Do geographic features impact pictures location shared on the Web? Modeling photographic suitability in the Swiss Alps.

Produit Timothée[1], Tuia Devis[1], de Morsier Frank[2], Golay François[1]
[1] LaSIG laboratory [2] Signal Processing Laboratory (LTS5)
Ecole Polytechnique Fédérale de Lausanne
1015 Lausanne
Switzerland

ABSTRACT

Nowadays, millions of landscape images are uploaded on photosharing platforms such as Flickr or Panoramio. More and more of these images are also accurately geotagged via GPS devices mounted on personal cameras. Each image results from a twofold spatial choice: the choice of the location and the choice of the picture subject.

In this study, our focus is on landscape images in large touristic areas. Firstly, our goal is to learn which geographic features play a role in the choice of the location of shared images. For our analysis, we extract a series of geographic features from a Digital Elevation Model (DEM) and a Topographic Landscape Model (TLM) and model the photographic suitability as a density estimation problem in the space of the geographic features. Secondly, each combination of geographic features of a region is associated with a probability to be a location suitable for a photography. The resulting map is useful to promote tourism, to evaluate the landscape attractiveness or with a more technical objective as a prior in close-range photogrammetry.

This study shows that databases of geotagged pictures can be used to understand tourists behaviour also in rural areas, even if most of current researches are adressed to cities. The application to a touristic region in the Swiss Alps shows that the proposed method suits well this Geographic One-Class Data problem and is more accurate than both standard KPCA and One-Class SVM to model the suitability for touristic photography at locations unseen during training. As expected, picture locations are mostly correlated with geographic features extracted from the digital elevation model, as well as with those related to accessibility (distance to roads, paths). However, the force of this study is the combination of the geographic features via a kernel method to model more accurately suitable picture locations.

Categories and Subject Descriptors

H.2.8 [**Database management**]: Database Applications Spatial databases and GIS

In: S. Vrochidis, K. Karatzas, A. Karpinnen, A. Joly (eds.): Proceedings of the International Workshop on Environmental Multimedia Retrieval (EMR 2014), Glasgow, UK, April 1, 2014, published at http://ceur-ws.org

Keywords

Geographic One-Class data (GOCD), image geotags, landscape images, density mapping

1. INTRODUCTION

With the advent of "Web 2.0", the number of collaborative databases has dramatically increased. Very often, the databases objects are geotagged, meaning that an attribute related to its geographic location is available.

In this paper, we focus on collective pictures databases. Pictures uploaded on the Web via a photosharing platform (Flickr, Panoramio, Instagram...) have a time stamp, some textual tags describing its content and sometimes a world coordinate representing the geographic position. We are specially interested to the geotag which can be attributed in several ways. Most of the images are located with a click on a map. The accuracy of the localization is related to the zoom level used and to the ability to recognize an area in an aerial view. On the contrary, more and more cameras have a GPS device able to track very accurately the latitude and longitude coordinates. As stated in [7], the location of a picture is the result of a spatial choices: the choice of the location (the place where the photograph stands) and the choice of the subject (the object at which the camera points). This study focuses on the first choice: it benefits from the GPS accuracy to learn which geographic features describe the locations chosen.

Such a study aims at providing a map of the suitability to be picture location. This map can be useful in several ways. First, it could be used to promote tourism in areas sharing similar geographic feature with famous regions. Second, it can be used to assess landscape attractiveness, a measure needed in environmental planning [13]. Finally, shared photographies in landscape areas could be a valuable source to extract information about natural phenomena (displacement of glaciers [18], landscape change [3], local meteorology...): to meet this goal they require to be located and oriented accurately. Recent research shows that landscape images pose estimation (computation of the location and orientation of a camera in the computer vision vocabulary) requires priors which can be provided by GIS data: horizon and 3D models [1, 2, 4], aerial views [17], sun position [11]. The proposed map can thus be used to extract the most probable picture locations in a neighborhood.

However, the collaborative database are more widely used to analyse people behaviour and general trends in the tourists movements in urban areas [23, 16, 7]. Picture locations drawn on a map are difficult to read and require a more appropriate geovisualization. In [16, 7], spatial density maps (also called heat maps) are used to extract easily the most attractive areas (see for instance the site sightsmap.com). The textual geotags associated to a world coordinate also give many opportunities. For instance, authors in [8] use the tags to draw the geographic boundaries of fuzzy regions. They compare how Kernel Density Estimation (KDE) and Support Vector Machines (SVM) are accurate in the extraction of areas such as the Alps. Heat and density maps are well suited for large scale mapping. However, once zooming in, the contours of the high density regions become inaccurate, mainly because of the inaccuracy of the clicked geotags and the use of a smoothing radius. For instance, locations such cliffs can be considered highly attractive for photographers just because they are within the influence radius of popular places, while, in reality, they are inaccessible.

In order to interpolate probability values in each location of the map and not only in areas where picture are found, we propose to compute density in a space generated from geographic features rather than the space of latitude and longitude. To this goal, we require precise image locations which are provided by the GPS embedded in recent cameras and appropriate geographic features. Such geographic features are extracted with a GI Software either from a Digital Elevation Model (DEM) or from a Topographic Landscape Model (TLM: roads, lakes, forests).

By modeling the density of pictures in the space of geographic features, we estimate the probability of being a picture location over all the territory considered. This setting is equivalent to a One-Class problem (OC) where there is a set of positive data but no negative data: the locations of landscape photographies compose the positive set, the set of "attractive places". However, for the map locations with no pictures, we don't know if the absence of photographies is due to the inappropriateness of the location (a "bad" or negative place) or simply to the lack of pictures in an "attractive" location. This type of problems is also common in geographic classification problems, such as change detection from satellite images [14, 5]. Since in our case we consider a OC classification problem applied to geographical information sciences, Guo [9] called these kinds of data Geographic One-Class Data (GOCD).

During the last years, Kernel methods have been widely used for OC problems. In this study, we consider the following OC models: the One-class SVM (OCSVM) [21], the Support Vector Data Description (SVDD) [24] and the Kernel PCA (KPCA) [10]. All of them use the kernel trick to project the original data on a hypersphere in $\mathbb{R}^d$. In the high dimensional space, the data are more likely separated by a linear model [20]. Guo [9] compares OCSVM, Maximum Entropy (MAXTENT) [12] and Positive and Unlabelled Learning (PUL) [6] for GOCD problems. The two last methods provided the most accurate results. However, in our case study, the hypothesis which states that "*the probability of a negative data being labelled is null*" is not valid, thus excluding the use of PUL models. Indeed, our set of image locations contains some outliers, for instance pictures taken from a cable car or pictures not related to landscapes. Since such outliers are present, some of the labelled data belong to the negative class.

We will therefore focus on non-parametric methods for distribution support estimation. Such methods do not require knowledge about the distribution of the data and fit well with our problem. Besides support vector methods, another straightforward method is the Kernel Density Estimation (KDE, also known as Parzen window) [15, 19]: KDE estimates the density of data by applying a local smoothing filter [26].

In this paper, we propose a KDE-based strategy to estimate the probability distribution function (PDF) of the data. We estimate the support of the data both in the original space of geographic features and in the feature space spanned by KPCA. The PDFs of the labeled vs the unlabeled data are used to define a Bayesian criterion, which measures the probability for a map location to be an image location (given its geographic features). To fit the free parameters of KDE and KPCA we define a performance measure to ensure that most of the locations with photographies are classified as positives and most of the unlabeled locations are classified as negative. We compare the proposed approaches with OCSVM and KPCA on a real dataset of touristic pictures taken in the Swiss Alps.

2. METHODOLOGY

2.1 Geographic features extraction

The geographic features are extracted using a GIS software for an ensemble of N cells on a grid with 100m resolution. The considered features are summarized in table 1 and are computed for each cell z_j of the grid forming the unlabeled set $z_u = \{z_j\}_{i=1}^N$ and for the n image locations z_i, forming the labeled set $z_l = \{z_i\}_{i=1}^n$. Each z_i and z_j are thus represented by a d-dimensional vector formed by the d geographic features.

The first set of features is extracted from the DEM (Altitude, slope, curvature and visible sky). Then, for the TLM based features, the distances from the cell to the nearest forest, lake, road and cable car are computed. All the geographic features are mean centered and scaled unit variance.

2.2 GOCD with Kernel density Estimation

The KDE function in equation (1) is used to evaluate the density of a data set from the observations of the positive class z_l. Once the density has been estimated, we can evaluate the density for an unknown location z_j.

$$\hat{f}(z_j, z_l) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{z_j - z_i}{h}\right) \tag{1}$$

where $K(x)$ is a local smoothing operator, or kernel function. Among the different kernel functions, we used the Gaussian function:

$$K(x) = (2\pi)^{-d/2} exp(-\frac{1}{2}x^2) \tag{2}$$

where h is the bandwidth of the kernel function. Scott's rule is used to compute the bandwidth [22]:

$$h = n^{\frac{-1}{d+4}} \sigma_{z_l} \tag{3}$$

where σ_{z_l} is the covariance of the positive dataset. This rule uses the dimension d of the dataset and the number of positive data n to estimate a reasonable h. The choice of the appropriate kernel function has less influence on the results than the choice of the proper bandwidth. Indeed, if the bandwidth is too small, the density is over-fitted to the positive data set and its generalization power is weak. On the contrary, if the bandwidth is too large, the density will be oversmoothed and its small peaks will disappear.

We consider the following scheme to describe our GOCD problem. Let $Y \in 0,1$ be the event (or class) "*is a picture location*": $Y = 1$ for a cell being on a picture location and $Y = 0$ otherwise. The probability of a cell being a "picture location", given its geographic features z_j, is

$$P(Y = 1|z_j) = \frac{p(z_j|Y = 1)P(Y = 1)}{p(z_j)} \tag{4}$$

$P(Y = 1|z_j)$ is the value we want to compute for each cell of the map. The data density $p(z_j)$ is estimated with a KDE on a random set of unlabeled cells: $\hat{f}(z_j, z_u)$, the conditional probability $p(z_j|Y = 1)$ is estimated from a KDE on the labeled data only: $\hat{f}(z_j, z_l)$ and $P(Y = 1)$ is a unknown constant c. We are observing:

$$\frac{P(Y = 1|z_j)}{c} = \frac{p(z_j|Y = 1)}{p(z_j)} \tag{5}$$

A threshold can be set on the cell probability from equation 5 under which it is assumed to be drawn from the generic distribution $p(z_j)$, while above it is assumed drawn from the distribution of the labeled data $p(z_j, Y = 1)$. In practice, we set the threshold to one.

To chose the best set of parameters for the method, a performance measure need to be defined. To this end, our labelled set is divided in three subsets, the first one to train the KDE and the second to select the parameters. We want the estimated densities of the training and testing sets to be similar, nevertheless both should differ from the random set density. Indeed, if the density of the random set and the one issued from the labeled set are similar, the ratio for equation 5 will tend to one. This means that the geographic features where badly chosen and are not able to distinguish properly the random data from the labelled ones. In [9] the positive and unlabeled score presented in equation 6 is maximized:

$$F_{pu} = \frac{r^2}{r_{pos}} \tag{6}$$

Where r is the recall (the proportion of correctly predicted data in the testing set) and r_{pos} is the ratio of positive predicted locations in the random set.

In this study, we propose another criterion, more adapted for the density estimation process: we want to ensure that the bandwidth h fits well for both the training and testing sets and thereby their density should be similar. We want to maximize C in equation 7:

$$C = \frac{s_R^2}{s_T^2 - s_t^2} = \frac{s_R^2}{(s_T - s_t)(s_T + s_t)} \tag{7}$$

Where s_R is $1 - r_{pos}$; $s_T = (1 - r_T)$ with r_T being the recall for the training set and $s_t = (1 - r_t)$, with r_t being the recall for the testing set. Thus, C is very similar to F_{pu} but the component $(s_T - s_t)$ ensures that the PDF estimation of the training and testing set are similar.

Table 1: Geographic features descriptions and abbreviations

Abbr.	Description	Source
Z	Altitude	DEM, Swisstopo
Slope	Slope in percent	DEM, Swisstopo
Curv	Curvature	DEM, Swisstopo
Sky	Visible sky ratio, unobstructed hemisphere as a percentage.	DEM, Swisstopo
DRoad	Distance to the nearest road.	TLM, Swisstopo
BRoad	Presence or absence of roads (paths) in the cell.	TLM, Swisstopo
DFor	Distance to the nearest forest boundary, negative if the location is within a forest.	TLM, Swisstopo
DLake	Distance to the nearest lake.	TLM, Swisstopo
DLift	Distance to the nearest cable car - chair lift.	TLM, Swisstopo

2.3 KDE and KDE(KPCA)

In order to take into account the possible correlation between the geographic features and the non-linearity in the data distribution, we propose to estimate the density of the data either in the original space or in a feature space spanned by the mapping $\phi(.)$ induced by the KPCA. Hoffmann [10] states that the density function is proportional to the spherical potential in the feature space. The spherical potential measures the distance between $\phi(z)$: the projection of the point z in the feature space, and the center of the data $\phi_0(z)$. However, the spherical potential can't be used to estimate the density of the unlabeled data, because if the kernel projection works well to separate the labeled set z_l from the unlabeled set z_u, their gravity centers $\phi_0(z_l)$ and $\phi_0(z_u)$ do not correspond. Consequently, we run the KDE on the nonlinear features extracted from KPCA.

2.4 Comparing approaches: One-Class SVM and KPCA

The kernel functions project the data on a hypersphere which has the same dimension number than the number of training data. OCSVM searches for a hyperplan with two constraints. First, the intersection between the plan and the sphere enclose a ratio of the points equal to $1 - \mu$, where μ is the ratio of the outliers. Second, the plan has to be as far as possible from the origin. SVDD [24] searches for the minimal sphere which encloses the points. For "RBF" functions, OCSVM and SVDD can produce similar results [10]. KPCA applies a PCA on the projection of the data on the sphere. The reconstruction error in the feature space is used to separate positive and negative data. It appears that this boundary encloses more tightly the data, resulting in best results than OCSVM and SVDD.

As in [9], we use the positive and unlabeled F-score presented in equation 6 to evaluate the performance and select the best set of parameters for these two methods. Indeed, for One-Class problems, the commonly used performance measures, based on the ratio of false positives cannot be applied.

3. RESULTS

First, we will present the data and geographic features considered and how the image locations differ from the distribution of the random locations. Then, we conducted some experiment using different combination and different amount of geographic features. Finally, we will present the resulting probability maps.

3.1 Data

The experiments consider one of the political regions of Switzerland located in the Swiss Alps: "Valais - Wallis". The area is bounded by the Geneva Lake to the West and by the highest summits of Switzerland. It encloses some of the most touristic spots in Switzerland: Zermatt, the Matterhorn and the Aletsch glacier. The altitude gradient is important between the lowest area on the Geneva lake shore (450m) and the highest summit the "Dufourspitze" (4634m). The area is a valley, whose bottom hosts small to mid-sized villages. Climbing the flanks, the low altitudes are generally covered with vineyards (500-700m); then forest, mountain villages and resorts are found in the range from 700 to 1400m; above 1400m pastures and slopes dedicated to ski give access to the highest peaks, playground for the alpinists.

The Swiss Topographic Agency (Swisstopo) provides a DEM; among the available resolutions of 25m and 200m, we retained the first. Swisstopo also provides a TLM containing vector layers of several territorial objects. For this study, we selected the roads, other transportation facilities, forest and lake layers. The image locations are extracted from the Flickr database and were provided by [25]. Those images are filtered to keep only those located outside built areas and geotagged with a GPS device. As stated above, we want to learn which are the good places in term of landscape features: images taken in built areas tend to capture the presence of a village or a touristic attraction rather than a natural landscape. The set of image locations retained contains 2683 points, which are then separated in the three subsets represented in figure 1. The first set, in blue, is used to train the methods; the second one, in green, is used to fit the free parameters of the methods; the third one, in yellow, is used for the validation to compute independent statistics of the results. To generate those three sets, a grid with a 5km side is generated over the area. Then, each grid cell is randomly attributed to one of the subsets, in order to obtain spatially uncorrelated set of approximatively equal size. Finally, we also select 10 sets of random locations with the same number of locations as in the training set. These random sets will be used in the next sections to select significant geographic features and to compare the distribution of the labeled and unlabeled data.

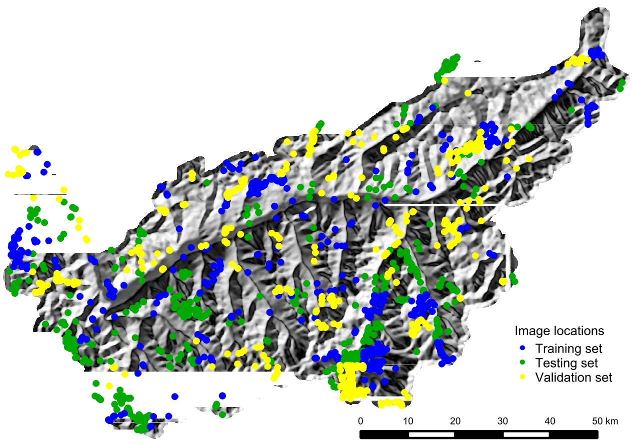

Figure 1: DEM hillshade rendering of the area under study overlaid by the Flickr image locations used: points colors depict whether the image is considered during training (blue), testing (green) or validation (yellow).

3.2 Geographic features selection

In order to measure the significance of these features, the unlabeled data and the picture location distributions are compared. For each of the geographic feature chosen, their distribution diverge (tested by a Kolmogorov-Smirnov test with $\alpha = 1\%$). Despite their statistical differences, some features are more discriminative than others. The following list presents the geographic features selected and explains how their value are different at true image locations.

- **Altitude (Z):** Since we are focusing on landscape photography, few images are taken between 450 and 1500m. In contrary, the range between 1800 to 2200m, where the ski slopes are found is very attractive. There is a small mode above 4000m representing pictures taken by alpinists at high altitude, figure 2 (a).
- **Curvature (Curv):** Positive curvatures (convex area) are more represented; indeed the ridges are more attractive than the valley.
- **Slope:** The flatter areas from 5% to 30%) are preferred.
- **Visible Sky (Sky):** This feature confirms the result observed with the curvature: cells with a high ratio of visible sky (>90%) are more often chosen.
- **Distance to nearest lake (DLake):** People tend to take more pictures in a radius of 200m around the lakes.
- **Distance to nearest road (DRoad) and presence of roads (BRoad):** The cells close to roads and paths are more active. Approximately 70% of the pictures are taken in a cell containing a road or a path, figure 2 (b).
- **Distance to the nearest forest (DFor):** The random and image distributions are very similar.
- **Distance to the nearest cable car (DLift):** Half of the pictures are taken within a range of 1500m surrounding a lift.

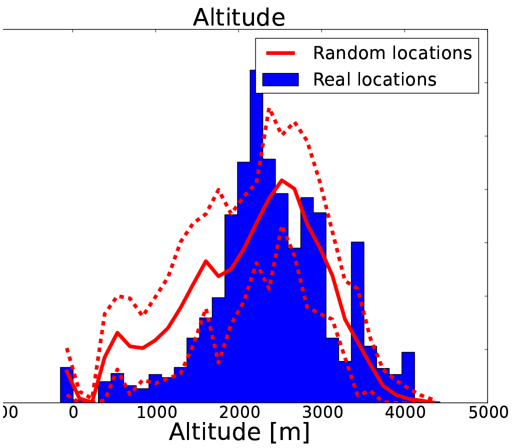

(a)

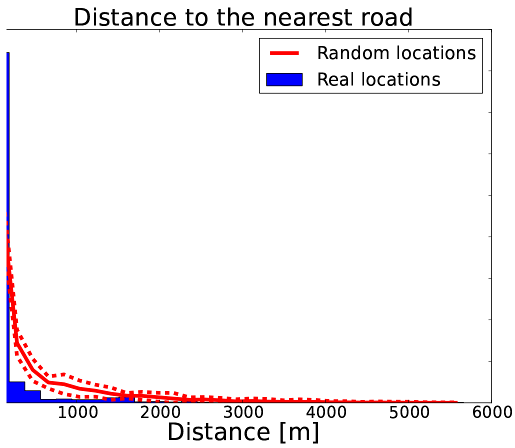

(b)

Figure 2: Distribution of the random and real locations for (a) Altitude, (b) Distance to the nearest road

3.3 Numerical results

In table 2, we report 10 experiments, obtained by considering different combination of geographic feature for GOCD. In this table only the best experiments are presented (C >= 20). Thus, the less-significant geographic features are less represented (DLift, DFor). The best combination of geographic features is altitude, slope, visible sky, curvature, distance to the nearest road and the binary roads (Exp. 9 in Table 2). The superiority of this experiment is observed for the two methods proposed but the KDE(KPCA) has the best result on the independent validation set, corresponding to a recall of 0.85. The KDE in the original space obtains slightly lower performance with a recall at 0.78. Thereby, it appears that the KDE and KDE(KPCA) have a similar behavior for the combination of few significant geographic feature (Exp. 1-4, 6, 7). However, if more features (Exp. 8 - 10) or binary features (Exp. 5) are added, the KDE(KPCA) is more suitable to describe the data relations between geographic features, resulting to better results. Both KDE and KDE(KPCA) are more suited than KPCA and OCSVM for this problem. By inserting the unlabeled data in the classification process, we ensure that less unlabeled data are classified positive and thus less data in the independent validation set are misclassified.

3.4 Evaluation on one of the most attractive area in the Alps: Zermatt.

On the map in figure 3 (a), probabilities for the Zermatt area are presented. The filled dots are the training locations, the triangles are the testing locations while the empty circles are the validation locations. Image locations are superposed to the results of KDE(KPCA) and correspond to the estimate of p(Y=1|z) at each location. A misclassification corresponds to a circle that would be located on an area of low probability (blue). The validation set is very specific in this area: the pictures on the south are above 4000m and accessible via a cable car (arrow **A** on the map in figure 3 (a)). This area is also a skiing region during winter and summer. First, some of the misclassified data are found along the cable car and are easily explicable. Indeed, these pictures aren't taken from the ground (but from the lift) and thus these locations aren't related to the geographic features. This also explains why the "DLift" geographic feature is not in the set of the best geographic features. Another set of locations, on the south of the "Breithorn" are badly classified (arrow **B** on the map in figure 3 (a)). They are shot on the way to this peak which is one of the most easily accessible 4000 summit in the Alps. However, the path to the summit is not in the roads / paths database. From this map, we can understand the geographic features related to the image locations. First, the mountain paths are easily recognizable in the figure. Indeed, they are always more attractive than the other locations. However, the paths on steep slopes are less probable than the other ones. Then, the ridges (Gornergrat), passes and summits (Matterhorn) are more attractive than other areas.

At the scale of the whole area, as seen in figure 4, it is interesting to note that the method is able to extract different behaviours. For instance, the paths are expected to have a large probability. By combining the geographic features, it appears that it is true at medium altitude. However, in the valley, where the roads have more traffic, locations between the roads are preferred. Moreover, at higher altitude, where ski slopes are found, people are more disposed to move away from the path. It is intuitive that the mountain peaks are good places to shot pictures. The strength of our method is to rank the peaks according to their altitude and shape. Finally, the less attractive locations are the steep slopes, bottom of deep valleys and flat areas at high altitude (glacier). Indeed, these regions are hardly accessible.

4. CONCLUSION

The GPS measured positions of shared images are more accurate than locations provided by the user. In this paper, we propose to use the GPS coordinates of a set of landscape pictures to train a classifier of likely and unelikely image locations. Every map location is described with geographic features extracted from a DEM and a TLM. The method proposed projects the geographic features in a space

Table 2: This table summarizes the results of our experiments. C is the score proposed in equation 7 to evaluate our methods: KDE and KDE(KPCA) , F_{pu} is the score from equation 6 used to evaluate KPCA and OCSVM. For the method using KPCA the number of Principal Component N_{pc} used is shown. nu is the ratio of outlier in OCSVM. $R(V)$ is the ratio of validation data correctly classified, this ratio is used to evaluate the methods. In bold, the best performing algorithm per experiment. Underlined the three best experiments over all tests.

	Geographic	**Proposed**					**Competing**					
Exp.	features	KDE		KDE(KPCA)			KPCA			OCSVM		
	considered	C	R(V)	N_{pc}	C	R(V)	N_{pc}	F_{pu}	R(V)	F_{pu}	ν	R(V)
1	Z, slope, Sky, DRoad	33.3	**0.83**	4	34.5	**0.83**	3	1.94	0.45	1.66	0.4	0.41
2	Z, slope, Sky, Broad	19.2	**0.73**	4	24.86	0.72	5	2.43	0.4	1.19	0.45	0.48
3	Z, slope, Sky, DRoad, Curv	20.43	0.83	7	35.6	**0.86**	6	1.59	0.57	1.23	0.35	0.52
4	Z, slope, Sky, DRoad, DLake	27.23	**0.78**	5	26.95	**0.78**	6	2.10	0.01	1.3	0.35	0.42
5	Z, slope, Sky, DRoad, BRoad	29.6	0.78	7	46.27	**0.84**	4	3.12	0.45	1.89	0.45	0.37
6	Z, slope, Sky, DRoad, DFor	18.8	**0.69**	4	25.7	0.55	4	1.52	0.55	1.18	0.25	0.64
7	Z, slope, Sky, DRoad, DLift	20.7	**0.75**	5	23.54	0.7	2	2.11	0.42	1.21	0.2	0.59
8	Z, slope, Sky, DRoad, Curv, DLake	21.85	0.79	7	34.44	**0.86**	6	1.52	0.53	1.16	0.25	0.63
9	Z, slope, Sky, DRoad, Curv, BRoad	35.26	0.78	7	51.74	**0.85**	5	2.31	0.55	1.44	0.35	0.52
10	Z, slope, Sky, DRoad, Curv, Broad, DLake	22.31	0.73	8	31.72	**0.81**	2	2.39	0.43	1.2	0.3	0.51

of higher dimension using a KPCA. Then, the spatial probability density function of the picture locations and random locations are estimated with a density function (KDE). The relation between them is used to compute the probability of each map cell to be a likely location given the geographic features. The recall on the independent validation set surpasses the KPCA and OCVM classification in their regular implementation.

This preliminary study could be refined in several ways. First, the geographic features computed were selected from *a priori* expected correlations. However, other geographic features could also be correlated to the image locations (orientation, DEM based features at finer or coarser scales, rivers, cable car departure and arrival stations only etc.). Second, our studies focus on a small area in the Alps, a similar approach could be applied to the entire Alps. Indeed, a bigger amount of picture locations would refine our results and increase their robustness and generalization power at a larger scale. Third, in this work we avoided the KDE bandwidth setting using the Scott's rule. A more appropriate bandwidth could improve the results. Finally, the improvement between the proposed method (KDE and KDE(KPCA)) and the standard ones (OCSVM, KPCA) is mainly due to the insertion of the unlabeled data in the classification process. Our method is easy to implement and shows good results without fine tuning.

We proved that the locations of a map are not equiprobable relative to landscape image locations. By describing each map location with geographic features, one can extract the most probable regions. The generated map differs from the density maps based on the Northing and Easting of the picture locations in two ways: firstly, using the geographic features, probabilities are also computed for region without image locations. Second, by taking into account the locations accessibility, the computed probabilities are more realistic. The force of our method is to learn attractive combinations of geographic features with the density estimation. Indeed, the relation between a geographic feature and an the picture location are intuitive (eg. paths are more probable, convex area are preferred...), but their combinations is more powerful to classify correctly picture locations.

Currently, there is a huge interest in computer vision for the pose estimation of shared images at a local or worldwide scale. Our results show that it is possible to use geographic

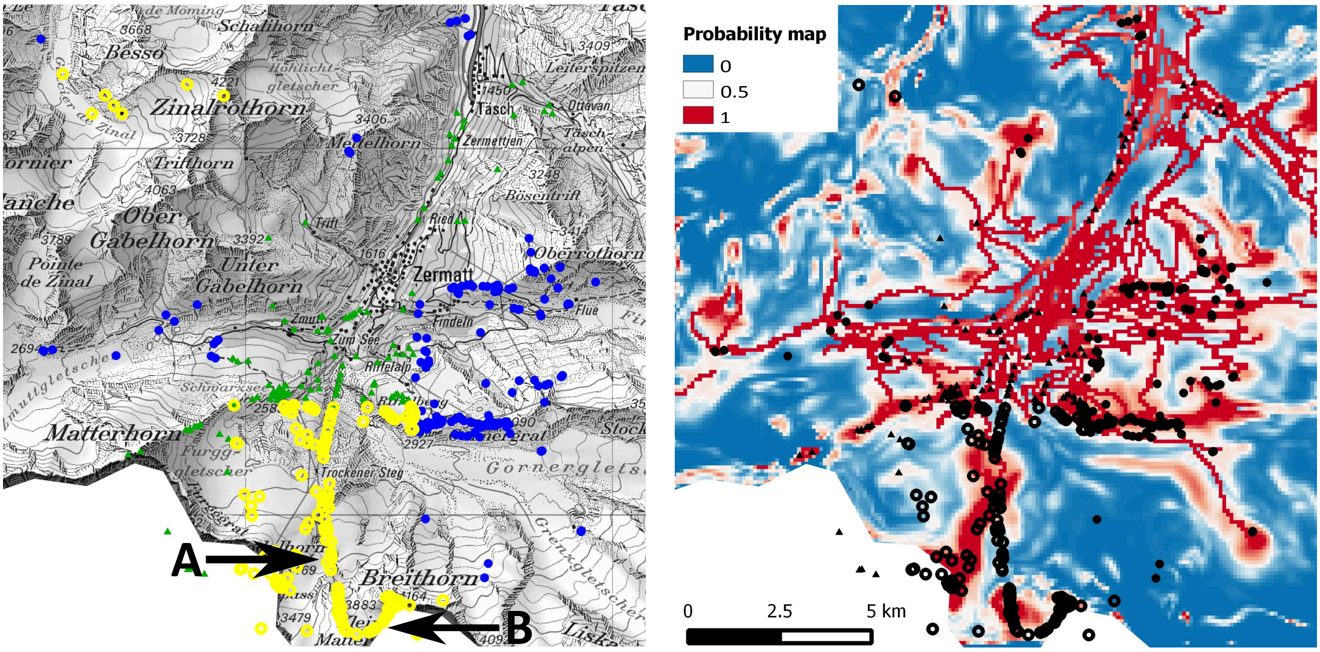

(a) (b)

Figure 3: Resulting maps: (a) Image locations are superimposed to the topographic map. Blue dots compose the training set, green triangles are the testing set, yellow circles are the validation set. "A" shows the cable car; "B" shows the way to the Breithorn. (b) Image locations are superimposed to probability map. Colour code is linked to the probability for a cell to be a picture location. Cells close to zero are similar to the random set, cells close to one are similar to the training set. Filled dots compose the training set, triangles are the testing set, empty circles are the validation set.

features as a prior knowledge to either discredit some unlikely poses or to promote the more probable ones.

5. ACKNOWLEDGMENTS

This work has been partly supported by the Swiss National Science Foundation (grant PZ00P2-136827).

6. REFERENCES

[1] G. Baatz, O. Saurer, K. Köser, and M. Pollefeys. Large scale visual geo-localization of images in mountainous terrain. In *ECCV*, pages 517–530. Springer, 2012.

[2] L. Baboud, M. Cadík, E. Eisemann, and H.-P. Seidel. Automatic photo-to-terrain alignment for the annotation of mountain pictures. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 41–48. IEEE, 2011.

[3] C. Bozzini, M. Conedera, and P. Krebs. A new tool for obtaining cartographic georeferenced data from single oblique photos. *Proceedings of the 23rd International CIPA Symposium*, 2011.

[4] P. Chippendale, M. Zanin, and C. Andreatta. Spatial and temporal attractiveness analysis through geo-referenced photo alignment. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, volume 2, pages 1116–1119, 2008.

[5] F. de Morsier, D. Tuia, M. Borgeaud, V. Gass, and J.-P. Thiran. Semi-supervised novelty detection using SVM entire regularization path. *IEEE Trans. Geosci. Remote Sens.*, 51(4):1939–1950, 2013.

[6] C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. In *International Conference on Knowledge Discovery and Data Mining*, 2008.

[7] F. Girardin, F. Calabrese, F. Fiore, C. Ratti, and J. Blat. Digital footprinting: Uncovering tourists with user-generated content. *Pervasive Computing, IEEE*, 7(4):36–43, 2008.

[8] C. Grothe and J. Schaab. Automated footprint generation from geotags with kernel density estimation and support vector machines. *Spatial Cognition & Computation*, 9(3):195–211, 2009.

[9] Q. Guo, W. Li, Y. Liu, and D. Tong. Predicting potential distributions of geographic events using one-class data: concepts and methods. *International Journal of Geographical Information Science*, 25(10):1697–1715, 2011.

[10] H. Hoffmann. Kernel PCA for novelty detection. *Pattern Recognition*, 40(3):863–874, 2007.

[11] N. Jacobs, S. Satkin, N. Roman, R. Speyer, and R. Pless. Geolocating static cameras. In *International Conference on Computer Vision*, pages 1–6. IEEE, 2007.

[12] E. T. Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

[13] P. S. Kane. Assessing landscape attractiveness: a comparative test of two new methods. *Applied Geography*, 1(2):77 – 96, 1981.

[14] J. Muñoz-Marí, F. Bovolo, L. Gómez-Chova,

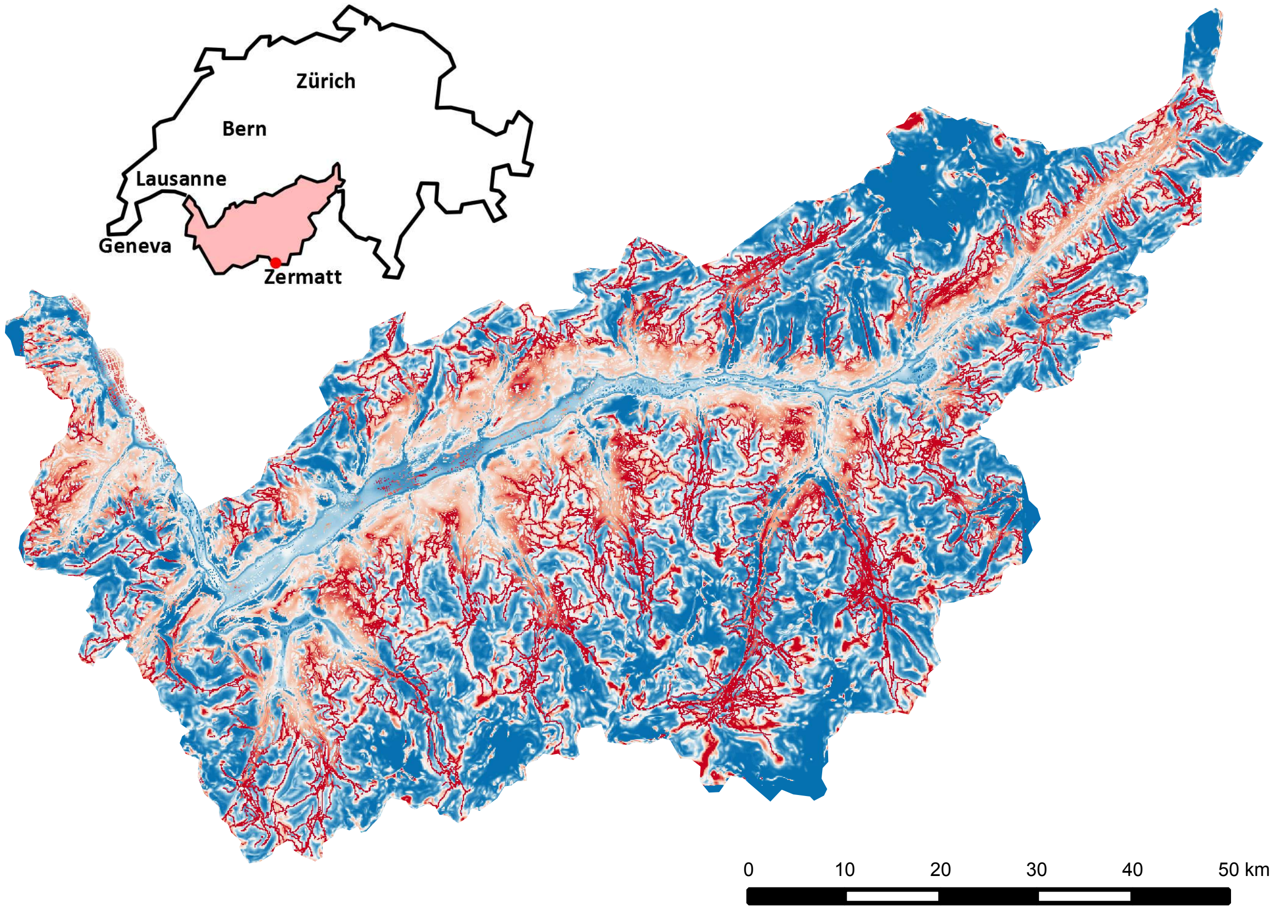

Figure 4: Resulting maps for the whole area. The case study area and Zermatt are highlighted on the Switzerland map. The colour code is the same as the previous map.

L. Bruzzone, and G. Camps-Valls. Semisupervised one-class support vector machines for classification of remote sensing data. *IEEE Trans. Geosci. Remote Sens.*, 48(8):3188–3197, 2010.

[15] E. Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 09 1962.

[16] A. Popescu and G. Grefenstette. Deducing trip related information from flickr. In *International Conference on World Wide Web*, WWW '09, pages 1183–1184, New York, NY, USA, 2009. ACM.

[17] T. Produit, D. Tuia, F. Golay, and C. Strecha. Pose estimation of landscape images using DEM and orthophotos. In *International Conference on Computer Vision in Remote Sensing*, pages 209–214. IEEE, 2012.

[18] T. Produit, D. Tuia, C. Strecha, and F. Golay. An open tool to register landscape oblique images and generate their synthetic model. In *Open Source Geospatial Research and Education Symposium*, 2013.

[19] M. Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, 27(3):832–837, 09 1956.

[20] B. Schölkopf and A. Smola. *Learning with Kernels.* MIT press, Cambridge (MA), 2002.

[21] B. Schölkopf, R. C. Williamson, A. J. Smola, J. Shawe-Taylor, and J. C. Platt. Support vector method for novelty detection. In *NIPS*, volume 12, pages 582–588, 1999.

[22] D. W. Scott Satk. On optimal and data-based histograms. *Biometrika*, 66(3):605–610, 1979.

[23] Y. Sun, H. Fan, M. Bakillah, and A. Zipf. Road-based travel recommendation using geo-tagged images. *Computers, Environment and Urban Systems*, (0), 2013.

[24] D. M. Tax and R. P. Duin. Support vector data description. *Machine learning*, 54(1):45–66, 2004.

[25] T. Wider, D. Palacio, and R. S. Purves. Georeferencing images using tags: application with flickr. In *AGILE International Conference on Geographic Information Science*, 2013.

[26] D.-Y. Yeung and C. Chow. Parzen-window network intrusion detectors. In *International Conference on Pattern Recognition*, volume 4, pages 385–388. IEEE, 2002.