

Accurate Image Localization Based on Google Maps Street View

Amir Roshan Zamir and Mubarak Shah*

University of Central Florida, Orlando FL 32816, USA

Abstract. Finding an image's exact GPS location is a challenging computer vision problem that has many real-world applications. In this paper, we address the problem of finding the GPS location of images with an accuracy which is comparable to hand-held GPS devices. We leverage a structured data set of about 100,000 images build from Google Maps Street View as the reference images. We propose a localization method in which the SIFT descriptors of the detected SIFT interest points in the reference images are indexed using a tree. In order to localize a query image, the tree is queried using the detected SIFT descriptors in the query image. A novel GPS-tag-based pruning method removes the less reliable descriptors. Then, a smoothing step with an associated voting scheme is utilized; this allows each query descriptor to vote for the location its nearest neighbor belongs to, in order to accurately localize the query image. A parameter called *Confidence of Localization* which is based on the Kurtosis of the distribution of votes is defined to determine how reliable the localization of a particular image is. In addition, we propose a novel approach to localize groups of images accurately in a hierarchical manner. First, each image is localized individually; then, the rest of the images in the group are matched against images in the neighboring area of the found first match. The final location is determined based on the *Confidence of Localization* parameter. The proposed image group localization method can deal with very unclear queries which are not capable of being geolocated individually.

1 Introduction

Determining the exact GPS location of an image is a task of particular interest. As there are billions of images saved in online photo collections - like Flickr, Panoramio etc. - there is an extant resource of information for further applications [1,2]. For example, in Agarwal et al. [1], a structure from motion approach is employed to find the 3D reconstruction of Rome using GPS-tagged images of the city. Many such applications need some sort of information about the exact location of the images; however, most of the images saved on the online repositories are not GPS-tagged. A system that is capable of finding an exact

* The authors would like to thank Jonathan Pooock for his valuable technical contributions and comments on various drafts of the submission, which have significantly improved the quality of the paper.

location using merely visual data can be used to find the GPS-tag of the images and thus make the huge number of non-GPS-tagged images usable for further applications.

However, there are many images which are incapable of being localized individually, due to their low quality, small size or noise. Many of these images are saved in albums or image groups; these groupings can act as clues to finding the exact location of the unclear image. For instance, images saved in online photo collections in an album usually have locations that are close to one another.

Visual localization of images is an important task in computer vision. Jacobs et al. [3] use a simple method to localize webcams by using information from satellite weather maps. Schindler et al. [4] use a data set of 30,000 images for geolocating images using a vocabulary tree [5]. The authors of [6] localize landmarks based on image data, metadata and other sources of information. Kalogerakis et al. [7] leverage images in a sequence to localize them in a global way. In their method, they use some travel priors to develop the chronological order of the images in order to find the location of images. Zhang et al. [8] perform the localization task by matching image key points and then applying a geometrical alignment. Hakeem et al. [9] find the geolocation and trajectory of a moving camera by using a dataset of 300 reference images. Although much research has been done in the area of localizing images visually, many other sources of information can be used alongside the visual data to improve the accuracy and feasibility of geolocation, such as used in Kalogerakis et al. [7]. To the best of our knowledge, image localization utilizing groups of images has not been investigated; as such, this paper claims to be the first to use the proximity information of images to aid in localization.

In our method, a query image is matched against a GPS-tagged image data set; the location tag of the matched image is used to find the accurate GPS location of the query image. In order to accomplish this, we use a comprehensive and structured dataset of GPS-tagged Google Maps Street View images as our reference database. We extract SIFT descriptors from these images; in order to expedite the subsequent matching process, we index the data using trees. The trees are then searched by a nearest-neighbor method, with the results preemptively reduced by a pruning function. The results of the search are then fed through a voting scheme in order to determine the best result among the matched images. Our proposed *Confidence of Localization* parameter determines the reliability of the match using the Kurtosis of the voting distribution function. Also, we propose a method for localizing group of images, in which each image in the query group is first localized as a single image. After that, the other images in the group are localized within the neighboring area of the detected location from the first step. A parameter called CoL_{group} is then used to select the rough area and associated corresponding accurate locations of each image in the query group. The proposed group localization method can determine the correct GPS location of images that would be impossible to geolocate manually. In the results section, we show how our proposed single and group image localization methods are significantly more accurate than the current methods.



Fig. 1. We use a dataset of about 100,000 GPS-tagged images downloaded from Google Maps Street View for Pittsburg, PA (Right) and Orlando, FL (left). The green and red markers are the locations of reference and query images respectively.

2 Google Maps Street View Dataset

Different type of image databases have been used for localization tasks. In Ha-keem et al. [9] a database of 300 GPS-tagged images is used, whereas Kalogerakis et al. [7] leverage a dataset of 6 million non-structured GPS-tagged images downloaded from internet, and Schindler et al. [4] use a data set of 30,000 street-side images. We propose using a comprehensive 360° structured image dataset in order to increase the accuracy of the localization task. The images extracted from Google Maps Street View are a very good example of such a dataset. Google Maps Street View is a very comprehensive dataset which consists of 360° panoramic views of almost all main streets and roads in a number of countries, with a distance of about 12m between locations. Using a dataset with these characteristics allows us to make the localization task very reliable, with respect to feasibility and accuracy; this is primarily due to the comprehensiveness and organization of the dataset. The following are some of the main advantages of using datasets such as Google Maps Street View:

- **Query Independency:** Since the images in the dataset are uniformly distributed over different locations, regardless of the popularity of a given location or object, the localization task is independent of the popularity of the objects in the query image and the location.

- **Accuracy:** As the images in the data set are spherical 360° views taken about every 12 meters, it is possible to correctly localize an image with a greater degree of accuracy than would be permitted by a sparser data set comprised of non-spherical images. The achieved accuracy is comparable to - and, in some cases, better than - the accuracy of hand-held GPS devices.

- **Epipolar Geometry:** The comprehensiveness and uniformity of the data set makes accurate localization possible without employing methods based on epipolar geometry [9]- methods which are usually computationally expensive and, in many cases, lacking in required robustness. Additionally, the camera's intrinsic parameters for both the query and the dataset images are not required in order to accurately localize the images.

- **Secondary Applications:** Using a structured database allows us to derive additional information, without the need for additional in-depth computation. For example, camera orientation can be determined as an immediate result of localization using the Google Maps Street View data set, without employing methods based on epipolar geometry. Since the data set consists of 360° views, the orientation of the camera can be easily determined just by finding which part of the 360° view has been matched to the query image - a task that can be completed without the need for any further processing. Localization and orientation determination are tasks that even hand-held GPS devices are not capable of achieving without motion information.

However, the use of the Google Maps Street View dataset introduces some complications as well. The massive number of images can be a problem for fast localization. The need for capturing a large number of images makes using wide lenses and image manipulation (which always add some noise and geometric distortions to the images) unavoidable. Storage limitations make saving very high quality images impossible as well, so a matching technique must be capable of dealing with a distorted, low-quality, large-scale image data set. The database's uniform distribution over different locations can have some negative effects - while it does make the localization task query-independent, it also limits the number of image matches for each query as well. For example, a landmark will appear in exactly as many images as a mundane building. This is in direct contrast to other current large scale localization methods like Kalogerakis et al. [7], which can have a large number of image matches for a location in their database - a fact especially true if a location is a landmark; this allows the localization task to still be successful on a single match. The small number of correct matches in our database makes the matching process critical, as if none of the correct matches - which are few in number - are detected, the localization process fails.

We use a dataset of approximately 100,000 GPS-tagged Google Street View images, captured automatically from Google Maps Street View web site from Pittsburgh, PA and Orlando, FL. The distribution of our dataset and query



Fig. 2. Sample Reference Images. Each row shows one placemark's side views, top view and map location.

images are shown in Fig. 1. The images in this dataset are captured approximately every 12 meters. The database consists of five images per placemark: four side-view images and one image covering the upper hemisphere view. These five images cover the whole 360° panorama. By contrast, Schindler et al.’s [4] dataset has only one side view. The images in their dataset are taken about every 0.7 meters, covering 20km of street-side images, while our dataset covers about 200km of full 360° views. Some sample dataset images are illustrated in Fig. 2.

3 Single Image Localization

Many different approaches for finding the best match for an image has been examined in the literature. Hakeem et al. [9] perform the search process by nearest-neighbor search among SIFT descriptors of a small dataset of about 300 reference images. Kalogerakis et al. [7] perform the task by calculating a number of low-level features - such as color histograms and texton histograms - for 6 million images while assuming that there is a very close match for the query image in their dataset. Schindler et al. [4] try to solve the problem by using the bag of visual words approach. In the results section, we show that the approach in Schindler et al. [4] cannot effectively handle large-scale datasets that are primarily comprised of repetitive urban features. In order to accurately localize images, we use a method based on a nearest-neighbor tree search, with pruning and smoothing steps added to improve accuracy and eliminate storage and computational complexity issues.

During training, we process the reference dataset by computing the SIFT descriptors [10] for all interest points detected by the SIFT detector [10,11]. Then, the descriptor vectors (and their corresponding GPS tags) are organized into a tree using FLANN [12]. As we show later, a well-tuned pruning method allows us to find very reliable descriptors; as such, we generally need to compute at most $\frac{1}{6}$ of the number of interest points that Schindler et al. [4]’s method

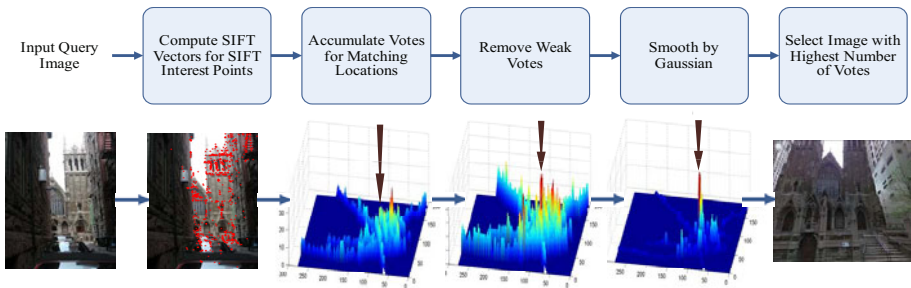


Fig. 3. Block diagram of localization of a query image. Lower row shows the corresponding results of each step for the image. Note the streets in the vote plots, as the votes are shown over the actual map. The dark arrow points toward the ground truth location. The distance between the ground truth and matched location is 17.8m.

requires. Fig. 3 shows the block diagram of the proposed method for localizing a query image. In the first step, the SIFT descriptors are computed for SIFT interest points in the same way as we process the dataset during training. Then, in the second step, the nearest-neighbors for each of the query SIFT vectors are found in the tree. Each of the retrieved nearest-neighbors vote for the image that they belong to. The votes can be shown as a plot over the actual map of the area covered by our reference dataset (as shown in third column of Fig. 3).

As noisy interest points are commonly detected in an image, a pruning step is essential. Lowe et al. [10] find reliable matches by setting a maximum threshold of 0.8 on the ratio of the distance between the query descriptor and the first and second nearest neighbors. For geolocation tasks in large-scale datasets, the pruning step becomes more important; this is primarily because many of the processed descriptors belong to non-permanent and uninformative objects (ie. vehicles, people, etc), or are detected on the ground plane - both cases where the descriptors become misleading for geolocation purposes. The massive number of descriptors in the dataset can add noisy, unauthenticated matches as well. Schindler et al. [4] find the more informative visual words by maximizing an information gain function, a process which requires reference images with significant overlap. Hakeem et al. [9] prune their dataset by setting the maximum SIFT threshold proposed in Lowe et al. [10] to 0.6 in order to keep more reliable matches. We propose using the following function in order to prune the matches:

$$V_{flag}(d_i) = \begin{cases} 1 & \frac{NN(d_i,1)}{NN(d_i,Min\{j\})} < 0.8 \\ 0 & otherwise \end{cases} \quad \forall j \rightarrow |Loc(NN(d_i,1)) - Loc(NN(d_i,j))| > D , \quad (1)$$

where $V_{flag}(d_i)$ is the flag of the vote corresponding to the query descriptor d_i . If the flag is 0, the descriptor is removed in the pruning step; if the flag is 1, it participates in the voting. $NN(d_i, k)$ is the k_{th} nearest-neighbor of d_i . $Loc(NN(d_i, k))$ is the GPS location of the k_{th} nearest-neighbor to descriptor d_i and $||$ represents the actual distance between the two GPS locations of the nearest neighbor. At its core, Eq. 1 may appear to be the SIFT ratio [10]; the changes we have made mean that the descriptor in the denominator is dynamically determined, based on actual GPS distance. This is an important difference, as allowing this ratio to be determined dynamically creates a great advantage over the simple ratio between first and second nearest-neighbors used in Lowe et al. [10] and Hakeem et al. [9], in that it allows the localization task to handle repeated urban structures more accurately. The importance of this method becomes clearer by considering the reference images shown in Fig. 2. The windows of the skyscraper shown in the 3_{rd} column, 3_{rd} row of the figure are identical, leading to very close nearest-neighbor results for a query descriptor of this window (as shown in bottom left corner image in Fig. 4). While the SIFT ratio used in Lowe et al. [10] and Hakeem et al. [9] removes this descriptor in the pruning step, the proposed method retains it, as the location of all of the very similar nearest neighbors are close to each other. In other words, even though we cannot necessarily determine which of the windows shown in the query image

correspond to each of the windows in the skyscraper, they will still be voting for the correct location, as the GPS-tag of all these very similar nearest-neighbors point to one location. To explain it in a less-anecdotal way, Eq. 1 removes a descriptor only if the descriptor in the denominator does not belong to any of the nearby locations of the first nearest-neighbor AND the ratio is greater than 0.8. As can be seen in the 4th column of Fig. 3, the votes around the ground truth location are mostly retained, whereas many of the incorrect votes are removed.

Since some of the objects in a query image may be in view of several reference images, we smooth the votes based on their locations in order to prevent the votes from being scattered using this equation:

$$V_{smoothed}(\lambda', \phi') = \sum_{\lambda=-\infty}^{+\infty} \sum_{\phi=-\infty}^{+\infty} e^{-\frac{(\lambda^2 + \phi^2)}{2\sigma'^2}} V(\lambda' - \lambda, \phi' - \phi) V_{flag}(\lambda' - \lambda, \phi' - \phi) , \tag{2}$$

where $V(\lambda, \phi)$ and $V_{flag}(\lambda, \phi)$ are the voting and flags function (respectively), for the GPS location specified by λ and ϕ , and the first coefficient is the 2D Gaussian function with a standard deviation of σ' . As each descriptor is associated with a GPS-tagged image, we can represent the voting function's parameter in terms of λ and ϕ . As can be seen in column 5 of Fig. 3, the smoothing step makes the peak which corresponds to the correct location more distinct.

As shown in the block diagram in Fig. 3, the location which corresponds to the highest peak is selected as the GPS location of the query image.

3.1 Confidence of Localization

There are several cases in which a query image may - quite simply - be impossible to localize. For instance, a query might come from an area outside of the region covered by the database; alternatively, the image might be so unclear or noisy that no meaningful geolocation information can be extracted from it. A parameter that can check for (and, consequently, prevent) these kind of positive errors is important. In probability theory, statistical moments have significant applications. The Kurtosis is a measure of whether a distribution is tall and slim or short and squat [13]. As we are interested in examining the behavior of the voting function in order to have a measure of reliability, we normalize it and consider it as a probability distribution function. Since the Kurtosis of a distribution can represent the peakedness of a distribution, we propose to use it as a measure of *Confidence of Localization*, since a tall and thin vote distribution with a distinct peak corresponds to a reliable decision for the location; correspondingly, a widely-spread one with a short peak represents a poor and unreliable localization. Our *Confidence of Localization* parameter is thus represented by the following equation:

$$CoL = Kurt(V_{smoothed}) = -3 + \frac{1}{\sigma^4} \sum_{\phi=-\infty}^{+\infty} \sum_{\lambda=-\infty}^{+\infty} [(\lambda - \mu_\lambda)^2 (\phi - \mu_\phi)^2] V_{smoothed}(\lambda, \phi), \tag{3}$$

where $V_{smoothed}$ is the vote distribution function (see Eq. 2). The above equation is the Kurtosis of the 2D vote distribution function, with random variables λ and ϕ , corresponding to the GPS coordinates. A high Kurtosis value represents a distribution with a clearer and more defined peak; in turn, this represents a higher confidence value. In the next section, we use this *CoL* parameter to localize a group of images.

4 Image Group Localization

We propose a novel hierarchical approach to localize image groups. The only assumption inherent in the proposed method is that all of the images in the group must have been taken within the radial distance R of each other; this radial distance R is a parameter that can be set in the method. In our approach, no information about the chronological history of the images is required.

To localize an image group consisting of images I_1 to I_N , we employ a hierarchical approach consisting of two steps:

- Step 1, Individual Localization of Each Image: In the first step of the approach, all of the images in the group are localized individually, independent from other images. In order to do this, we use the Single Image Localization method described previously in section 3; thus, each one of the single images in the group returns a GPS location.

- Step 2, Search in Limited Subsets: In the second step, N subsets of reference images which are within the distance R of each of the N GPS locations found in step 1 are constructed. Following that, a localization method - similar to the method defined in section 3 - is employed for localizing the images in the group; however, in this case, the dataset searched is limited to each of the N subsets created by the initial search. We define the *CoL* value for each of the secondary, sequential search processes done in each of the limited subsets as:

$$CoL_{group}(S) = \sum_{i=1}^N \frac{CoL_i}{N} , \quad (4)$$

where S represents each of the secondary search processes. Once the CoL_{group} value for each of the limited subsets is calculated, the subset that scores the highest value is selected as the rough area of the image group. From there, each query image is assigned the GPS location of the match that was found in that limited subset.

Since this proposed approach to image group localization requires multiple searches in each step, the computational complexity of the method is of particular interest. The number of necessary calculations for localizing a single query image in our method is dependent on the number of detected interest points in the image. If we assume C is a typical number representing the number of required calculations for localizing an image individually, the number of required calculations to localize a group of images using the proposed approach is:

$$CE_{group}(N, \delta) = C(N + \frac{(N-1)N}{\delta}) , \quad (5)$$

where N is the number of images in the group and δ is a constant that is determined by the size of the limited subsets used in the step 2 of section 4. δ ranges from 1 to ∞ , where 1 means each limited subset is as large as the whole dataset and ∞ means each subset is extremely small. Since the number of required calculations to localize an image individually is C , the number of required calculations to localize N images individually will be $N \times C$, so the percentage increase in computational complexity using the proposed group method vs. the individual localization method is :

$$PE(N, \delta) = \frac{CE_{group}(N, \delta) - N \times C}{N \times C} \times 100 , \quad (6)$$

i.e.,

$$PE(N, \delta) = \frac{N - 1}{\delta} \times 100 , \quad (7)$$

For 4 and 50 - both typical values for N and δ , respectively - the increase in computational complexity is 24%, garnering a roughly three-fold increase in system accuracy.

5 Experiments

Our test set consists of 521 query images. These images are all GPS-tagged, user-uploaded images downloaded from online photo-sharing web sites (Flickr, Panoramio, Picasa, etc.) for Pittsburgh, PA and Orlando, FL. Only indoor images, privacy-infringing images and irrelevant images (e.g. an image which only shows a bird in the sky), are manually removed from the test set. In order to ensure reliability of results, all the GPS tags of the query images are manually checked and refined, as the user-tagged GPS locations are usually very noisy and inaccurate. Fig. 4 depicts some of the images.

311 images out of the 521 query images are used as the test set for the single-image localization method; 210 images are organized in 60 groups of 2,3,4 and 5 images with 15 groups for each as the test set for group image localization method.

5.1 Single Image Localization Results

Fig. 5 shows the results of the localization task for the test set of 311 images. In order to avoid computational issues of indexing the large number of images in a single tree, we construct 5 individual trees spanning the whole dataset. The final nearest-neighbor selected is chosen from among the 5 nearest-neighbor results retrieved across each tree. In these experiments, the queries and reference images of both of the cities are used. In order to make the curves in Fig. 5 invariant with respect to differing test sets, we randomly divide the single image localization method's test set into ten smaller test sets; likewise, we divide the group image localization method's test set into 5 smaller test sets. The curves in Fig. 5 are the average of the result curves generated for each of the smaller test sets. As




Fig. 4. Sample query images in our test set

can be seen in Fig. 5, all of the steps proposed in Fig. 3 improve the accuracy significantly. The smoothing step unifies the votes, leading to a more distinct correct peak, while attenuating the incorrect votes. Dynamic pruning removes the wrong matches, bringing about a more accurate localization task; this enables us to calculate and save fewer SIFT descriptors per image. By comparison, we have found (on average) 500 SIFT interest points per image; in Schindler et al. [4], the implementation used about 3000 interest points. As can be seen in Fig. 5, our method shows a significant improvement over the bag of visual words method used by Schindler et al. [4]. This is mostly due to the fact that, in the very similar and repeated structures of an urban area, the information lost in the quantization becomes critical. Additionally, the method proposed in Schindler et al. [4] requires reference images with significant overlap to maximize the information gain function, an assumption which can lead to significant issues in large scale localization. As can be seen in Fig. 5, about 60% of the test set is localized to within less than 100 meters of the ground truth; by comparison, this number for the method by Schindler et al. [4] is about 22%. However, our method fails when images are extremely cluttered with non-permanent objects (e.g. cars, people) or objects of low informative values (e.g. foliage).

In order to examine the performance of the proposed *CoL* function, the distribution of the *CoL* values of the localization of the test set consisting of 311 images is shown in Fig. 6 versus the error distance. The 311 *CoL* values are grouped into 8 bins based on the *CoL* values; the mean error of each of the bin members are shown on the vertical axis. As observed in the figure, higher *CoL* values - due to distinct peaks in the vote distribution function - correspond to lower error, meaning the localization is more reliable. Since theoretically the value of the Kurtosis is not limited, we normalize the *CoL* values and show them ranging from 0 to 1 on the plot.




Fig. 5. The left figure shows the single image localization method results vs. Schindler et al.’s method, along with the curves representing the effect of each step. The right figure shows the localization results using the proposed image group localization method.




Fig. 6. The left figure shows the distribution of *CoL* values for the localization of the test set of 311 images. The *CoL* values are organized in 8 bins; the vertical axis shows the mean error value in meters for each bin. The right figure shows the breakdown of the results from the test set of the group image localization method based on the number of images in each group.

In order to show the importance of a parameter which represents the reliability of the localization task, we performed another experiment on *CoL* by using a test set of 62 query images. 34 of the images are from the Pittsburgh query set; 28 are from the Orlando query set. In this experiment, we grow one tree for each city, allowing the *CoL* function to determine the correct tree to use for each query. We localize each query image using each tree. Since a low *CoL* value for the tree to which the query image does not belong is expected, we select the location returned by the tree with higher *CoL* value as the final location of the query images. By this method, the proposed *CoL* parameter selected the correct location for 53 images out of the 64 test images - an accuracy of 82%. This

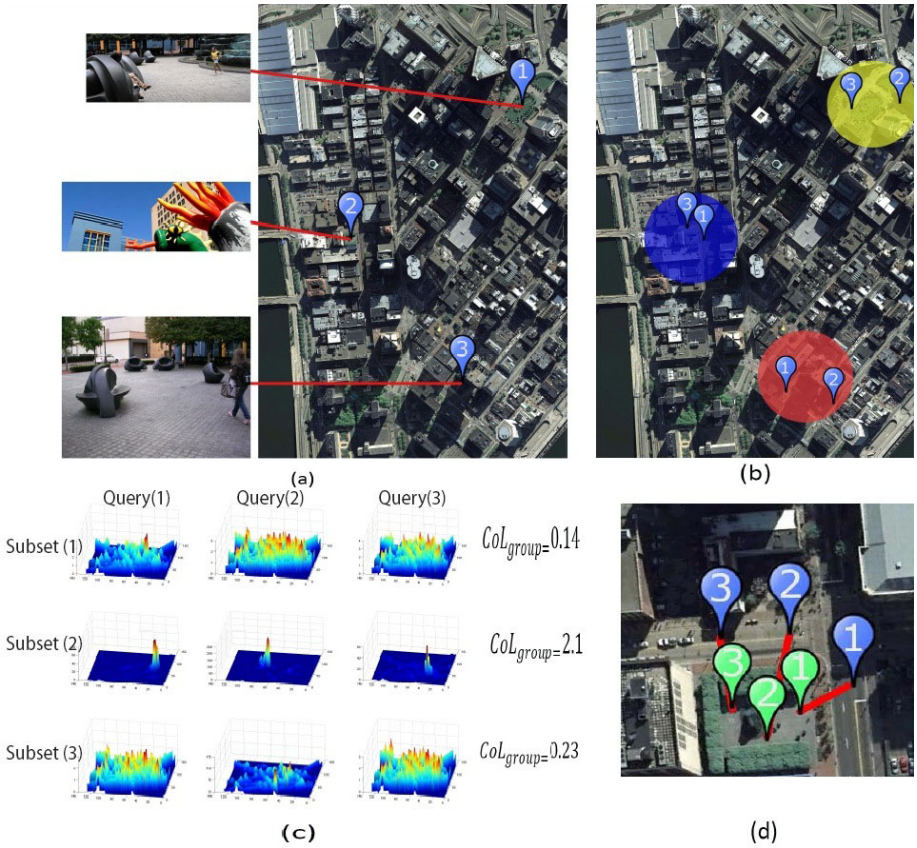


Fig. 7. An Example of Image Group Localization. (a):Query Images and Single Localization Results (b): Results of Search in Limited Subset. Each colored region is a different limited subset (c): Voting Surfaces and CoL_{group} for each query in each subset. (d): Blue Markers: Matched locations in the specific limited subset. Green markers represent the corresponding ground truth of queries. The red lines connect the ground truth with the respective correct match. The distances between the ground truth and final matched location are 10.2m, 15.7m and 11.4m, for queries 1, 2, and 3 respectively.

shows how a parameter representing the confidence of the localization can be of great assistance in preventing positive errors. More importantly, it can assist in extending the reference dataset as it may make reconstruction unnecessary.

5.2 Image Group Localization Results

Fig. 7 shows an example of localizing a set of images using the proposed method for geolocating image groups. The image group has 3 images, which are depicted on the left-hand side of Column (a). As discussed in Section 4, the first step of the proposed method is localization of images individually, resulting in a GPS

location for each image. Each query’s individual localization is displayed on the map in Column (a). Column (b) shows the result of applying a search within the limited subset created by the initial search in step 1; the other two query images are localized around the initial points found in Column (a). Column (c) shows the voting surfaces for each query in each subset. As can be seen, Subset (2) has the most distinct peaks across all three queries; correspondingly, Subset (2) also has the highest CoL_{group} value and is thus selected as the correct set of matches. Finally, Column (d) shows an inset of the map corresponding to Subset (2) with the matched images represented by blue markers and the ground truth locations for the queries represented by green markers.

As discussed earlier, there are 210 images in our test set for group image localization. Most of the images were selected as they are (individually) very unclear and therefore challenging to localize; this was done in order to show how proximity information can be extremely helpful in localizing images that are incapable of being geolocated individually. We set the parameter R to 300 meters for our tests; this is a conservative assumption. This means that we assume that the images in one group are all taken within 300 meters of each other. The right column of Fig. 5, compares the performance of Schindler et al. [4]’s method, our proposed single image localization method, and the group image localization method. As can be seen, the use of proximity information results in a drastic improvement. The right plot in Fig. 6 shows the breakdown of the results of the test set from the group image localization method based on the number of images in the groups. As mentioned earlier, this set consists of groups of 2, 3, 4 and 5 images. As can be seen in Fig. 6, the accuracy of localization for groups with a larger number of images is greater, due to the fact that groups with a larger number of images will search more limited subsets. Consequently the chance of finding the correct location is higher.

6 Conclusion

In this paper we addressed the problem of finding the exact GPS location of images. We leveraged a large-scale structured image dataset covering the whole 360° view captured automatically from Google Maps Street View. We proposed a method for geolocating single images, specifically examining how the accuracy of current localization methods degenerates when applied to large-scale problems. First, we indexed the SIFT descriptors of the reference images in a tree; said tree is later queried by the SIFT descriptors of a query image in order to find each individual query descriptor’s nearest neighbor. We proposed a dynamic pruning method which employed GPS locations to remove unreliable query descriptors if many similar reference descriptors exist in disparate areas. Surviving descriptors votes were then smoothed and then voted for the location their nearest neighbor reference descriptor belonged to. The reliability of the geolocation was represented by a proposed parameter called CoL , which was based on the Kurtosis of the vote distribution. Finally, a novel approach - using the proximity information of images - was proposed in order to localize groups of images. First, each image

in the image group was localized individually, followed by the localization of the rest of the images in the group within the neighborhood of the found location. Later, the location of each image within the rough area (Limited Subset) with the highest CoL_{group} value was selected as the exact location of each image.

References

1. Agarwal, S., Snavely, N., Simon, I., Seitz, S.M., Szeliski, R.: Building rome in a day. In: ICCV (2009)
2. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. ACM Trans. Graph. 25, 835–846 (2006)
3. Jacobs, N., Satkin, S., Roman, N., Speyer, R., Pless, R.: Geolocating static cameras. In: ICCV (2007)
4. Schindler, G., Brown, M., Szeliski, R.: City-scale location recognition. In: CVPR, pp. 1–7 (2007)
5. Nister, D., Stewenius, H.: Scalable recognition with a vocabulary tree. In: Proceedings of the, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2006, vol. 2, pp. 2161–2168. IEEE Computer Society, Los Alamitos (2006)
6. Crandall, D., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the world's photos. In: International World Wide Web Conference (2009)
7. Kalogerakis, E., Vesselova, O., Hays, J., Efros, A., Hertzmann, A.: Image sequence geolocation with human travel priors. In: ICCV (2009)
8. Zhang, W., Kosecka, J.: Image based localization in urban environments. In: 3DPVT 2006: Proceedings of the Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT 2006), pp. 33–40 (2006)
9. Hakeem, A., Vezzani, R., Shah, M., Cucchiara, R.: Estimating geospatial trajectory of a moving camera. In: International Conference on Pattern Recognition, vol. 2, pp. 82–87 (2006)
10. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. IJCV 60 (2004)
11. Vedaldi, A., Fulkerson, B.: VLFeat: An open and portable library of computer vision algorithms (2008), <http://www.vlfeat.org/>
12. Muja, M., Lowe, D.G.: Fast approximate nearest neighbors with automatic algorithm configuration. In: VISAPP (2009)
13. Balanda, K.P., MacGillivray, H.L.: Kurtosis: A critical review. The American Statistician 42, 111–119 (1988)