

Robust crack detection strategies for aerial inspection

Emanuel Aldea^a and Sylvie le Hégarat^a

^aAutonomous Systems Department, IEF, Paris Sud University, France

ABSTRACT

In this work, we evaluate the relevance of current state of the art algorithms widely employed in the detection of cracks, for the specific context of aerial inspection, which is characterized by image quality degradation. In this study we focus on minimal cost path and on Marked Point Process algorithms, and we test their resilience to motion blur. The results show that the current strategies for defect detection are sensitive to the quality of input images; alternatively, we suggest some improvements based on a-contrario methods that are able to cope with significant motion blur.

Keywords: Crack detection, fissure detection, a-contrario strategies, FFA, MPP, aerial inspection

1. INTRODUCTION

Either for road quality assessment, for the examination of buildings or for the examination of critical infrastructure, an automated framework aiming at the detection of surface defects (cracks, pitting corrosion, rutting etc.) may have a significant positive impact in terms of cutting down maintenance costs and enhancing security. One major reason supporting automation is that human supervision is prohibitively costly for large scale, frequent tasks of infrastructure surveillance. Among surface defects, cracks exhibit a high level of pervasiveness on different mediums which are affected by various phenomena (thermal phenomena, fatigue etc). However, their relatively similar appearance makes image processing an ideal tool for detection, all the more so as cameras are affordable sensors, which are also ideal for non destructive evaluation (NDE).

For these reasons, there is a large body of work aiming at automating the problem of crack detection. More recently, robotic systems and unmanned aerial vehicles (UAV) further increased the quantity of data to be processed since they have the ability to approach scenes which were otherwise difficult to access or even unreachable in safe conditions. However, this approach has the main disadvantage of degrading the image quality as a result of various phenomena specific to image recording in the presence of significant camera dynamics, such as motion blur, defocus or fast illumination variations. These degradation represents a significant drawback since it has a direct impact on the salience of cracks, which are often challenging to detect. We note in the literature two main trends handling crack detection from different perspectives; the first body of work deals with crack detection as a pattern recognition problem (see for example,^{1,2} or³ for a review), and assumes ideal or at least adequate image quality. The second type of work is focused on validating the functionality of a robotic examination system,⁴⁻⁷ and few works⁸ consider in their experimental validation the correlation between acquisition factors and the deterioration in detection performance; usually the experimental evaluation is performed using basic processing tasks on images acquired in favorable conditions.

Our work is aimed at bridging the current gap between robust pattern recognition algorithms which are suited for complex detection tasks, and their evaluation in conditions which are more representative for NDE using automated systems. For crack detection, we consider three fundamentally different algorithms. These methods have been selected as belonging to the state-of-art either of the considered application (crack detection) or on other applications (winkles detection) but that, from our point of view, is very close in terms of pattern recognition objective and assumptions. Indeed, as mentioned by several authors, the developed methodology for crack detection also applies to problems as different as road extraction in remote sensing images, or blood vessel detection in medical images.

Further author information: (Send correspondence to E.A.)

E.A.: E-mail: emanuel.aldea@u-psud.fr, Telephone: +33 169 154 127

S.l.H.: E-mail: sylvie.le-hegarat@u-psud.fr, Telephone: +33 169 154 036

2. RELATED WORK

2.1 Minimum path based strategies

Free form anisotropy (FFA) is a defect detection approach² which is based on the assumption that a defect pixel may be characterized by a minimal path traversing the pixel. Furthermore, for a given pixel location l we consider a set of four constrained minimal paths traversing the location, i.e up-down, left-right, upper left - lower right and upper right - lower left. The directional paths are imposed each time by two opposing pairs of 3-connectivity oriented neighborhoods, and we assume that all the paths are of length N and centered in l . If we assume that a fissure corresponds to a minimal path which exhibits statistical properties which are different from those of the other three “regular” minimal paths, we dispose of a crack detection strategy which does not require any training data.

Assuming a normal distribution of radiometric values of the pixels along the minimal paths, let us denote by $s_1 = (\mu_1, \sigma_1^2)$ and by $s_2 = (\mu_2, \sigma_2^2)$ the mean and variance for the minimal cost path (among the four), and respectively for the three others combined. The authors of² rely on a degree of coherence h between two sources, adopted from the possibility theory. A low degree of coherence, close to 0, indicates that the first source is different from the second, and thus in our situation a low value of h implies that the pixel may very well be located on a genuine fissure.

A final step is necessary in order to get a binary assignment at pixel level related to the presence of a defect. For the input image locations $l_i \in I$, $h(l_i) \in [0, 1]$ and thus a thresholding step is required. The authors of² apply a hysteresis threshold which is performed using a multi-class variant of the Otsu algorithm.

2.1.1 Adjustments for our problem

Firstly, the post-processing of the thresholding step does not fall within the scope of our current work. We are here more interested in estimating the robustness of the main step of the algorithm rather than in the thinning of the connected components or in interconnecting them even further. For this main part however, we replaced the degree of coherence as a distance between two distributions may be approached straightforwardly by computing the Bhattacharyya distance $D_B(s_1, s_2) \in [0, \infty)$ and then the Bhattacharyya coefficient $BC(s_1, s_2) = e^{-D_B(s_1, s_2)}$ that we will employ for the same purposes as h .

2.2 Marked Point Process models

A direction of research which has emerged for the application of wrinkle detection is the exploitation of Marked Point Processes (MPP), a method previously used in various applications such as road extraction, tree detection or crowd analysis.⁹⁻¹¹ This strategy models an observation using a set of base elements denoted as points, which belong to a compact space $\mathcal{F} \in \mathbb{R}^d$. A point process \mathbf{X} defines a probability distribution onto the configuration space $\Omega = \cup_{n=0}^{\infty} \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n\}, \mathbf{p}_i \in \mathcal{F}$. For MPP, additional cues denoted as marks are associated to the elements of \mathcal{F} ; these marks aim to assimilate geometrical information which is relevant for the appearance of the object that we intend to perceive within the observed data. In order to be able to identify the optimal solution in the configuration space, the probability distribution is required to favor configurations which agree more with a prior appearance, by means of an energy term, and for which at the same time the location of the points is validated by the data profile, by means of a data fidelity term.

Thus, crack or fissure detection may be performed using MPP if one defines accordingly the characteristics and the interactions among the marked points for this specific task. In the literature, a related problem, namely the detection of facial wrinkles, has been approached successfully using this strategy. The proposed algorithms^{12,13} model wrinkles as sequences of interconnected line segments $w_i = (\mathbf{p}_i, \mathbf{m}_i)$, where $\mathbf{p}_i = (x_i, y_i)$ encodes the center of the segment, while the mark $\mathbf{m}_i = (l_i, \theta_i)$ defines the relevant features of the respective segment, its length and orientation. The approach is well adapted for the task of crack detection since the constraints imposed by the prior model for wrinkles are easily transposable to cracks. Let us recall briefly below the probability distribution which is used over the MPP configuration space.

2.2.1 The probability density for the MPP based crack detection

Generally, the probability density for a given configuration \mathbf{w} of segments is introduced as the product between a data likelihood term and a prior energy:

$$f(\mathbf{w}) \propto h_d(\mathbf{w})h_p(\mathbf{w})$$

Similarly to the work of¹³ we prefer to employ a steerable filter for the data fidelity term in the following way. A basis set for a pixel location l is given by

$$\mathbf{G}(l) = \left[\frac{\partial^2 g(l)}{\partial x^2} \quad \frac{\partial^2 g(l)}{\partial x \partial y} \quad \frac{\partial^2 g(l)}{\partial y^2} \right]$$

where $g(l)$ denotes a Gaussian function. A steerable filter associated with a location l and with a direction θ is given by

$$g_\theta(l) = [\cos^2 \theta - \sin 2\theta \sin^2 \theta]^T \mathbf{G}(l) \quad (1)$$

Finally, the data fidelity term for a segment w will be simply the accumulation of the filter response of the constituent pixels with respect to a given direction: $U_d(w) = -\sum_{l \in w} g_\theta * I(l)$.

The energy term we employ follows the works of^{12,13} and basically it defines the interaction between adjacent segments by encouraging smoothly connected lines, and also by rejecting occurrences of abrupt intersection or overlapping through a repulsive term.

For our work, we have simplified the segment selection procedure by adopting a guided (biased) drawing process, that would select segments initially based on seeds which are strongly correlated to the data fidelity term, and which would iteratively augment and fusion them according to the energy term.

3. A-CONTRARIO SELECTION

When considering the general task of object detection, one realizes that as the target variability increases, as the prior knowledge regarding their appearance becomes more imprecise, it may be more effective to model the background data rather than to model accurately the category of interest. Thus, an a-contrario strategy proposes typically a “naive” model, which characterizes the absence of interesting structures in the data. A particular event, such as a particular type of alignment, a change etc., will be detected as a very unlikely realization under the initial assumption of the naive model. This phenomenon was underlined by the Helmholtz principle in the following way: “*Any structure too regular to appear by chance attracts the attention and becomes a perception*”.

In order to measure the degree of significance of an observation which would seem as ‘too structured’,¹⁴ introduces a data-driven measure directly related to the number of occurrences associated to this observation. If this number is large, the significance of that particular sample is low. The approach allows for grouping a potentially large set of detection parameters (thresholds, etc.) under a single criterion: the expected number of false alarms (NFA).

Under these circumstances, we consider a crack as an ‘unexpected’ configuration of pixels exhibiting a radiometric property which differentiates them from the background model. The difficulty of the detection process arises from the limited amount of information characterizing each pixel in the input data, and also from the fact that the recording conditions (indoor/outdoor variations, embedded cameras etc.) might degrade even further this information.

Typically, one would perform thresholding on the gray level values in order to determine the values of pixels which would be surprisingly low. However, the choice of the threshold is difficult and suffers from robustness; in contrast, NFA boils down to the automatic decision of a threshold. The strategy we propose above assists us in determining the threshold automatically by estimating the significance of the analyzed pixels with respect to the following naive model:

Definition 1 (*naive model H_1*): An image representing the unstructured scene (defect free) is a random field of independent Gaussian variables $\mathcal{N}(\mu, \sigma^2)$, with σ being the standard deviation estimated from the observations of defect free environments.

Considering an observation (acquired image) I , the residual error δ_D^2 between I and the unstructured scene model on a sub-domain $D \in I$ is difficult to interpret in absolute terms. Therefore, we approach the decision in terms of the probability $\mathbb{P}_{H_1}(\delta_D^2)$ of observing δ_D^2 by chance. Consequently, the number of false alarms on that specific sub-domain D is defined as:

$$NFA(D) = \eta(D)\mathbb{P}_{H_1}(\delta_D^2) = \eta(D)\frac{1}{\Gamma(\frac{|D|}{2})} \int_0^{\delta_D^2/2\sigma^2} e^{-t}t^{\frac{|D|}{2}-1}dt \quad (2)$$

where Γ is the Gamma function and $\eta(D)$ is a normalization term (equivalent to the number of tests) that controls the average number of false alarms.

In our work, we use the a-contrario strategy in two stages, namely the detection of seeds and respectively the creation of uncommon paths, which are detailed below.

3.1 The detection of seeds

Assuming that defect pixels are characterized by a significantly lower intensity value in the images, we may rely on this property in order to identify pixels which have a significant chance of belonging to a fissure. The NFA measure will thus take into account in a straightforward manner a residual error based on the radiometric difference between the pixel and the naive model. This first step will provide a set of locations $\mathcal{S} \in I$ corresponding to locations of potential fissure pixels, which we denote as seeds.

Although this stage may seem relatively simple, it has a significant impact on the performance of the following processing steps for the following reason. Given two seeds lying on the same fissure, it is trivial to locate the fissure accurately by searching for a min-cost path between the two seeds. However, the automated detection process is not aware a priori of any seed locations, thus the robustness of this first seed detection stage is essential for the rest of the process. Secondly, given a pair of seeds, we do not know a priori whether the seeds belong to the same fissure, and the min-cost path connecting them might not represent a defect. The following step of the detection deals with this issue.

3.2 The creation of uncommon paths

Rather than selecting seeds and attempting to compute min-cost paths, we adopt an alternative strategy which will select a seed and attempt to build iteratively a fissure originating at the respective seed, based equally on the assumption that along the fissure gray level values will be low. In order to constrain a specific fissure profile for the path, we consider connectivity types as in.² For a path $\mathbf{s} = \{l_1, l_2, \dots, l_n\}$ of pixel locations l_i (l_1 being the seed), with any consecutive pair (l_i, l_j) respecting the considered connectivity constraint, the cost of the path may be formulated as $c(\mathbf{s}) = \sum_{i=1}^n I(l_i)^2$. The use of a thresholding criterion is necessary in order to prune the potential paths which exhibit a high cost, which would require careful tuning. Our proposal is to use a second NFA measure which will integrate an error residual (equally based on the naive model H_1) along the currently considered path. Since the number of paths to track increases exponentially with the path length, we perform pruning whenever the number of paths becomes too high by ordering the paths with respect to the residual and discarding the least promising half of them.

4. EXPERIMENTAL RESULTS

4.1 Datasets and procedure

For our experiments, we considered two fissure datasets. A very recent dataset is publicly available and contains good quality high-resolution images,¹⁵ but currently it does not provide the ground truth related to the defects. A second dataset has been used in¹⁶ and contains lower resolution images, but it also contains ground truth data which allows for quantitative performance analysis. We have also blurred images from the latter dataset using a realistic algorithm for generating sudden motion blur,¹⁷ and used the results as a third evaluation element.

4.2 Standard crack images - the high resolution dataset

We have applied the FFA algorithm as well as the a-contrario based strategies on the high resolution dataset (Figure 1). FFA seems to isolate roughly the well contrasted parts of fissures, but we found that on these images the detection would output a significant amount of false positives, for a large spectrum of detection windows d . In contrast, the seed selection guarantees that cracks will be initialized from highly promising locations; at the same time, false seeding locations are also pruned by the path creation, since this stage of the a-contrario algorithm favors longer paths. The only drawback of this mechanism is that close defects may be connected using an insignificant path in order to obtain a longer result. Still, the interpretation of this result is open to debate as the FFA has also a specific post-processing step performing this task, and even human operators may choose to interconnect close defects.

4.3 Standard crack images - the low resolution dataset

We have also tested our algorithms on the dataset provided by.¹⁶ On these images, we tested the FFA, the a-contrario detection and the MPP algorithm, and all of them perform significantly well. The main remarks regarding the FFA is that the detection window d must be carefully tuned (albeit only once per dataset) and that occasionally spurious clusters of false positives are noticeable in the output. The NFA criterion performs very well as long as the seeding process does not encounter a very difficult context. Finally, the MPP algorithm outputs reliable solutions, and although this solution, at least in the current stage, is not as fine-grained as the NFA-based result, it does have the advantage of highlighting locally the main defect direction, and of encoding in a compact manner the characteristics of the defect. In Figure 2 (the upper row) we present some detection results for the a-contrario strategy; owing to space limitations, we cannot display the detection results for the other two algorithms considered above but they are satisfactory as well.

4.4 Blur degradation

We have applied to the previous dataset a motion blur degradation which simulates random motion trajectories featuring with a lower probability abrupt perturbations, thus mimicking sudden movements specific to servoing or to trajectory micro-perturbations.¹⁷ We have performed the same tests on the resulting dataset and we note that the FFA strategy is extremely sensitive to blur; as soon as the fissure contrast is deteriorated by the blur phenomenon, FFA is not able to identify one of the free-form local paths as statistically different with respect to the other three, and it will output very few or no defect pixels at all. The other strategies cope much better with blur since, among other things such as avoiding hard thresholding, a final exclusion decision for a pixel is never taken based on the pixel local neighborhood alone. The last three rows of Figure 2 illustrate the performance of the a-contrario strategy under blur motion; MPP performs quite well too, while the FFA is unable to cope with this type of data.

4.5 Discussion

The experiments we performed seem to point to the fact that FFA is significantly more sensitive to blur than the alternative solutions we propose. We assume that this is due to the fact that in blur conditions, the degree of coherence between fissure and non-fissure paths grows to the point where they are indistinguishable. This would suggest that either deblurring should be employed as a preprocessing step, or that the coherence and also the subsequent thresholding process should be estimated using a more sensitive approach. Although the other two methods are more robust in the absence of any pre-processing, an advantage of FFA is represented by its relatively low computational complexity.

5. CONCLUSION

We have considered below two different families of fissure detection algorithms, and we proposed some improvements in order to cope better with blur in input images. The results show that on both simulated and real images, the image quality degradation has a significant impact on the detection of the fine crack structures; however, the solutions we propose provide a graceful degradation of performance, and could further be improved by pre-processing steps (i.e. deblurring).

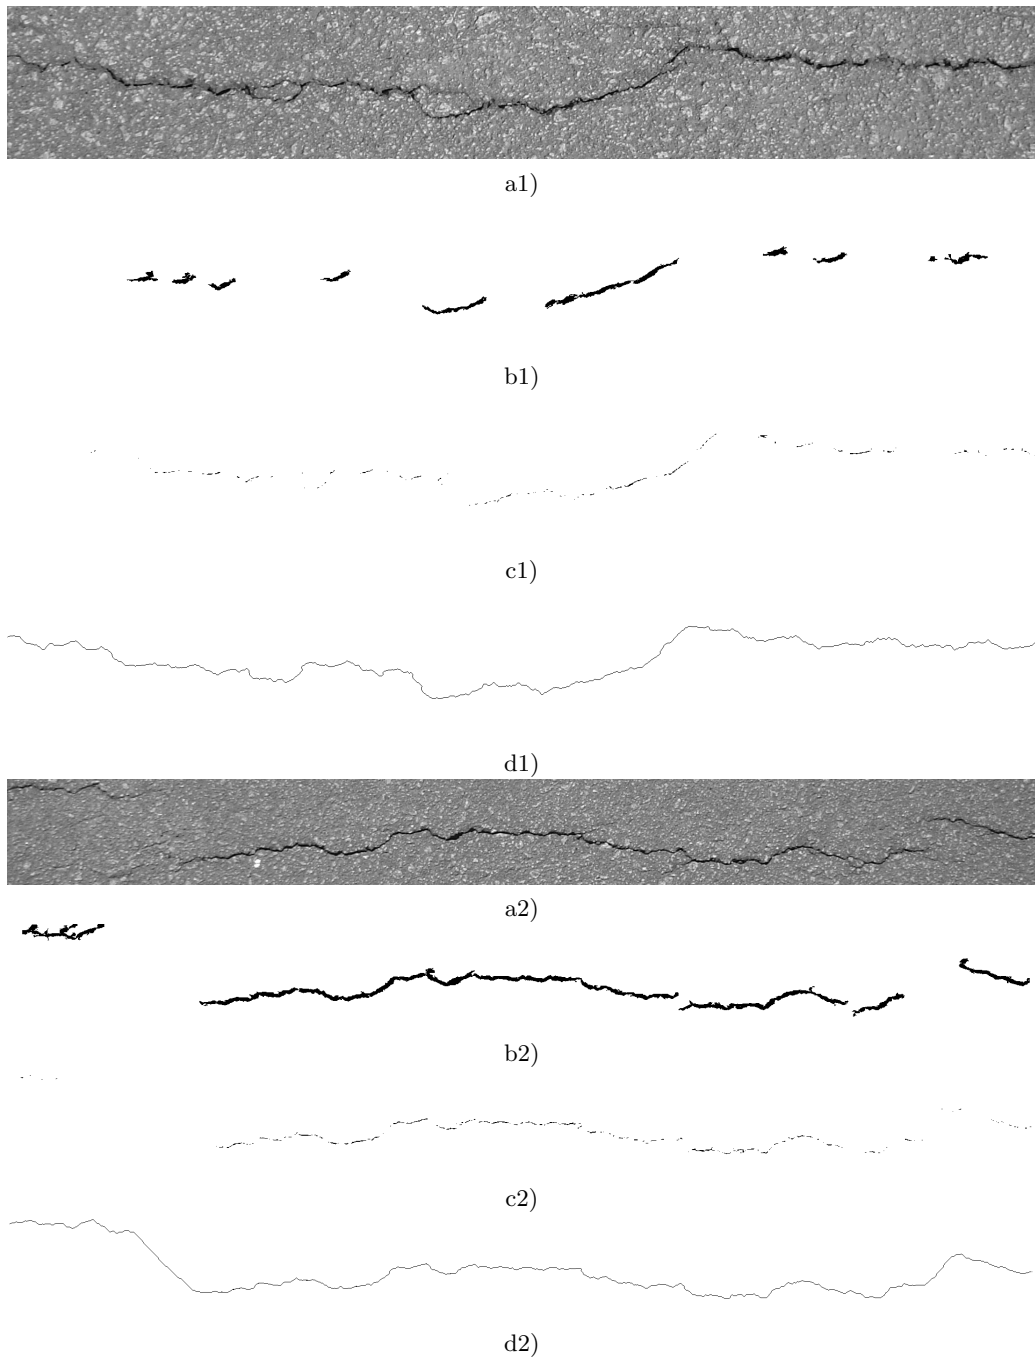


Figure 1. Qualitative comparison on a public dataset.¹⁵ a) highlighted regions of interest (presented for size constraints) b) FFA detection c) seed detection d) uncommon path built based on the seeds. Outside the region of interest, FFA outputs a significant amount of false detections. The a-contrario strategy will reliably identify the defect, although it may choose to fusion smaller paths into a longer one (see d2) as it favors the latter category of defects.




Figure 2. Detection performance and resilience to blur for the dataset.¹⁶ The first row illustrates an a-contrario detection for initial images. The following three rows show the performance for successively stronger motion blurs (equivalent to 2x longer motion exposure each time). The red channel encodes the detected defect; the lower left image indicates a failure.

ACKNOWLEDGMENTS

We would like to thank the authors of¹⁶ for providing us their dataset and ground truth information for testing.

REFERENCES

- [1] Gavilan, M., Balcones, D., Marcos, O., Llorca, D. F., Sotelo, M. A., Parra, I., Ocana, M., Aliseda, P., Yarza, P., and Amirola, A., “Adaptive road crack detection system by pavement classification,” *Sensors* **11**(10), 9628–9657 (2011).
- [2] Avila, M., Begot, S., Duculty, F., and Nguyen, T. S., “2D image based road pavement crack detection by calculating minimal path and dynamic programming,” in [*ICIP*], (2014).
- [3] Chambon, S. and Moliard, J.-M., “Automatic road pavement assessment with image processing: Review and comparison,” *International Journal of Geophysics* **2011** (2011).
- [4] Metni, N. and Hamel, T., “A UAV for bridge inspection: Visual servoing control law with orientation limits,” *Automation in Construction* **17**(1), 3 – 10 (2007).
- [5] Adhikari, R., Moselhi, O., and Bagchi, A., “Image-based retrieval of concrete crack properties for bridge inspection,” *Automation in Construction* **39**(0), 180 – 194 (2014).
- [6] Eschmann, C., Kuo, C., Kuo, C., and Boller, C., “Unmanned aircraft systems for remote building inspection and monitoring,” in [*Proceedings of the sixth European workshop on structural health monitoring*], (2012).
- [7] Varadharajan, S., Jose, S., Sharma, K., Wander, L., and Mertz, C., “Vision for road inspection,” in [*IEEE Winter Conference on Applications of Computer Vision, Steamboat Springs, CO, USA, March 24-26, 2014*], 115–122 (2014).
- [8] Oh, J.-K., Jang, G., Oh, S., Lee, J. H., Yi, B.-J., Moon, Y. S., Lee, J. S., and Choi, Y., “Bridge inspection robot system with machine vision,” *Automation in Construction* **18**(7), 929 – 941 (2009).
- [9] Lacoste, C., Descombes, X., and Zerubia, J., “Point processes for unsupervised line network extraction in remote sensing,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **27**(10), 1568–1579 (2005).
- [10] Perrin, G., Descombes, X., and Zerubia, J., “A marked point process model for tree crown extraction in plantations,” in [*Image Processing, 2005. ICIP 2005. IEEE International Conference on*], **1**, I–661, IEEE (2005).
- [11] Ge, W. and Collins, R. T., “Marked point processes for crowd counting,” in [*Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*], 2913–2920, IEEE (2009).
- [12] Batool, N. and Chellappa, R., “Modeling and detection of wrinkles in aging human faces using marked point processes,” in [*Computer Vision - ECCV 2012. Workshops and Demonstrations - Florence, Italy, October 7-13, 2012, Proceedings, Part II*], 178–188 (2012).
- [13] Jeong, S.-G., Tarabalka, Y., and Zerubia, J., “Marked point process model for facial wrinkle detection,” in [*ICIP*], (2014).
- [14] Desolneux, A., Moisan, L., and Morel, J., “Meaningful alignments,” *International Journal of Computer Vision* **40**(1), 7–23 (2000).
- [15] Oliveira, H. and Correia, P., “Crackit an image processing toolbox for crack detection and characterization,” in [*Image Processing (ICIP), 2014 IEEE International Conference on*], (Sept 2014).
- [16] Fujita, Y. and Hamamoto, Y., “A robust automatic crack detection method from noisy concrete surfaces,” *Machine Vision and Applications* **22**(2), 245–254 (2011).
- [17] Boracchi, G. and Foi, A., “Modeling the performance of image restoration from motion blur,” *Image Processing, IEEE Transactions on* **21**, 3502 –3517 (aug. 2012).