

Visual Recognition and Fault Detection for Power Line Insulators

Markus Oberweger, Andreas Wendel¹, and Horst Bischof

Institute for Computer Graphics and Vision
Graz University of Technology, Austria

markus.oberweger@student.tugraz.at, {wendel, bischof}@icg.tugraz.at

Abstract *The inspection of high voltage power lines is an important task in order to prevent failure of the transmission system. In this work, we present a novel approach to detect insulators in aerial images and to analyze them automatically for possible faults. Our detection algorithm is based on discriminative training of local gradient-based descriptors and a subsequent voting scheme for localization. Further, we introduce an automatic extraction of the individual insulator caps and check them for faults by using a descriptor with elliptical spatial support.*

We demonstrate our approach on an evaluation set of 400 real-world insulator images captured from a helicopter and evaluate our results with respect to a manually created ground-truth. The performance of our insulator detector is comparable to other state-of-the-art object detectors and our insulator fault detection outperforms existing methods.

1 Introduction

High voltage power lines and transmission systems become more and more important with the raising demand of energy, especially in context of renewable resources. Pre-emptive inspection is an essential maintenance procedure in order to keep the downtime of a power line low, but it is time- and money-consuming, requiring much manual labor. Therefore we propose a machine-aided method for insulator inspection by automatically analyzing the images taken along a power line in order to determine faulty insulators, which are among the most common problems in transmission networks [21]. While there are different types of insulators, we focus on the most common *cap and pin insulator* with its characteristic stacked caps. The insulator can be applied to the pylon in different orientations, different sizes, or combined parallel or serial. We thus do not assume a certain orientation, combination or size in our method.

In this paper we present a novel recognition method for insulators in highly cluttered images (see Fig.1), and we introduce an automatic insulator fault detector. An overview illustrating the key-features of our work is shown in Fig. 2. The main contributions of this work are

1. an insulator detector, which is invariant to insulator orientation, size and combinations, partial overlap, illumination, and background clutter based on a circular descriptor and a noise-tolerant voting scheme, and




Figure 1: Our method detects insulators in highly cluttered aerial images and performs an automatic fault analysis. The faults are highlighted in red.

2. an automatic insulator fault detector, which automatically partitions each insulator into its individual caps and subsequently analyzes each cap for faults based on an elliptical descriptor.

We demonstrate the performance of our approach using an image set taken from a helicopter inspection, and evaluate the quality using a manually created groundtruth, which is, to the best of our knowledge, the first systematic evaluation.

2 Related Work

While different methods for detecting insulator faults exist, e.g. visual inspection or electrotechnical measurement, our method can be used complementary with other methods and especially for identifying mechanical damage and flashover marks. However, there is no inspection method or measurement device that is able to detect all possible insulator faults [14].

For insulator detection there are several works, as e.g. [4, 9], who use the detection as enabling method for further processing, but these methods' localization is too inaccurate for our work and restricted to a certain scenario, e.g. untextured background or a camera facing the sky, and thus not working well for our highly cluttered background. Opposing to these inaccurate methods, Kawamura et al. [11] published an approach based on 3D template matching for accurate 3D localization of insulators for robot interaction. The method

¹Current affiliation: Google Inc.

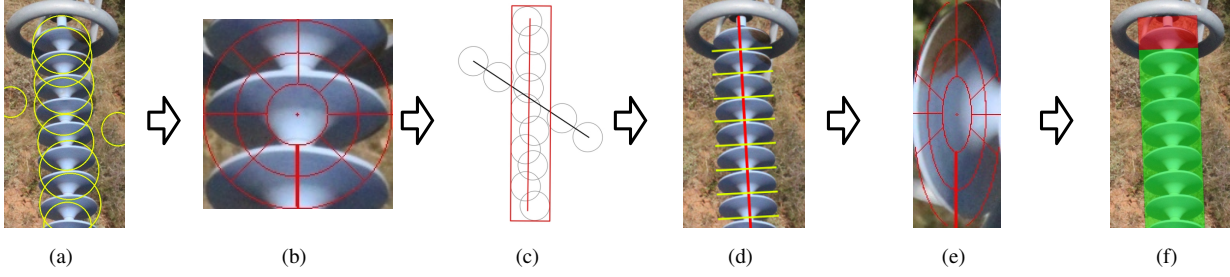


Figure 2: An overview of our method. Based on detected keypoints (a) we extract our proposed descriptor (b) and classify them as insulator cap or background clutter. On the classified keypoints we perform a RANSAC-based voting scheme to locate insulator detections (c). From these detections we compute the insulator partitioning (d) and extract our elliptical descriptor on individual caps (e). From these descriptors we determine faulty caps, which are then highlighted in the original image (f). Best viewed in color.

has high computational costs and is specifically designed for their task of 3D localization, which is not necessary and not possible in our case as we neither have range images nor a suitable model readily available.

A simple method to detect insulators is provided by Zhang et al. [26] who use color thresholding, but their method is only capable of detecting tempered glass insulators due to their characteristic color, and it needs a well adjusted threshold parameter which limits this method considerably.

Two approaches [24, 13] use edge descriptors for insulator detection. Both calculate the descriptors on a dense grid, which creates a high number of false positive detections resulting in a limited applicability in a cluttered environment. Further, edge descriptors are not discriminative enough, which is indicated by a high false positive rate.

A completely different method was proposed by Zhao et al. [27] who use a modified Markov Random Field to model the repetitive geometric structure of an insulator, which is more invariant to clutter. They have only shown their method for combined insulators in groups of two or four. In our case most of the insulators are attached solely to the power line and thus their method cannot initialize the geometric models and fails to detect them.

For the task of insulator fault detection, [16] and [8] proposed methods especially for dirt detection based on high resolution images which cannot be applied in our case, because our images are taken from a greater distance and thus the spatial resolution is too low.

For the detection of missing caps from aerial images Zhang et al. [25] proposed a method predicated on an accurate binary segmentation of the insulator provided by color thresholding, which is limited by the choice of the threshold. Further, they split the insulator into ten parts, but this static partitioning does not incorporate differently sized insulators or partially visible insulators.

Up to now, no work that provides a proper evaluation which could be used as baseline for our evaluation has been published, but only practical demonstrations. Next to the wide applicability of our method, this is another reason to provide a well documented baseline for further work.

3 Insulator Recognition

We first detect the insulators in the image and based on the detections we perform the fault detection. From a recognition point of view, insulators are weakly textured objects and in our case surrounded by clutter, which makes it hard to detect them. In contrast, insulators have a rigid form with repetitive geometric structure and a distinctive circular shape of each cap, which are properties that can be exploited. Therefore we use a part-based model with a tailored circular descriptor, where each insulator cap is one part of the model. The model geometry is a line segment (the major axis), where all caps belonging to the insulator must lie close to it and near other detected caps of the insulator. Therefore we detect Difference of Gaussians (DoG) [15] keypoints in the image and extract a square image patch around the keypoint according to the size of the keypoint, which fits very well to the actual cap size. From the patch we calculate our Circular GLOH-like (CGL) descriptor. It is similar to the GLOH descriptor [17], which is in turn a circular implementation of the prevalent SIFT descriptor [15]. Our descriptor is based on image gradients, which are derived by the Scharr operator [22]. This operator exhibits better rotational invariance than other gradient operators, which is beneficiary in our case as the caps are circular. The gradients are assigned to 17 spatial regions as visualized in Fig. 2b. Note that the central radial bin does not have any angular bins. Each gradient casts a vote according to its gradient magnitude in the 16-bin orientation histogram of its spatial region, resulting in 272 dimensions. For rotation and illumination invariance the methods of [15] are implemented. For scale invariance we enlarge the spatial support according to the size of the keypoint.

We employ Principal Component Analysis (PCA) to reduce the dimensionality of our descriptor for speedup but without loss of classification performance [23]. We calculate the orthogonal eigenspace from 11.3k cap descriptors and project the descriptors on the reduced space spanned by the 192 components with the largest eigenvalues.

We train a k-Nearest Neighbors (kNN) classifier with the descriptors of detected DoG keypoints from the training set. Keypoints within the ground-truth mask are positive sam-

ples and randomly selected keypoints from the background negative samples. For recognition we query the classifier with the descriptors of the detected keypoints in order to distinguish between insulator caps and clutter.

From the classified keypoints we determine the bounding boxes for the insulators. We group the keypoints by their scale and apply an adapted RANSAC [7] approach on all keypoints of each scale to robustly fit the insulator model to the detected keypoints. We have to modify the original algorithm to handle multiple insulator instances in an image [19]. Therefore we determine two random initial points $\mathbf{p}_1, \mathbf{p}_2$ from all available keypoints of one scale which satisfy the proximity constraint $\|\mathbf{p}_1 - \mathbf{p}_2\| < 4 \cdot sz(\mathbf{p}_1)$, where $sz(\mathbf{p})$ is the keypoint size. If no point combination satisfies the proximity constraint, the algorithm terminates as no model can be initialized. From these two points we create an initial model, i.e. the line segment connecting them. We add keypoints as inliers to the model, whose distance from point to line is smaller than half the keypoint size. All inliers have to be in relative proximity to already added inliers. This ensures that we do not add false positive detections located on the line. Each valid model must contain at least five inliers. We use the best model, thus having most inliers, as detection. The bounding box is created from the inliers by calculating the minimum bounding rectangle. The inliers are removed from the available keypoints, thus each keypoint is only used for one model. We repeat the process as long as a single iteration generates a valid model, else all insulators have been detected and the algorithm terminates. Using the described procedure we first optimize the recall of the detection. In the next step we use the estimation of the fundamental period of the insulator partitioning to verify the detections by evaluating the repetitive structure within the insulator, which is not present in false positive detections.

4 Insulator Fault Detection

The detected insulators are then analyzed for faults. Each insulator is described by its major axis and divided into its individual caps along this axis, as shown in Fig. 2d. Further, we calculate an elliptical descriptor from each cap to generate a score, which serves as a level of faultiness. By using a cap-wise partitioning of the insulator we can localize faults more accurately and invariant to differently sized and truncated insulators.

4.1 Insulator Partitioning

The first step of the insulator partitioning is the estimation of the overall orientation Θ . This orientation estimation and correction is important in order to provide a fixed layout for the following cap extraction. Therefore we use a binary segmentation mask obtained from the actual detection using GrabCut [20]. On the mask we apply an image moment-based method which exploits the elongated shape of the insulator. The moment M of order $(i+j)$ is calculated as [10]

$$M_{ij} = \sum_x \sum_y x^i y^j I(x, y) \quad (1)$$

where x and y are the image coordinates and $I(x, y)$ the pixel intensity. Using M we calculate the insulator orienta-

tion as

$$\Theta = \frac{1}{2} \text{atan2} \left(2 \left(\frac{M_{11}}{M_{00}} - \frac{M_{10}}{M_{00}} \frac{M_{01}}{M_{00}} \right), \frac{M_{20}}{M_{00}} - \left(\frac{M_{10}}{M_{00}} \right)^2 - \frac{M_{02}}{M_{00}} - \left(\frac{M_{01}}{M_{00}} \right)^2 \right). \quad (2)$$

The orientation is derived from the covariance matrix of the normalized second order image moments of the image mask.

In the second step of the partitioning we separate the insulator into its individual caps. Therefore we detect separation candidates by extracting Canny edges [3] and intersect them with the major axis of the insulator. Because these candidates are noisy, we employ signal processing methods to estimate and extract the separations. The separation candidates are formed as an impulse in a 1D signal $i[u]$ along the major axis u using

$$i[u] = \begin{cases} 1, & \text{if major axis intersects with edge} \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

The signal consists of impulses which are not properly aligned and thus create high frequency responses. Therefore we apply a Gaussian filter which provides a low pass filter and serves as noise estimate. For the Gaussian kernel we use $\sigma = \frac{1}{3} \lfloor \frac{1}{N} \sum_{i=1}^{N-1} d_{i,i+1} \rfloor$ where $d_{i,i+1}$ is the distance between two consecutive impulses. This kernel has been chosen to adaptively minimize the leverage on neighboring impulses due to filtering, which is necessary because of the high variance of cap sizes.

The fundamental period f of the filtered signal $x[u]$ is used to estimate the repetitive structure within the insulator. From $x[u]$ we calculate the N_{fft} -point Fast Fourier Transform (FFT) as $\mathbf{X} = \mathcal{F}\{x\}$ where $N_{fft} = 2^{\lceil \log_2(\text{length}(x[u])) \rceil}$. In order to improve the period estimate we use a weighting function $\mathbf{w}[k]$ to suppress unwanted frequency parts, which can be caused by wrong separation candidates. $\mathbf{w}[k]$ is again a Gaussian distribution with $\sigma = \frac{k_c}{\sqrt{2}}$ centered at $k_c = \lfloor \frac{N_{fft}}{w/2} \rfloor$ where w is the width of the insulator masks bounding box which serves as an estimate. Further, $f = \frac{N_{fft}}{k_{max}}$ and k_{max} is determined by

$$k_{max} = \arg \max_k (|\mathbf{X}| \cdot \mathbf{w}[k] \mid k > 3) \quad (4)$$

where $k > 3$ is used to suppress the dominant constant component of the signal. We estimate f on the major axis and on two lines parallel to this axis and use the median of the detected frequencies, which improves the robustness.

The alignment of the partitioning within the insulator is calculated using cross-correlation of the filtered input signal $x[u]$ and an idealized partitioning created from the fundamental period. The determined offset is used for matching the calculated to the nearest detected separations. If a separation is missing, e.g. not detected, a separation is thus automatically added at the most likely position.

In Fig. 3a the input signal $i[u]$ is compared to the resulting partitioning. Note the noisy input signal with multiple responses for each separation and with noise within the caps, which are removed in the result. In Fig. 3b the shown spectrum exhibits a strong peak at $k_{max} = 67$ for a $N_{fft} = 1024$ point FFT, which correctly results in $f = 15px$.

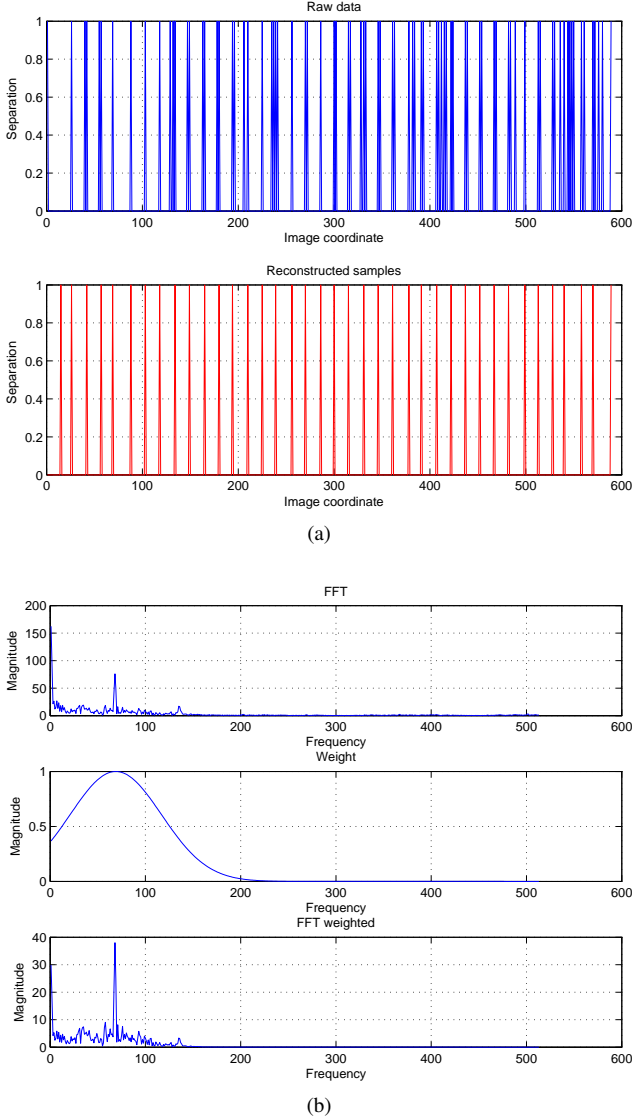


Figure 3: Insulator partitioning (measured $f = 15px$). (a) shows the reconstruction of the partitioning and (b) shows the estimation of the fundamental period.

4.2 EGL Descriptor

Based on the partitioning we extract an elliptical descriptor from each cap, which fits an insulator cap very well and thus minimizes the influence of neighboring caps or background clutter. In contrast to the CGL descriptor our Elliptical GLOH-like (EGL) descriptor contains gradient histograms for elliptical spatial regions (planar elliptical coordinates [18]). The spatial layout of the descriptor is shown in Fig. 2e, which again is divided into 17 regions with a 16-bin histogram of gradient orientations for each region. The axes of the ellipse are given by the size of the insulator cap. The EGL descriptor is not rotational invariant, which is not needed, because the insulator tilt can be corrected by using the insulator orientation. The illumination invariance and scale invariance are the same as with the CGL descriptor.

4.3 Fault Detection

Based on the extracted descriptors from each insulator cap we determine outliers, which are the faulty caps. We have to use an unsupervised outlier detector due to the lack of faulty caps for training and due to high intraclass variations of the insulators (background, viewpoint). Further, we do not treat the faults as binary classification problem, but we assign a score to each cap of an insulator. A higher score characterizes a higher dissimilarity to the other caps of the insulator and thus a more likely faulty cap. Therefore we use the Local Outlier Factor (LOF) approach proposed by Breunig et al. [2], which provides a score for the dissimilarity by using the distance of a descriptor to the k nearest neighbors as an estimate for the local descriptor density. We assume that an outlier has a lower local density in high dimensional space. The relation of this distance to the distance of its neighbors is used to identify outliers. As distance measure $d(A, B)$ we use the normalized L1 distance between two descriptors A and B in D -dimensional space,

$$d(A, B) = \sum_{i=1}^D \frac{|A_i - B_i|}{\max_i - \min_i} \quad (5)$$

where \max_i and \min_i are the minimal and maximal values of dimension i over all descriptors. We use $k = 3$ neighbors, initialize the distances with all descriptors of an insulator, and add each test descriptor separately to the set in order to preset the distances to a task-specific range [12].

The scores exhibit different ranges for each insulator. For a global representation we normalize the scores for each insulator to a range of $[0 \dots 1]$. Further, we enhance the distinctiveness by thresholding the score sequence and set all values that do not exceed the confidence level of $\frac{1}{\sqrt{2}}$ of the standard deviation above mean to 0. This provides a more discriminative visualization as shown in Fig. 2f.

5 Evaluation

In this section we present and discuss the results of the evaluation. To the best of our knowledge, there is currently no publicly available dataset for insulator detection. Therefore we use our own evaluation set, which contains 400 images (size $2352 \times 1568px$) with 375 labeled insulators, whereat 20

out of over 11.3k caps are labeled faulty (4 flashover damages and 16 cracked caps). For evaluation and training we use a segmented ground-truth, which is generated by manually segmenting the insulators using a GrabCut [20] based labeling tool. Truncated insulators with less than five visible caps are not labeled. We evaluate on a computer with a Core 2 Duo with 2.6GHz and 4GB of RAM.

5.1 Insulator Recognition

We first evaluate the keypoint classification and subsequently the recognition based on the classified keypoints. For the evaluation of the recognition there has not been published a work that provides proper evaluation metrics which could be used as baseline, but only practical demonstrations.

5.1.1 Keypoint Classification The evaluation of the keypoint classification is based on the true positive rate (TPR), the fraction of correctly identified caps to all caps, and the true negative rate (TNR), the fraction of correctly identified clutter to all clutter samples. For the keypoint classification a high TPR and TNR are required in order to distinguish between cap and clutter keypoints. From our evaluation set we automatically extract 11.3k caps located on DoG keypoints within the ground-truth mask and 11k clutter samples randomly selected from the background as ground-truth.

Tab. 1 shows the classification results for different descriptor types. The classification rates are obtained by using a kNN classifier with 2-fold cross-validation. Our CGL descriptor scores the highest TPR and TNR rate, thus it is the most suitable descriptor for this task. The high TNR of 99.7% is essential for efficient clutter suppression, but a slightly smaller TPR of 92.7% can be tolerated as not all caps of an insulator must be detected in order to detect the insulator itself due to our part-based model.

Descriptor	CGL	SIFT [15]	SURF [1]
TPR	92.7%	92.5%	75.8%
TNR	99.7%	92.4%	90.5%
Average runtime	1.4ms	1.9ms	0.25ms

Table 1: Average keypoint classification rates. Our descriptor (CGL) performs best.

The detection of keypoints works best with a DoG [15] detector. A dense grid (e.g. 7 scales from 20 to 100px) is not applicable because of the high runtime and the high number of false positive detections. Other detectors as e.g. an approximated Hessian detector [1] cannot locate the insulator caps accurately, or detectors that detect corner-like structures rather locate the keypoints on the boundaries than in the center of the caps.

5.1.2 Recognition We evaluate the insulator detector on our evaluation set, where each connected component in the ground-truth mask is determined as insulator and the minimal bounding rectangle is used as ground-truth bounding box. Note that these are rotated rectangles, therefore they fit the insulators very well.

In order to evaluate the localization of our method we use

the well-known Pascal score [5], which is calculated from the overlap of our generated bounding box B_c to the ground-truth B_{gt} by

$$p(B_c, B_{gt}) = \frac{\text{area}(B_c \cap B_{gt})}{\text{area}(B_c \cup B_{gt})}. \quad (6)$$

An object is considered detected if $p(B_c, B_{gt}) > 0.5$.

As objective we want to maximize the number of correct detections and minimize the number of false detections. The two used evaluation metrics are precision, the fraction of correct detections to the total number of detections made by our detector, and recall, the fraction of correctly detected objects to the number of annotated objects. Our detector provides a score for each detection, which is used to vary the trade off between these two metrics. The score is calculated by

$$\text{score}(\mathbf{p}, L) = \left(1 - \frac{\sum_{i=1}^N d_{\mathbf{p},L}(\mathbf{p}_i, L)}{\sum_{i=1}^N \text{sz}(\mathbf{p}_i)/2} \right) \cdot N \quad (7)$$

where \mathbf{p} is the inlier set, L the estimated model, N is the number of inliers and $d_{\mathbf{p},L}$ the Euclidean distance from inlier to model. The score is higher if a model contains more inliers or better fitting inliers.

We use the Precision-Recall Curve (PRC) as performance measurement, which is shown in Fig. 4 for our detector. The curve is constructed in accordance to [5] by using the interpolated precision. We perform 2-fold cross-validation and plot the averaged precision and recall values. The recall reaches a maximum of 98% and drops rapidly for higher recall scores. Most false positives are caused by insufficient overlap of the detected to the ground-truth bounding box. Only about 7% of the false positives are actually located on the background or on the pylons, thus showing efficient clutter suppression.




Figure 4: PRC for insulator detection. The average precision is over 56%.

Badly detected boxes are identified by a low overlap value as shown in Fig. 5. For our evaluation the required Pascal criterion is very strict due to the fact that we use rotated bounding boxes, which is not originally intended [5].

Subjectively speaking, an overlap score of 0.5 already fits the insulator very well, whereas a lower threshold might be concerned, e.g. $p(B_c, B_{gt}) > 0.4$ already improves the recall to 100%. Using an overlap threshold of 0.5 the insulators in Fig. 5 (a)-(c) would be considered correct detections, and Fig. 5 (d)-(f) false detections although the first two fit well. Only the detection in Fig. 5f completely fails, because the line model is initialized badly and further the detected insulator caps in the upper part are not included into the model due to a large distance from the model.




Figure 5: Detection results (red) with overlap score $p(B_c, B_{gt})$ stated. Best viewed in color.

5.2 Insulator Fault Detection

For the fault detection, a proper insulator partitioning is essential. Therefore we first evaluate the insulator partitioning and then the fault detection itself. Fig. 1 shows the fault detection results for a mechanical damage and Fig. 2f for a flashover damage.

5.2.1 Insulator Partitioning For the evaluation of the partitioning we compare the calculated orientation to the orientation of the minimum bounding rectangle of the ground-truth segmentation mask. We have implemented the Hough-based method of [25] as baseline. As evaluation criterion we use the angular error $e(\Theta, \Theta_{gt}) = |\Theta - \Theta_{gt}| \pmod{\frac{\pi}{2}}$, where Θ_{gt} is the orientation of the ground-truth and Θ the orientation provided by the method. Only the absolute orientation matters due to the symmetric shape of the insulator. The mean $\mu(e)$ and the standard deviation $\sigma(e)$ of the angular error are used to compare the methods. A smaller mean indicates a more accurate orientation estimation and a smaller standard deviation shows, that the errors are less scattered.

The evaluation results in Tab. 2 show, that our moment-based method is more accurate and requires less computation time than the Hough-based approach. The orientation of the bounding box of the actual detection requires no additional computation, but it is less accurate than our proposed method and prone to outliers. A more precise orientation provides a higher accuracy of the partitioning.

For the evaluation of the partitioning we measured the fundamental period from the images and use it as ground-truth. The fundamental period is correctly computed for 84% of all insulators, w.r.t. a tolerable deviation of $\pm 5px$ or 10%. For the cap extraction the creation of a ground-truth

Method	$\mu(e)$	$\sigma(e)$	Average runtime
Moment-based	0.33°	0.48°	12ms
Hough-based [25]	1.57°	2.03°	430ms
Bounding box	1.02°	1.24°	0ms

Table 2: Insulator orientation evaluation.

is not feasible, thus resulting in manual checks. A proper partitioning requires a cap width of at least $10px$, otherwise the separation features between the individual caps vanish. The partitioning works for different perspectives as shown in Fig. 6a and 6b, but fails if there are no separation features as depicted in Fig. 6c. For that sample also humans fail to separate the caps.

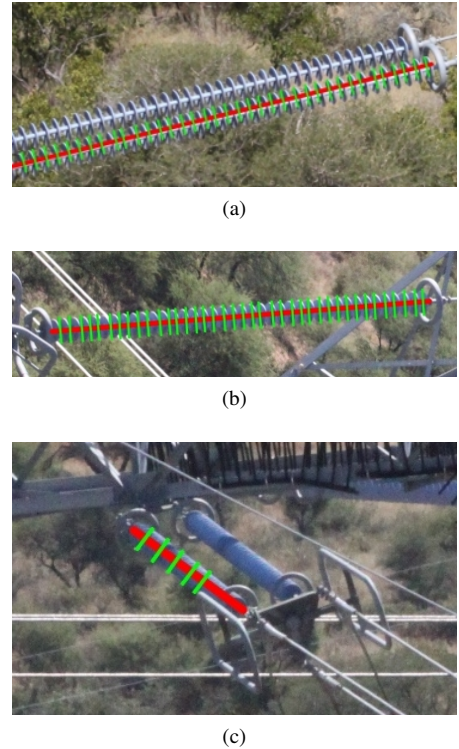


Figure 6: Partitioning results. Note the different sizes (red line is $5px$ wide). Best viewed in color.

5.2.2 Fault Detection We compare our fault detection method to the approach of [25] based on Gray-Level Co-occurrence Matrices (GLCM). In order to make their scores comparable we use our partitioning to calculate their features, instead of constantly ten parts.

A detailed view of a score sequence is given in Fig. 7, which shows the scores for a flashover damage (see Fig. 2e) at cap 1. The scores are normalized, but not thresholded as described in Section 4.3 for a better illustration. The ground-truth has a value of 1 for a faulty cap and 0 for non-faulty caps. Both methods score high values for the faulty caps, but the values scored for non-faulty caps are quite different. By using our EGL-based approach, we efficiently suppress background clutter and are invariant to faults in the segmentation mask, in contrast to the GLCM-based method, which

cannot distinguish between a faulty cap and a defect in the segmentation mask.




Figure 7: Score sequence for a flashover damage at cap 1. Our LOF-based approach can clearly distinguish between faulty and non-faulty caps.

For the evaluation of our fault detection we use the Receiver Operating Characteristic (ROC), which is independent of class skew [6] and thus advantageous in the case of fault detection because the faulty samples make up only a fraction of all samples. Although we use a continuous score for the fault detection, the evaluation can be seen as a binary classification problem: faulty caps are the desired positive class, non-faulty caps are the negative class. As performance measurements we use the TPR, the fraction of successfully detected faulty caps to all faulty caps, and the false positive rate (FPR), the fraction of mistakenly identified faulty caps to all non-faulty caps. A comparison of different ROC for insulator fault detection is shown in Fig. 8. The ROC is created by using the fault score as threshold. Our EGL descriptor with LOF performs best and achieves a TPR of 95% at a FPR of 12%.

6 Conclusion

In this work we have presented a novel approach for insulator recognition and a subsequent automatic fault detection from aerial images. We introduced a method for insulator recognition using a part-based model with local image features and a RANSAC-based clustering approach, which enables the detection of insulators in a highly cluttered environment regardless of their orientation, size or combinations.

Further, we proposed a method for insulator fault detection based on a descriptor with elliptical spatial support. We used LOF to assign a score of faultiness to each insulator cap extracted by our automatic partitioning algorithm. This provides an accurate cap-wise fault assessment of the insulator under different photometric and geometric conditions.

Both methods have been evaluated thoroughly and we have

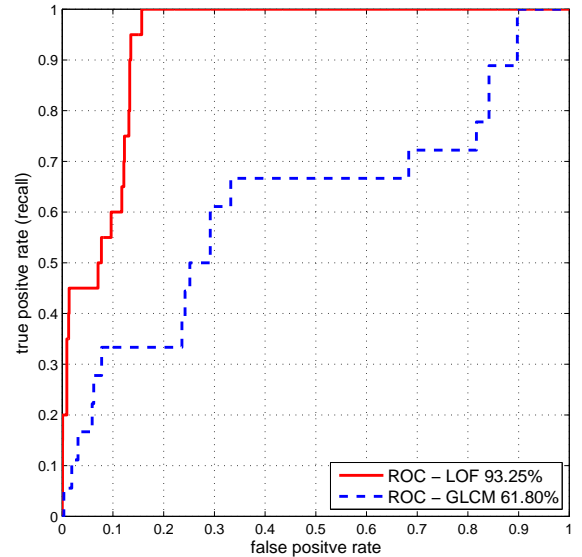


Figure 8: ROC for fault detection, with area under the curve stated in the caption. Our proposed method (LOF with EGL descriptor) significantly outperforms GLCM [25].

established a baseline, which can be used for comparison in future works.

Acknowledgement

This work has been supported by the Austrian Research Promotion Agency (FFG) project FIT-IT Pegasus (825841/10397).

References

- [1] Herbert Bay, Tinne Tuytelaars, and Luc J. Van Gool. SURF: Speeded Up Robust Features. In *Proc. ECCV*, 2006.
- [2] Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. LOF: Identifying density-based local outliers. In *Proc. SIGMOD*, 2000.
- [3] John Canny. A computational approach to edge detection. *Trans. PAMI*, 8(6):679–698, 1986.
- [4] Jaime Del-Cerro, Antonio Barrientos, Pascual Campoy, and Pedro J. García. An autonomous helicopter guided by computer vision for inspection of overhead power cable. In *Proc. IROS Workshops*, 2002.
- [5] Mark Everingham, Luc J. Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *IJCV*, 88(2):303–338, 2010.
- [6] Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.
- [7] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM*, 24(6):381–395, 1981.
- [8] Yumin Ge, Baoshu Li, Shutao Zhao, and Chengzong Pang. Detection of the insulator dirtiness based on the computer vision. In *Proc. CICED*, 2006.

- [9] Irene Y. H. Gu, Unai Sistiaga, Sonja M. Berlijn, and Anders Fahlström. Online detection of snowcoverage and swing angles of electrical insulators on power transmission lines using videos. In *Proc. ICIP*, 2009.
- [10] Ming-Kuei Hu. Visual pattern recognition by moment invariants. *Trans. Information Theory*, 8(2):179–187, 1962.
- [11] Kentaro Kawamura, Mark D. Wheeler, Osamu Yamashita, Yoichi Sato, and Katsushi Ikeuchi. Localization of insulators in electric distribution systems by using 3D template matching from multiple range images. In *Proc. IROS*, 1998.
- [12] Aleksandar Lazarevic and Vipin Kumar. Feature bagging for outlier detection. In *Proc. ICKDDM*, 2005.
- [13] Weiguo Li, Gaosheng Ye, Feng Huang, Shikun Wang, and Wenzhi Chang. Recognition of insulator based on developed MPEG-7 texture feature. In *Proc. ICISP*, 2010.
- [14] Zheng Li and Yi Ruan. Fault diagnosis system for the inspection robot in power transmission lines maintenance. In *Proc. OIT*, 2009.
- [15] David G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [16] Xin Mei, Tiecheng Lu, Xiaoyun Wu, and Bo Zhang. Insulator surface dirt image detection technology based on improved watershed algorithm. In *Proc. PEEC*, 2012.
- [17] Krystian Mikolajczyk and Cordelia Schmid. A performance evaluation of local descriptors. *Trans. PAMI*, 27(10):1615–1630, 2005.
- [18] Parry Hiram Moon and Domina Eberle Spencer. *Field Theory Handbook: Including Coordinate Systems, Differential Equations and Their Solutions*. Springer, 1971.
- [19] Julien Rabin, Julie Delon, Yann Gousseau, and Lionel Moisan. MAC-RANSAC: a robust algorithm for the recognition of multiple objects. In *Proc. 3DPVT*, 2010.
- [20] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM TOG*, 23(3):309–314, 2004.
- [21] Murari Mohan Saha, Jan Jozef Izykowski, and Eugeniusz Rosolowski. *Fault Location on Power Networks*. Springer, 2010.
- [22] Hanno Scharr. *Optimale Operatoren in der digitalen Bildverarbeitung*. PhD thesis, Ruprecht-Karls-Universitt Heidelberg, 2000.
- [23] Simon Winder, Gang Hua, and Matthew Brown. Picking the best DAISY. In *Proc. CVPR*, 2009.
- [24] Jian Zhang and Ruqing Yang. Insulators recognition for 220kv/330kv high-voltage live-line cleaning robot. In *Proc. ICPR*, 2006.
- [25] Xinye Zhang, Jubai An, and Fangming Chen. A method of insulator fault detection from airborne images. In *Proc. GCIS*, 2010.
- [26] Xinye Zhang, Jubai An, and Fangming Chen. A simple method of tempered glass insulator recognition from airborne image. In *Proc. ICOIP*, 2010.
- [27] Jingjing Zhao, Xingtong Liu, Jixiang Sun, and Lin Lei. Detecting insulators in the image of overhead transmission lines. In *Proc. ICIC*, 2012.