TRANSFORMATION OF ADAPTIVE THRESHOLDS BY SIGNIFICANCE INVARIANCE FOR CHANGE DETECTION

Til Aach

Institute of Imaging and Computer Vision
RWTH Aachen University, D-52056 Aachen
Germany
til.aach@lfb.rwth-aachen.de

Alexandru P. Condurache

Institute for Signal Processing
University of Lübeck
Ratzeburger Allee 160, 23538 Lübeck
Germany
condura@isip.uni-luebeck.de

ABSTRACT

The detection of changes in image sequences often is the first essential step to video analysis, e.g. for the detection, classification and tracking of moving objects. As a binary classification problem, change detection is afflicted by the trade-off between two class error probabilities, viz. the rates of false positives and false negatives. In this contribution, we derive an adaptive two-threshold scheme to improve on this trade-off. The threshold selection for each pixel in the current frame is controlled by the previous detection result for this pixel. Since the test statistics are calculated from samples comprising several pixels within a local sliding window, a transformation of the thresholds from the single-pixel observations to decisions based on larger samples is required. Based on the fact that we can only model the null hypothesis, i.e., absence of motion, realistically, we suggest to transform the threshold under the constraint of a constant false-positive rate, or significance invariance. The resulting detection algorithm is only marginally more complex than a straightforward global thresholding procedure, while providing visibly improved results.

Keywords: motion detection, adaptive thresholding, significance test, significance invariance.

1. INTRODUCTION

Detection of moving objects in image sequences often relies on the detection of temporal grey level changes [1, 2, 3, 4, 5, 6, 7]. The underlying reasoning is that structural changes in the depicted scene, such as motion of the objects to be detected, but also background motion or changes in the viewing angle by egomotion of the camera, induce temporal changes of the 3D world point projected respectively onto any given point in the image plane, which in turn result in a noticeable grey level change. However, such grey level variations may also originate from other sources, particularly from illumination changes and camera noise. Assuming a static camera or that global background motion has been compensated, and a temporally only slowly varying scene illumination, the observed grey level changes may then be attributed to object motion or camera noise [7, 8]. In addition, if the condition of an only slowly varying scene illumination is not met, the corresponding unwanted temporal grey level variations can be filtered out to a certain degree by approaches such as homomorphic filtering [9, 10, 11] or motion analysis using an intrinsically less illumination-sensitive test statistic based on a total least squares approach [12, 13, 14].

On the other hand, object motion does not always generate noticeable grey level variations. Motion of, e.g., spatially homogeneous regions with low gradients may result in only low grey level variations (cf. [15] for analysis of optic flow). Object motion is therefore neither a sufficient nor a necessary condition for temporal grey level changes to occur. In consequence, we regard the reverse inference from grey level variations to object motion as a two-class decision problem. The detection of object motion via temporal grey level changes is therefore afflicted by the trade-off between the two class error probabilities, viz. the rates of false positives and false negatives. Moreover, since motion detection as a binary segmentation problem seeks to infer the underlying structure from observed noisy image data, we added it [16, 17] to the list of ill-posed problems in computer vision [18, 19, 20]. When solving such problems, the space of possible solutions is constrained by appropriate regularization [21, 22, 23, 24], which may also be expressed in a statistical manner within Bayesian approaches [25, 7]. In the context of motion detection, this translates into adaptive algorithms, which considerably improve the above trade-off. Successful adaptive techniques often rely on an appropriate modelling of the prior knowledge within a Bayesian framework, e.g. in terms of Gibbs/Markov random fields (GMRF) [26, 27]. Spatial GMRFs are suitable to express the prior expectation of compactly shaped moving objects, and consequently sup-

press the emergence of spurious, noise-like detection results [16].

For some applications, however, such models may overconstrain the detection procedure. An example is the analysis of a bolus of contrast agent flowing through the blood vessels in digital X-ray angiography. The fine, filigran-like and anisotropic nature of the smaller vessels may not be appropriately captured by a spatial GMRF. This kind of overconstraining is vividly illustrated by the almost isotropically compact nature of label fields drawn from GMRFs by a Gibbs sampler [26, 28].

In this contribution, we therefore restrict ourselves to modelling the temporal behaviour of the states of a pixel. An injection of a contrast bolus, for instance, extends over a certain finite time interval. While flowing through the vessels, the contrast bolus is further diluted. This dispersion is often described by a convolution of a vessel impulse response with the injection profile. It is therefore reasonable to assume that, once the moving contrast agent is detected at a certain pixel, this pixel remains in a moving state for a certain time interval [29]. This translates into a certain degree of similarity between subsequent change masks, which is in the following exploited.

2. CONTEXT-ADAPTIVE CHANGE DETECTION

2.1. Statistical Modelling

To assess temporal changes, we calculate the pixel-wise difference image $D = \{d(k)\}$ between the two grey level images $G_j = \{g_j(k)\}$, $j = 1, 2$, to be compared; these may be two successive video frames, or one reference and one current image. In the case of two successive frames, this differencing corresponds to a temporal highpass which eliminates slow illumination changes from further analysis. For each pixel k, the difference is computed as $d(k) = g_2(k) - g_1(k)$. We assume the camera noise as additive, zero-mean, Gaussian distributed and spatiotemporally independent. Note that, when processing X-ray images, the noise is actually Poisson distributed [30], what, for reasonable quantum counts, may be approximated by a Gaussian pdf with signal-dependent variance [31]. In the current motion mask $Q = \{q(k)\}$, we seek to assign to each pixel k a binary label $q(k) \in \{u, c\}$, with u denoting *unchanged* (null hypothesis H_0), and c denoting *changed* (hypothesis H_1). Since we process image sequences, the previous motion mask $R = \{r(k\}$, $r(k) \in \{u, c\}$ shall also be available. For the decision at pixel i, we observe a sample $\mathbf{d}_i = \{d(k)|k \in w_i\}$ of differences within a small sliding window w_i centred at i. Use of an observation window makes the decision more robust, at the cost, though, of a slight loss of spatial resolution [8].

In addition, within w_i we observe a sample $\mathbf{r}_i = \{r(k)|k \in w_i\}$ of previous labels. Resorting to the likelihood ratio as a most powerful test statistic [32], a possible decision approach is

$$\frac{p(\mathbf{d}_i, \mathbf{r}_i|H_1)}{p(\mathbf{d}_i, \mathbf{r}_i|H_0)} \underset{u}{\overset{c}{\gtrless}} t \tag{1}$$

with this notation indicating that $q(i) = c$ if the left quantity exceeds the right-hand side, and $q(i) = u$ if the left-hand side is less than the right-hand side. The fixed threshold t depends on the chosen criterion, such as maximum a posteriori estimation, and the global prior probabilities of the two hypotheses.

It is reasonable to assume that the statistical properties of the observed grey level difference vector $\mathbf{d}_i$ depend only on the underlying hypothesis (or "state") H_0 or H_1. In other words, given the state, the old label constellation $\mathbf{r}_i$ does hence not influence the probability density function (pdf) of $\mathbf{d}_i$, i.e. $p(\mathbf{d}_i|H_j, \mathbf{r}_i) = p(\mathbf{d}_i|H_j)$, $j = 0, 1$. The observations $\mathbf{d}_i$ and $\mathbf{r}_i$ are thus *conditionally* independent:

$$p(\mathbf{d}_i, \mathbf{r}_i|H_j) = p(\mathbf{d}_i|H_j, \mathbf{r}_i)p(\mathbf{r}_i|H_j) = p(\mathbf{d}_i|H_j)p(\mathbf{r}_i|H_j) \ . \tag{2}$$

Decision (1) can then be written as

$$l(\mathbf{d}_i) = \frac{p(\mathbf{d}_i|H_1)}{p(\mathbf{d}_i|H_0)} \underset{u}{\overset{c}{\gtrless}} t \cdot \frac{p(\mathbf{r}_i|H_0)}{p(\mathbf{r}_i|H_1)} = \hat{t}(\mathbf{r}_i) \ . \tag{3}$$

Here, $\hat{t}(\mathbf{r}_i)$ is an adaptive threshold depending on $\mathbf{r}_i$. Let the variance of the differences $d(k)$ under the null hypothesis be σ^2, which be known from calibration measurements (a simple but in many cases sufficient model is that σ^2 is twice the camera noise variance). Under the assumption that given H_1, the observed grey level differences obey another zero-mean Gaussian distribution with considerably larger variance σ_c^2 [16], we obtain from (3)

$$\left(\frac{\sigma}{\sigma_c}\right)^N \exp\left\{\frac{1}{2}(1 - \frac{\sigma^2}{\sigma_c^2})\Delta_i\right\} \underset{u}{\overset{c}{\gtrless}} \hat{t}(\mathbf{r}_i) \ . \tag{4}$$

Here, $\Delta_i = 1/\sigma^2 \sum_{k \in w_i} d^2(k)$, and N is the number of pixels in w_i (practically, $N = 25...49$). With $\sigma_c^2 >> \sigma^2$ (measurements showed that typically, $\sigma_c^2 \approx 100\sigma^2$), and taking the logarithm, (4) becomes

$$\Delta_i \underset{u}{\overset{c}{\gtrless}} 2 \log\left[\left(\frac{\sigma_c}{\sigma}\right)^N t\right] + 2 \log \frac{p(\mathbf{r}_i|H_0)}{p(\mathbf{r}_i|H_1)} = \tilde{t}_N(\mathbf{r}_i) \ . \tag{5}$$

The decision threshold adapts to the label constellation $\mathbf{r}_i$ in the previous image. The first part of the sum corresponds to a fixed threshold t_s, which, due to σ_c^2 being unknown, cannot be directly evaluated. Given H_0, $p(\Delta_i|H_0)$ is a χ^2-distribution with N degrees of freedom. Setting an acceptable false alarm rate or significance α, we determine t_s from $\text{Prob}(\Delta_i > t_s|H_0) = \alpha$. For the adaptive part of the

threshold, however, there are 2^N different label constellations possible in w_i, which, even for small windows, result in impracticably many parameters. In the next section, we show how the number of adaptive thresholds can be reduced to two, which depend only on the previous detection result $r(i)$.

2.2. Threshold Transform by Significance Invariance

We start from a window w_i with just one pixel, i.e., $N = 1$. The observed vectors $\mathbf{d}_i$ and $\mathbf{r}_i$ reduce to a scalar difference $d(i)$ and a single previous label $r(i)$. The fixed threshold part t_s is then determined from a χ^2-distribution with only one degree of freedom for a given significance α. From the four "reverse" transition probabilities $p(r(i) = u|H_1)$, $p(r(i) = c|H_1)$, $p(r(i) = u|H_0)$, $p(r(i) = c|H_0)$, only one per hypothesis needs to be estimated or set from experience, the other two are easily determined from the total probability. Evidently, because of temporal coherence, $p(r(i) = u|H_0) > p(r(i) = c|H_0)$, and $p(r(i) = c|H_1) > p(r(i) = u|H_1)$. For $r(i) = u$ and $r(i) = c$, the threshold pair is from (5)

$$\tilde{t}_1(u) = t_s + 2\log\frac{p(r(i) = u|H_0)}{p(r(i) = u|H_1)} > t_s \qquad (6)$$

and

$$\tilde{t}_1(c) = t_s + 2\log\frac{p(r(i) = c|H_0)}{p(r(i) = c|H_1)} < t_s \qquad (7)$$

In this notation, the subscript "1" of the threshold variables $\tilde{t}_1(r(i))$ reflects the validity of these thresholds for a window w_i of size 1×1.

To transform this threshold pair towards being used with larger window sizes N, we determine the significance for each threshold by

$$\text{Prob}(d^2(i) > \tilde{t}_1(u)|H_0) = \alpha_u \qquad (8)$$

and

$$\text{Prob}(d^2(i) > \tilde{t}_1(c)|H_0) = \alpha_c \qquad (9)$$

Again, the probabilities are evaluated based on a χ^2-pdf with one degree of freedom. Requiring these significances to be valid for other window sizes N as well (*significance invariance*), the threshold pair $\tilde{t}_N(u)$, $\tilde{t}_N(c)$ can be obtained by reversing (8) and (9), now based on a χ^2-pdf with N degrees of freedom. Table 1 illustrates this threshold transformation for $N = 25$, $\alpha = 5 \cdot 10^{-4}$, $p(u|H_0) = 0.99$, and $p(c|H_1) = 0.98$.

2.3. The Two-Threshold Algorithm

The above strategy can be implemented as a two-threshold algorithm, which is only marginally more complex than a global, non-adaptive algorithm: the values of the threshold pair are determined off-line in the beginning, and stored in a LUT. Before processing pixel i, we only check its previous label $r(i)$ to extract the corresponding threshold from the LUT. Since checking on $r(i)$ is done before actually deciding on the current label $q(i)$, $r(i)$ can be overwritten by $q(i)$, additional storage for the previous change mask R is therefore not needed. Also, the algorithm can be fully parallelized.

α	t_s	$\tilde{t}_1(u)$	$\tilde{t}_1(c)$
$5 \cdot 10^{-4}$	11.9	19.6	2.7
α_u	α_c	$\tilde{t}_{25}(u)$	$\tilde{t}_{25}(c)$
10^{-5}	0.1	67.5	34.3

Table 1. Threshold transform based on significance invariance.

2.4. Initialization of the Two-Threshold Algorithm

Obviously, when computing a motion mask, the threshold selection is only applicable if a previous motion mask is indeed available. For the first motion mask of an image sequence, though, a predecessor mask does not exist. The initial mask can therefore only be determined by a non-adaptive algorithm with a global threshold, which is reasonably taken from the set $\{\tilde{t}_N(u), \tilde{t}_N(c)\}$. The probably best choice appears to be the lowest of these, viz. $\tilde{t}_N(c)$. This results in a high number of false positives in the first mask. These vanish, however, quickly, since they are caused by outliers of the camera noise, the statistical bindings of which extend over only a short temporal interval determined by the length of the differencing operation. Alternatively, initialization by the high threshold $\tilde{t}_N(u)$ would result in a high false negative rate (i.e. "holes" in moving objects). Since these false negatives are caused by spatially homogeneous regions of moving objects, their correlation length is much longer. The rate of decay is therefore slower.

3. RESULTS

Fig. 1 compares two nonadaptive thresholding results with the dual-threshold method. Fig. 1 b) shows a nonadaptive thresholding using the higher threshold $\tilde{t}_{25}(u)$ in the entire image, resulting in a low false positive rate, while c) shows the same using the lower threshold $\tilde{t}_{25}(c)$, resulting in less false negatives, but at the cost of a higher false positive rate. The dual-thresholding result in d) combines the low false positive rate in b) with the low false negative rate in c). Despite the lack of explicit modelling of spatial compactness, the detected objects appear smoothly shaped. A similar observation holds for the digital angiogram showing a contrast bolus flowing through vessels in Fig. 2. While in Fig. 1, the comparison was carried out between two suc-

cessive frames, in Fig. 2 the algorithm was always applied to the current frame and the first frame, which served as a reference without contrast agent showing up. Fig. 3 shows a moving car recorded by a static camera, and the detection results. Here, the square-shaped detection errors result from outliers caused by digitization errors of the originally analog sequence, which are amplified by the squaring operation. As in Fig. 1, the algorithm was applied to successive frames. In all examples, the window size was $N = 25$; all other parameters were as in table 1. The variances σ^2 were estimated prior to the detection procedure from calibration measurements. Fig. 4 shows the adaptation behaviour of the algorithm after initialization with the lower threshold $\tilde{t}_{25}(c)$. Evidently, the high false-positive rate in the first mask drops swiftly over the next few frames.

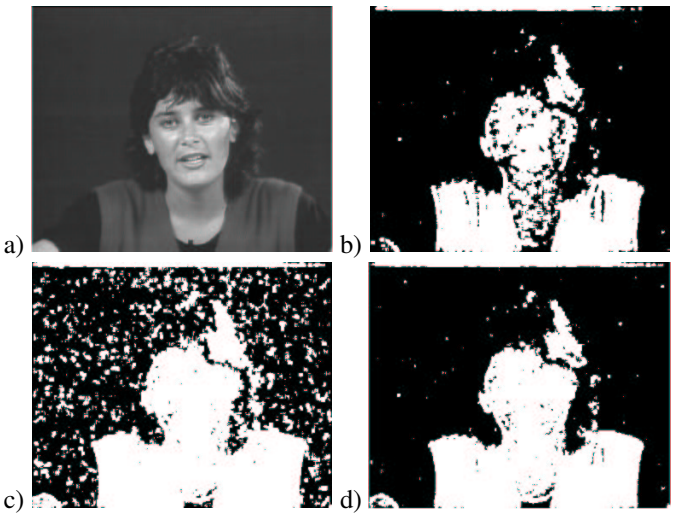

Fig. 1. a) Frame 20 of a video sequence, b) non-adaptive result, $\alpha = \alpha_u = 10^{-5}$, c) non-adaptive result, $\alpha = \alpha_c = 0.1$, d) dual-threshold result.

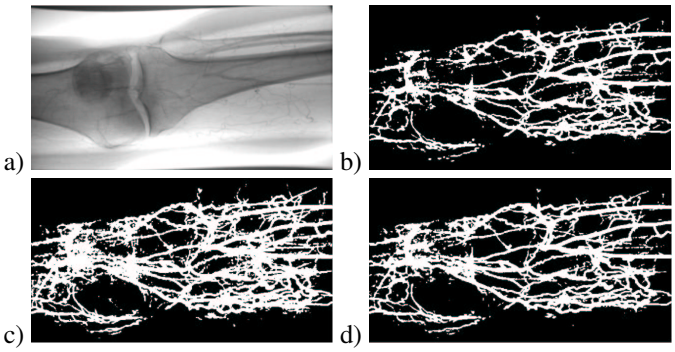

Fig. 2. a) Frame 11 of a digital angiography, b) non-adaptive result, $\alpha = \alpha_u = 10^{-5}$, c) non-adaptive result, $\alpha = \alpha_c = 0.1$, d) dual-threshold result.

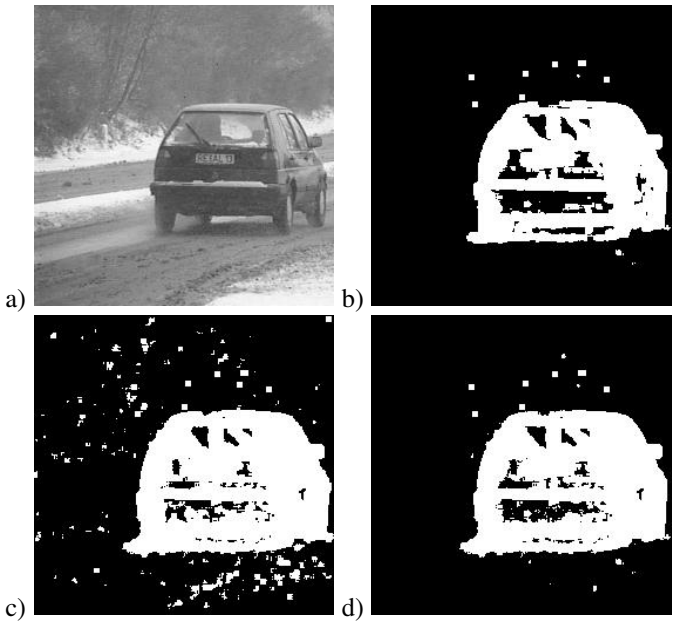

Fig. 3. a) Frame 10 of a traffic scene, b) non-adaptive result, $\alpha = \alpha_u = 10^{-5}$, c) non-adaptive result, $\alpha = \alpha_c = 0.1$, d) dual-threshold result.

4. DISCUSSION

We have described an adaptive dual-threshold motion detection algorithm. The algorithm is non-iterative, and only marginally more complex than a non-adaptive algorithm, but provides a visibly improved trade-off between the class error probabilities. The algorithm can be fully implemented in parallel, and, once the (camera) noise variance σ^2 for the null hypothesis H_0 is estimated, needs only four parameters, viz. the window size N, the initial significance level α, and two transition probabilities. Decision thresholds are calculated based on error probabilities and transition probabilities. For the results shown in this paper, the transition probabilities were determined experimentally. Essentially, the transition probabilities determine how different the two thresholds are. Thus, though their values are not quite uncritical, our results show that they do not depend strongly on the type of image sequences they are applied to. One next step would be to compare the values used in the experiments with transition measurements from sequences of motion masks.

When deriving the decision rule (5) from the general expression in (3), we have assumed independent Gaussian distributions in the difference images. Alternative models taking into account correlations [33] or other marginal distributions, such as a Laplacian pdf to describe frame differences, have been put forward. Our earlier experiments show that exploiting correlations in the test statistics for motion detection has only a minor influence, if at all, on the result [7,

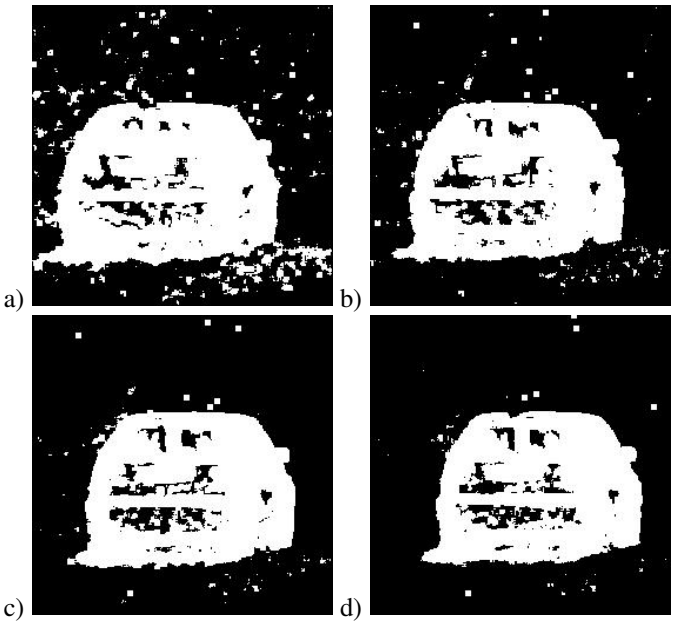

Fig. 4. Adaption of the algorithm after initialization with the lower threshold. a) motion mask 1, b) motion mask 2, c) motion mask 3, d) motion mask 4.

Sec. 5.1], [16, App. A]. The Laplacian model can be used in our framework as well. The major differences are that the square sum Δ_i of differences in (4) is replaced by the sum of absolute differences, and that the number of degrees of freedom of the χ^2-pdf used to relate thresholds and significance values needs to be doubled. Details can be found in [7, 8, 16]. In practice, decision rules based on a Gaussian and a Laplacian model perform almost equivalently. A similar observation was made for image restoration in [34]. The only noticeable difference between these is that the Laplacian model is less sensitive to outliers, since their values are not squared as in the Gaussian model. Furthermore, note that we need distributions *conditioned* on the hypotheses of motion (H_1) and no motion (H_0), for which the Gaussian model appears quite appropriate. The Laplacian model with its narrow peak around zero and longer tails may be viewed as an unconditional mixture model of the two hypotheses.

While the algorithm in its current form is sensitive to quick illumination changes, it can also be used within the illumination-insensitive frameworks mentioned in the introduction, e.g., [10, 12].

The decision procedure in this paper essentially relies on the observation of only a single previous label $r(i)$. Our approach in (3), though, allows to consider a whole set of previous labels $\mathbf{r}_i$. A promising direction of future research is to seek to exploit this in practice. One way to do this would be to make the detection of, say, line-like vessels filled by a contrast agent dependent on previously detected line-like structures captured within $\mathbf{r}_i$, by specifying the transition probabilities $p(\mathbf{r}_i|H_j)$ appropriately.

5. REFERENCES

[1] Y. Z. Hsu, H.-H. Nagel, and G. Rekers, "New likelihood test methods for change detection in image sequences," *CVGIP*26, pp. 73–106, 1984.

[2] A. Mitiche and P. Bouthemy, "Computation and analysis of image motion: A synopsis of current problems and methods," *IJCV*19(1) pp. 29–55, 1996.

[3] G. Tziritas and C. Labit, *Motion Analysis for Image Sequence Coding*, Elsevier, Amsterdam, 1994.

[4] S.-Z. Liu, C.-W. Fu, and S. Chang, "Statistical change detection with moments under time-varying illumination," *IEEE T IP* 7(9), pp. 1258–1268, 1998.

[5] P. Doubek, T. Svoboda, and L. van Gool, "Monkeys - a software architecture for viroom - low cost multicamera system," in *Proceedings 3rd International Conference on Computer Vision Systems*, J. L. Crowley, J. H. Piater, M. Vincze, and L. Paletta, Eds., LNCS, Vol. 2626 2003, pp. 386–395, Springer Verlag.

[6] R. Mech and M. Wollborn, "A noise robust method for segmentation of moving objects in video sequences," in *Proceedings International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Munich, Germany, April 1997, pp. 2657–2660, IEEE.

[7] T. Aach, *Bayes-Methoden zur Bildsegmentierung, Änderungsdetektion und Verschiebungsvektorschätzung*, Fortschrittberichte VDI Reihe 10, Nr. 261, VDI Verlag, Düsseldorf, 1993.

[8] T. Aach, A. Kaup, and R. Mester, "Statistical model-based change detection in moving video," *Signal Processing*, vol. 31, no. 2, pp. 165–180, 1993.

[9] A. V. Oppenheim, R. W. Schafer, and T. G. Stockham Jr., "Nonlinear filtering of multiplied and convolved signals," *Proceedings of the IEEE*, vol. 56, no. 8, pp. 1264–1291, 1968.

[10] D. Toth, T. Aach, and V. Metzler, "Bayesian spatio-temporal motion detection under varying illumination," in *European Signal Processing Conference*, M. Gabbouj and P. Kuosmanen, Eds., Tampere, Finland, September 3-8 2000, pp. 2081–2084, EURASIP.

[11] D. Toth, T. Aach, and V. Metzler, "Illumination-invariant change detection," in *4th IEEE Southwest Symposium on Image Analysis and Interpretation*, Austin, TX, April 2 - 4 2000, pp. 3–7, IEEE Computer Society.

[12] T. Aach, L. Dümbgen, and R. Mester, "Bayesian illumination invariant change detection using a total least squares test statistic," in *Actes/Proceedings 18e Colloque GRETSI sur le Traitement du Signal et des Images*, Toulouse, France, Septembre 10–13 2001, pp. 587–590.

[13] T. Aach, L. Dümbgen, R. Mester, and D. Toth, "Bayesian illumination-invariant motion detection," in *Proceedings IEEE International Conference on Image Processing (ICIP), Vol. III*, ISBN 0-7803-6725-1, Thessaloniki, October 7–10 2001, pp. 640–643, IEEE.

[14] T. Aach, D. Toth, and R. Mester, "Motion estimation in varying illumination using a total least squares distance measure," in *Picture Coding Symposium*, Saint Malo, France, April 23–25 2003, pp. 145–148.

[15] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

[16] T. Aach and A. Kaup, "Bayesian algorithms for change detection in image sequences using Markov random fields," *Signal Processing: Image Communication*, vol. 7, no. 2, pp. 147–160, 1995.

[17] T. Aach, A. Kaup, and R. Mester, "Change detection in image sequences using Gibbs random fields," in *Proceedings International Workshop on Intelligent Signal Processing and Communication Systems*, Sendai, Japan, October 1993, pp. 56–61, IEEE.

[18] M. A. Bertero, T. Poggio, and V. Torre, "Ill-posed problems in early vision," *Proceedings of the IEEE*, vol. 76, no. 8, pp. 869–889, 1988.

[19] T. Poggio, V. Torre, and C. Koch, "Computational vision and regularization theory," *Nature*, vol. 317, pp. 314–319, September 1985.

[20] T. Poggio, "Early vision: From computational structure to algorithms and parallel hardware," *Computer Vision, Graphics, and Image Processing*, vol. 31, pp. 139–155, 1985.

[21] A.N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*, Winston & Sons, Washington,D.C., 1977.

[22] A. N. Tikhonov and A. V. Goncharsky (Eds.), *Ill-Posed Problems in the Natural Sciences*, Mir, Moscow, 1987.

[23] N. B. Karayiannis and A. N. Venetsanopoulos, "Regularization theory in discrete image restoration," in *Proceedings Visual Communications and Image Processing 88*. 1988, vol. 1001, pp. 25–36, SPIE.

[24] V. Torre and T. Poggio, "On edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 2, pp. 147–163, 1986.

[25] R. M. Haralick, "Decision making in context," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 4, pp. 417–429, 1983.

[26] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, no. 6, pp. 721–741, 1984.

[27] J. Besag, "On the statistical analysis of dirty pictures," *Journal Royal Statistical Society B*, vol. 48, no. 3, pp. 259–302, 1986.

[28] G. R. Cross and A. K. Jain, "Markov random field texture models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 1, pp. 25–39, 1983.

[29] T. Aach, A. Condurache, K. Eck, and J. Bredno, "Statistical-model based identification of complete vessel-tree frames in coronary angiograms," in *Electronic Imaging 2004: Computational Imaging II*, C. A. Bouman and E. L. Miller, Eds., San Jose, USA, January 18–22 2004, pp. 283–294, SPIE Vol. 5299.

[30] T. Aach, U. Schiebel, and G. Spekowius, "Digital image acquisition and processing in medical x-ray imaging," *Journal of Electronic Imaging*, vol. 8, no. Special Section on Biomedical Image Representation, pp. 7–22, 1999.

[31] T. Aach and D. Kunz, "Bayesian motion estimation for temporally recursive noise reduction in x-ray fluoroscopy," *Philips Journal of Research*, vol. 51, no. 2, pp. 231–251, 1998.

[32] C. W. Therrien, *Decision, Estimation, and Classification*, John Wiley, New York, 1989.

[33] D. J. Connor and J. O. Limb, "Properties of frame-difference signals generated by moving images," *IEEE Transactions on Communications*, vol. 22, no. 10, pp. 1564–1575, 1974.

[34] H. Derin, "Comments on "Restoration of noisy images modeled by Markov random fields" (with reply)," *IEEE Transactions on Circuits and Systems*, vol. 38, no. 5, pp. 566–567, 1991.