

A POINT-EVENT DETECTION ALGORITHM FOR THE ANALYSIS OF CONTRAST BOLUS IN FLUOROSCOPIC IMAGES OF THE CORONARY ARTERIES

Alexandru Condurache and Alfred Mertins

Institute for Signal Processing, University of Luebeck
 Ratzeburger Allee 160, D-23538, Luebeck, Germany
 phone: + 49(0)451 5005801, fax: + 49(0)451 5005802, email: condura@isip.uni-luebeck.de
 web: www.isip.uni-luebeck.de

ABSTRACT

The paradigm of event detection is general enough to be used in a multitude of applications from various domains. In this contribution we describe an improved method to mark the appearance of blood mixed with contrast agent in a sequence of fluoroscopic images of the coronary arteries. This is needed for various computer-vision based technologies aimed at supporting the physician during a Percutaneous Transluminal Coronary Angioplasty (PTCA). PTCA is a surgical intervention conducted for the purpose of reopening blocked coronary arteries. We show how to extract a feature describing the amount of contrast agent present in each fluoroscopic image and how to establish a threshold over this feature, to separate the event of contrast-agent appearing from the normal case, when no contrast agent is present. For this purpose we estimate the likelihood of feature-values given the normal case and decide to mark the event for images whose feature has a very small likelihood. We test our algorithm on a number of sequences acquired in clinical routine.

1. INTRODUCTION

The problem of event detection can be divided into three groups [4]: point events, context event and collective events. To detect an event one analyzes single observations or groups of observations relative to a certain process. In the case of point events, the observations we try to select are very different from the rest of the data and appear very rarely. Under this definition looking for point events is in some cases similar to detecting outliers. For context events, one no longer analyzes single observations, but groups of successive observations and marks an event if a single observation is strikingly different from the others in its group, even if the observation itself is no outlier with respect to entire data. In the case of collective events, we again analyze groups of observations, and mark an event if a group exhibits certain characteristics, usually other than those specific to a context event. In this contribution we describe a point-event detection algorithm.

Coronary arteries disease occurs when the vessels supplying oxygenated blood to the heart muscle narrow as a consequence of plaque buildup. Treatment is facilitated by Percutaneous transluminal coronary angioplasty. During such an intervention, the narrowing is eliminated usually by inflating a balloon at that position in the vessel. The balloon is brought in place with the help of a fluoroscopic imaging system. The imaging system functions in imaging sessions, which usually differ from one another by the position of the imaging device. To make the vessels visible under X-ray for a while, a bolus of radio-opaque contrast agent is injected through a catheter positioned into the vessels. Algorithms

have been developed [1], [6] to use the contrasted images, also called coronary angiograms, to build a dynamic vessel roadmap such that the physician still sees the vessels even after the contrast agent has washed out, and can thus navigate better. In this context we need to find out when contrast agent first appears in the fluoroscopic images during an imaging session. Considering that previous to that moment only images without vessels have been recorded, this is a point-event detection problem.

In [5] we have described a method to detect the first image of a contrast agent injection by means of a significance test. In this paper we show how to improve this method by estimating the likelihood of the null hypothesis by means of nonparametric probability density function (pdf) estimation, rather than parametric under the Gaussian assumption. We also provide a detailed description of the feature-extraction process, test our algorithm on a larger set of sequences acquired in clinical routine and quantitatively evaluate its performance.

2. METHODS

For our vessel-event detection algorithm, we start by first computing a feature related to the proportion of image area occupied by vessels. We then establish a threshold over this feature and mark the event at the moment when we record a fluoroscopic image whose feature is above the threshold.

2.1 Vessel enhancement and feature extraction

To directly estimate the vessel area in fluoroscopic images we need vessel segmentation, which is difficult due to the small signal-to-noise ratio in most of the analyzed images. Thus, instead of computing the vessel area, we introduce a vessel-area-related feature. The vessel feature is computed from a vessel map whose histogram can be assumed to consist of two distinct distributions, one from background and one from potentially occurring vessels. The vessel map is the result of several enhancement steps aimed at improving separability between background and vessels. As we seek to detect the presence of contrast agent from the vessel-map histograms rather than accurate segmentation or enhancement geared toward the human observer, issues like border accuracy and preservation of a certain "harmony" in the vessel map are of less concern [2], [9].

The first step towards enhancing vessel contrast exploits the property that, being filled by contrast agent, vessels are locally darker than their immediate surroundings [2]. We select such structures and equalize the varying background by a Bothat transform [7], thus reducing its standard deviation.




Figure 1: Coronary angiogram (a) and result of morphological processing (b).

The Bothat transform preserves though the grey level difference between background and vessels. This transform applied to image I is defined as:

$$B(I) = (I \bullet \mathcal{B}) - I \quad (1)$$

the difference between the closing of the original image by the structuring element \mathcal{B} and the image itself. \mathcal{B} has a circular shape slightly larger than the largest vessel. In the end we invert the image, such that vessels are again darker than background. A processing example is shown in Figure 1. The purpose of the next enhancement steps is to transform the additional absorption of contrasted vessels such that its mean is considerably larger than the background mean.

An additional clue to vessels is their mostly strong and jerky motion, which is caused by the beating heart [2] and shifts them quickly over a distance of several vessel diameters. When the vessel moves to a new position, absorption at this position will increase due to the contrast agent within the vessel. Therefore, when calculating the pixel-wise difference image between any given current image I_C and its predecessor I_P , pixels with the new vessel positions will tend to exhibit negative differences, while the vessel positions in the previous image will tend to show positive differences. This effect is even more pronounced in rescaled Bothat-transformed angiograms, since the Bothat transform reduces background structures that would otherwise cause clutter in the difference images. In a second motion-based vessel-enhancement step, we therefore clip positive values in the difference images $D(I_C, I_P) = B(I_C) - B(I_P)$ to zero, and add this result to the current transformed image:

$$M(I) = B(I_C) + D_{\leq 0}(I_C, I_P) \quad (2)$$

Consequently, the moved vessel-tree regions will become darker than before. Local minima of small extent present in the clipped difference image $D_{\leq 0}(I_C, I_P)$, which are unlikely to be caused by moving vessels, are removed by a morphological closing with a 3×3 structuring element. A block diagram of the processing chain and a processing result are shown in Figure 2.

The Bothat transform selects all dark background structures comparable in size with the vessels including some which are not genuine vessels but vessel-like artifacts (i.e. vessel noise). As the motion-based enhancement step is effective only for moved vessels, such noise still affects motionless vessels. These noise structures have a small surface, are dark but somewhat less contrasted than comparable vessels, have mostly diffuse edges, and their shape is mostly




Figure 2: Block diagram of spatiotemporal filtering (a) and filtration result (b).




Figure 3: Block diagram of derivative processing (a) and final vessel map (b).

patch-like although some of them can be elongated and thin. The patch-like vessel noise is due to the soft-tissue that the X-rays encounter on their way through the body and is thus ubiquitous. The surface it covers varies from image to image as a consequence of the motion induced by the heart and by the respiration. The elongated vessel noise is mainly due to bone-tissue, it does not cover such a large area, it is localized, and its surface is approximately constant. By contrast to the patch-like noise, the influence of the elongated noise on our algorithms can be ignored.

In the third and last vessel enhancement step, we attenuate the patch-like noise. For this purpose, starting from the result of the Bothat transform, we compute first the gradient-norm and add it to the second derivative [10, 13] using empirically determined weights:

$$E(I) = 0.55 \|\nabla B(I)\| + 0.45 \Delta B(I) \quad (3)$$

Then, we multiply this result pixelwise (\odot) with the result of the motion-based enhancement to finally obtain the vessel map:

$$V(I) = E(I) \odot M(I) \quad (4)$$

$E(I)$ is equivalent to a logical OR operation between the results of the gradient-norm, which responds to contrasted structures with sharp edges but does not respond to vessel-like ridge-profiles, and the second derivative, which does. Each operation responds to different vessel characteristics, by combining them we increase the robustness of our method. $V(I)$ is the equivalent of a logical AND operation between the result of the derivatives-based enhancement and that of the motion-based enhancement, and is justified by the fact that each procedure is based on different vessel properties, which are thus brought together. To correctly apply the AND operation, the result of the motion-based enhancement is rescaled before multiplication so that vessels are brighter than background. A block diagram of the processing chain and a vessel-map result are shown in Figure 3.

To detect interventional images where contrast agent is also present, we desire a feature that reaches its peak once

the slightest trace of contrast agent appears. As the area still influenced by vessel noise in an angiogram is about one percent, we choose the 98 percentile as feature. Then, for images where contrast agent is present, where more than two percents of all pixels in the analyzed image have high intensities, the feature will exhibit higher values than for images that show no vessels.

2.2 A nonparametric maximum-likelihood method for vessel-event detection

The 98-percentile feature ideally builds a separable feature space, as an image either shows some vessel structures or shows only background. The 98 percentile is linked to the vessel area. When contrast agent is injected into the vessels, the 98 percentile increases rapidly as it is computed from gray levels that are characteristic for pixels belonging to the large, well-contrasted vessels.

Our problem can be also formulated like a binary classification problem, where we try to separate vessel images from background images. To find the optimal percentile threshold T that separates these two classes we need to know something about them. This knowledge is usually extracted from a sample of our feature space that covers both classes. However, it is very difficult to obtain such a sample here, because:

1. the strong intervention variability requires each case to be treated independently, and thus it is pointless to gather data from several patients, design a classifier and use it on other patients.
2. the analysis of fluoroscopic on-line images clearly requires a strictly causal processing, and thus one cannot wait until the end of the sequence to analyze it.

All images recorded until the vessels appear show only the background, therefore we formulate our problem like an event-detection problem and consider that an image showing no vessels represents the normal case, and an image with vessels represents the event. Usually there are 10 to 15 seconds from the moment a recording begins until vessels appear. By a frame rate of at least 12 fps, this gives us quite a large sample of background images that in turn makes possible the computation of the likelihood of the vessel feature given the background class. We can thus model the pdf of the “normal” case and mark the event as soon an image is recorded whose feature has a small likelihood.

This likelihood approach to point-event detection is similar to a significance test [11] whose null hypothesis H_0 is that the investigated image shows no vessels [5].

$p(y(n)|H_0)$ is estimated from the samples acquired during the first seconds of an imaging session. Instead of assuming this pdf to be Gaussian and estimate it parametrically, we propose here to estimate it nonparametrically by kernel smoothing. For this purpose we use the Epanechnikov kernel, that is defined as [8]:

$$K_E = \begin{cases} \frac{3}{4} \cdot (1 - x^2) & \text{for } |x| \leq 1 \\ 0 & \text{else} \end{cases} \quad (5)$$

A plot of the kernel is known in Figure 4.

The nonparametric estimate of $p(y(n)|H_0)$ is then computed by the function [14]:

$$f(x) = \frac{1}{n \cdot h} \cdot \sum_{i=1}^n K_E \left(\frac{x - x_i}{h} \right) \quad (6)$$




Figure 4: The Epanechnikov kernel.

where n is the total numbers of components in the sample and h is the kernel bandwidth.

The bandwidth is fixed and is computed as [14]:

$$h = \left[\frac{8 \cdot \sqrt{\pi} \cdot R(K_E)}{3 \cdot \mu_2(K_E) \cdot n} \right]^{\frac{1}{5}} \cdot \hat{\sigma} \quad (7)$$

with

$$\hat{\sigma} = \frac{SIQR}{\Phi^{-1} \left(\frac{3}{4} \right) - \Phi^{-1} \left(\frac{1}{4} \right)} \quad (8)$$

where $SIQR$ is the sample interquartile range, $\Phi^{-1}(p) = \sqrt{2} \cdot \text{erf}^{-1}(2 \cdot p - 1)$, $p \in (0, 1)$ is the quantile function of the normal distribution, $R(K_E)$ is defined as $R(K_E) = \int K_E^2(z) dz$ and $\mu_2(K_E)$ as $\mu_2(K_E) = \int z^2 \cdot K_E(z) dz$. h was shown to minimize the estimator error for kernels with fixed bandwidth.

The event-threshold T is determined such that the probability of $y(n)$ exceeding T given H_0 is α , the significance level. A typical value for α is $\alpha = 1 \cdot 10^{-4}$. T is computed by inverting $\Pr(y(n) > T | H_0) = \alpha$ based on $p(y(n)|H_0)$. As soon as the 98-percentile $y(n)$ exceeds this threshold, we mark the vessel-event at the corresponding image.

To improve detection, before hypothesis testing we attenuate the inherent variation of the feature curve by filtering with a time-varying causal recursive first-order low-pass filter. A feature curve before and after filtering is shown in Figure 5. The difference equation characterizing this filter is:

$$y(n) = a(n)x(n) + (1 - a(n))y(n-1) \quad (9)$$

$$\text{where: } a(n) = \begin{cases} k \in [0, 1) & \text{if } x(n) - x(n-1) \leq 3\hat{\sigma} \\ 1 & \text{else} \end{cases}$$

with $\hat{\sigma}$ the standard deviation estimated for the null hypothesis from the first images of the sequence, as described in equation (8). The filter thus smoothes within stationary time intervals, but preserves what it assumes to be a transition.

3. EXPERIMENTS AND DISCUSSION

We have detected contrast-injections in 18 sequences. The sequences were recorded during nine different interventions. All sequences were acquired under similar conditions, i.e., constant X-ray dosage, no patient nor table movement. Thus, the percentile-based feature responds only to variations of the observable vessel surface. The patients were mostly allowed to breath freely, but in some cases they were required to hold their breath. Some patients had also an open-heart surgery,




Figure 5: 98 percentile feature curve before (a) and after (b) filtration for an interventional sequence. The vessel event is marked by a bullet on the filtered feature curve.

and the sewing wires were visible in the analyzed sequences. The analyzed sequences contained images with a resolution of 512×512 pixels and were quantized to eight bits. To compute the vessel map the same parameters set was used for all sequences. In our experiments, the maximum vessel diameter was 19 pixels. The window size of the Bothat operator was thus chosen to be 21×21 pixels. We chose empirically $k = 0.1$ for the low-pass filter.

The significance level was: $\alpha = 3 \cdot 10^{-3}$. The size of the kernel was: $h = 2.3684$, computed as average value over six sequences. During testing, α and h were held constant for all sequences. Alternatively we have also conducted experiments where h was computed for each analyzed sequence independently. The classification results do not change, however such an approach is more robust, as the algorithm adapts better to the data.

For the learning phase we take 72 images, which are acquired over a period of six seconds, with a frame rate of 12fps. This is justified by the need to sample at least one complete respiration cycle, considering that under normal circumstances, a human breathes between 15 to 20 times a minute.

The algorithms described in this contribution give the solution to a classification problem – i.e. for each new image decide if it shows any vessels. We evaluate them using discrepancy methods [3], [15], by analyzing the frequency of correct and incorrect decisions relative to a golden standard [12]. The golden-standard was achieved after detecting per expert-visual analysis the first image which shows vessels. Taking into consideration the variability of such a visual segmentation, we consider that the automatic detection was successful if it falls in an interval of plus/minus two images from the expert-reference. In all sequences the detection was successful.

We have segmented the 18 sequences currently available to us also with the algorithm described in [5], where the likelihood of the normal case is estimated parametrically under the Gaussian assumption. It detected the vessel-event correctly only in 15 sequences, which means that on this data set, the new algorithm represents a 20% improvement over the old one.

4. CONCLUSIONS

Our vessel-event detection algorithm works under a set of assumptions. The most important one is that the first six seconds of an imaging session show no vessels. As in this

interval we collect our “training” data, it should include a good sample of the normal case and cover a few heart beats and a full respiration cycle. This is particularly needed for projection angles which permit also the visualization of the diaphragm, as in this case, the 98-percentile follows the variation of the projected diaphragm area. We believe that this is a consequence of the automatic gain control unit trying to compensate the variation of the image brightness caused by the moving diaphragm.

The contrast agent that reaches the vessels mixes with blood and is carried away by this, thus at the end of the imaging session, there is usually no contrast agent left into the vessels. However, the contrast agent present in the catheter is not washed away by blood. If a new imaging session begins with the catheter with contrast agent in place, then there will be something similar to vessels visible from the very beginning. By handling each imaging session independently, we are also robust against this case and can detect the moment when contrast agent reaches the vessels, even if some parts of the catheter with contrast agent in it are visible in the analyzed sequence.

In this contribution, we estimate the pdf of the feature given the normal case nonparametrically. Nonparametric estimation provides a better, robuster statistical model for the specific problem we are dealing with. This leads to better results in comparison to a parametric approach. Using the Epanechnikov kernel, the estimation result does no longer have an infinite support. We have chosen this specific kernel, also because it was proven to fulfill some optimality criteria for nonparametric density estimation [14]: it minimizes the mean squared error of the estimator.

REFERENCES

- [1] T. Aach, A. P. 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*, volume 5299, pages 283–294, San Jose, CA, USA, January 18–22 2004. SPIE.
- [2] T. Aach, C. Mayntz, P. Rongen, G. Schmitz, and H. Stegehuis. Spatiotemporal multiscale vessel enhancement for coronary angiograms. In *Medical Imaging: Image Processing*, volume 4684, pages 1010–1021, San Diego, CA, USA, February, 23–28 2002. SPIE.
- [3] J. S. Cardoso and L. Corte-Real. Toward a generic evaluation of image segmentation. *IEEE Transactions on Image Processing*, 14(11):1773–17872, 2005.
- [4] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. Technical report, University of Minnesota - Computer Science and Engineering Technical Report, 2007.
- [5] A. P. Condurache, T. Aach, K. Eck, and J. Bredno. Fast detection and processing of arbitrary contrast agent injections in coronary angiography and fluoroscopy. In *Proceedings of BVM-2004*, pages 5–9, Berlin, Germany, March 29–30 2004. Springer.
- [6] A. P. Condurache, T. Aach, K. Eck, J. Bredno, and T. Stehle. Fast and robust diaphragm detection and tracking in cardiac X-ray projection images. In *Medical Imaging 2005: Image Processing*, volume 5747,

pages 1766–1775, San Diego, CA, February 12–17 2005. SPIE.

- [7] E. R. Dougherty. *An introduction to morphological image processing*. SPIE Optical Engineering Press, 1992.
- [8] V. Epanechnikov. Nonparametric estimation of a multivariate probability density. *Theory of Probability and its Applications*, 14:153–158, 1969.
- [9] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. *Lecture Notes in Computer Science*, 1496:130–137, 1998.
- [10] D. Marr and E. Hildreth. Theory of edge detection. In *Proceedings of the Royal Society London B*, volume 207, pages 187–217, 1980.
- [11] A. Papoulis. *Probability & statistics*. Prentice-Hall International, Englewood Cliffs, 1990.
- [12] M. Sonka and J. M. Fitzpatrick. *Handbook of Medical Imaging*. SPIE Press, 2000.
- [13] V. Torre and T. Poggio. On edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):147–163, 1986.
- [14] M. Wand and M. Jones. *Kernel Smoothing*. Chapman and Hall, 1995.
- [15] Y. J. Zhang. A survey on evaluation methods for image segmentation. *Pattern Recognition*, 29(8):1335–1346, 1995.