyze the reprojection error for different correspondences. To account for both of these problems, we converted the reprojection error visualization to lines drawn at the reprojected x-coordinate in the color matching the image point.

Real Time Updates The impact a particular correspondence has on the overall calibration is not intuitive, especially to individuals who are not exceedingly familiar with the geometry of camera calibration. For example, moving a correspondence that was close to the camera has a larger impact on the final calibration than moving points far away. While that may be obvious to individuals who are well versed in camera calibration, Geocalibration.org has been designed for anyone seeking to calibrate an image regardless of their geometric background—as part of that effort, we implemented real time updates to both the reprojection error and camera frustum visualizations. As a user moves a correspondence on either the image or the map, the impact that change has on the final calibration is immediately shown in both visualizations.

Saving & Browsing (Figure 6) The last improvement made on Geocalibration.org as a result of lessons learned in locating the lost grave of the young Jane Doe was the ability to save good calibrations and return to them at a later date, rather than redoing the entire process every time a change is made or the calibration needs to be retrieved. Once a calibration on Geocalibration.org is geometrically valid (i.e., no points reprojecting onto the wrong sides of each other, no extreme focal lengths, sufficiently low reprojection error), those calibration results are saved to a database and available to the user on the Browse page of Geocalibration.org, where every calibrated image is linked from an icon at the calibrated location on a world map. The user can make subsequent changes to the calibration at any time.

Case Study 2: A Second Lost Grave

Following our successful identification of the lost Jane Doe grave, our lab was approached by a family that had lost their mother's grave in Washington Park Cemetery. They had photographs of the funeral, seen in Figure 11(a), but after many attempts, could not find their mother's headstone where they remembered it being. The photographs show the family and the casket, as well as the building and telephone poles that are still present at the site today, and can be seen on Google Maps aerial imagery. While many of the same difficulties from the Jane Doe case existed in here regarding the passage of time and difficulties of working in what is essentially dense forest, the building and telephone poles are more reliable correspondences than the trees and headstones in the Jane Doe case. Another difference is that we were seeking an existing headstone rather than an unmarked grave. Because of this, it was sufficient to find a reasonably good camera geocalibration and then look for the headstone in the vacinity of that location.

In the photographs provided by the family, there are evenly spaced vertical concrete beams on the exterior of the building in the background. By piecing together two photographs that view both ends of the side of the building closest to the funeral as well as a light in both images which served as an anchor point between the two images, we were able to determine that there were seven evenly spaced beams. Using Google Maps we were able to measure the length of that side of the building and divide that by seven to get the distance between concrete beams. This process is shown in Figure 11(b). In Figure 11(a), three of these beams are visible, as well as the light that falls between two beams and the two telephone poles that are visible on Google Maps. These are the points that we used in Geocalibration.org, which resulted in the camera frustum shown in Figure 11(c). We returned to the cemetery with the family and used our cell phones to get as close as possible to the location output by Geocalibration.org. This location was within several feet of the lost headstone, shown in Figure 11(d).

While locating this grave was technically simpler than the Jane Doe case, it demonstrated the usefulness of this tool for individuals outside of the Computer Vision community. Locating this family's lost grave site was done using only Google Maps and Geocalibration.org, both tools that are available to and usable by the general public without any significant training.

VI. Discussion & Conclusion

Our goal is to make it possible for anyone to exploit the strong geometric constraints available within photographs. These constraints are inherent in the images and may be useful to understand where, exactly, someone was standing or where

(a) A photo from the burial of a woman at Washington Park Cemetery.

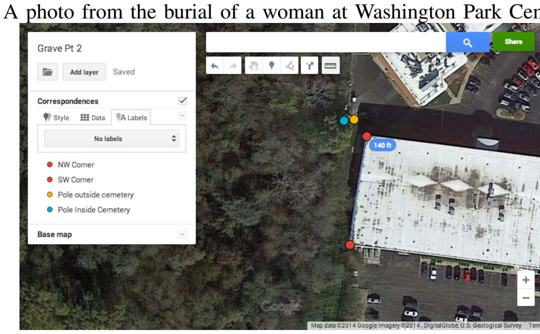

(b) Finding points and making measurements using Google Maps.

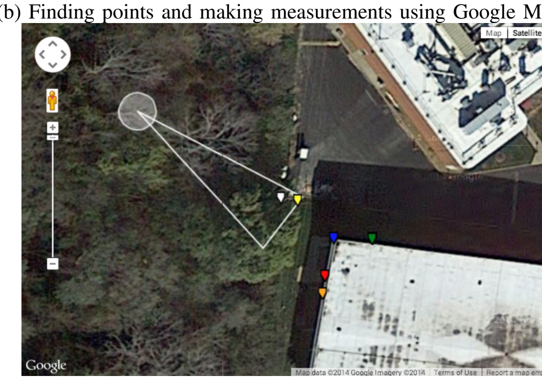

(c) Geocalibration.org output showing the camera location.

(d) The headstone of the woman buried in (a), located using only Geocalibration.org and Google Maps.

Fig. 11: Following our succesful identification of the lost grave of a young Jane Doe in Washington Park Cemetery, a family contacted us to help locate their mother's lost headstone.

an object was in a scene. These questions arise in many contexts, and by sharing these tools broadly, we hope to give everyone access to tools they can use to answer the questions that are most relevant to them.

Thus, our work focused on creating an intuitive interface, based on lessons learned in building this tool for the particular forensic application of geo-locating a grave from 30 year old photos. Anecdotally, we find that people have a very good intuition for matching points from 2D images to google maps overhead view, once they have approximately figured out the right region in the map. We find that people have good intuitions about points in the scene changes, such as trees that might have grown. We find the people, even experienced computer vision researchers (us), have very bad intuition at understanding how important a single corresponding point is in defining the optimization solution. This led us to explicitly build an interactive interface that makes it possible to ”jitter” one of the point correspondences and see its effect on the solution. Finally, in real scenes, especially with a 30 year gap, there are often points that are ”maybe correct”, such as a billboard which may or may not be in the same location, so it is important to give tools that let a user explore which sets of points to include.

ACKNOWLEDGMENT

This material is based upon work partially supported by the National Science Foundation under Grant No. NSF-IIS1111398, and NSF-DEB1053554, and by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory, contract FA8650-12-C-7212. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL, or the U.S. Government.

REFERENCES

[1] bellingcat, by and for citizen investigative journalists. *https://www.bellingcat.com/*.

[2] Ukraine@war. *http://ukraineatwar.blogspot.com/*.

[3] A. Abrams and R. Pless. Webcams in context: Web interfaces to create live 3D environments. In *ACMMM*, pages 331–340, Jun 2010.

[4] A. Y. I. Aziz and H. M. Karara. Direct linear transformation into object space coordinates in close-range photogrammetry. In *Proc. of the Symposium on Close-Range Photogrammetry*, pages 1–18, Urbana, Illinois, 1971.

[5] D. C. Brown. Close-range camera calibration. *Photogrammetric Engineering*, 37(8):855–866, 1971.

[6] E. Guillou, D. Meneveaux, E. Maisel, and K. Bouatouch. Using vanishing points for camera calibration and coarse 3d reconstruction from a single image. *The Visual Computer*, 16(7):396–410, 2000.

[7] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*, volume 2. Cambridge Univ Press, 2000.

[8] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.

[9] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: Exploring photo collections in 3D. In *SIGGRAPH Conference Proceedings*, pages 835–846, 2006.

[10] A. Stylianou, A. Abrams, and R. Pless. Finding jane doe: a forensic application of 2d image calibration. *IET Conference Proceedings*, pages 1.17–1.17(1), January 2013.

[11] R. Y. Tsai. An efficient and accurate camera calibration technique for 3D machine vision. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 364–374, Miami, June 1986.

[12] L. Xie and S. Newsam. IM2MAP: deriving maps from georeferenced community contributed photo collections. In *Proceedings of the 3rd ACM SIGMM international workshop on Social media*, pages 29–34, 2011.

[13] W. Zhang and J. Kosecka. Image based localization in urban environments. In *International Symposium on 3D Data Processing, Visualization and Transmission*, 2006.