ELSEVIER

Computerized Medical Imaging and Graphics 29 (2005) 343–355

Computerized Medical Imaging and Graphics

www.elsevier.com/locate/compmedimag

Vessel enhancement in digital X-ray angiographic sequences by temporal statistical learning

András Lassó[a,*], Emanuele Trucco[b]

[a]Department of Control Engineering and Information Technology, Budapest University of Technology and Economics, Magyar tudósok körútja 2, Budapest 1117, Hungary

[b]Department of Electrical, Electronic and Computer Engineering, School of Engineering and Physical Sciences, Heriot Watt University, Riccarton, Edinburgh EH14 4AS, UK

Received 20 October 2004; revised 28 February 2005; accepted 28 February 2005

Abstract

In this paper, we present a vessel enhancement method, SVM temporal filtering (STF), for X-ray angiographic (XA) images using Support Vector Machine (SVM). We show that the linear SVM applied to vessel enhancement can be regarded as a matched linear filter optimizing the contrast-to-noise ratio in XA images. We propose a non-linear kernel function for the SVM leading to good enhancement with noisy, varying grey-level dynamics at vessel pixels. One key advantage over the matched filters is that an optimal filter is learnt from images, not estimated at design stage. Results on clinical XA images show that learning-based enhancement achieves better results compared to simple subtraction and other image stacking methods.

Keywords: X-ray angiography; Image enhancement; Matched filters; Learning systems; Support vector machines

1. Introduction

This paper presents a method for vessel enhancement in X-ray angiographic (XA) image sequences. A Support Vector Machine (SVM) learns the function giving grey-level at vessel pixels in time, due to the flow of a contrast medium. Such function is then used for SVM classification of pixels as vessel or non-vessel.

Morphological or anatomical information about the vascular system is essential for enhancement of vascular diseases and planning of surgical procedures or catheter interventions. Vessels, organs of interest and tools (e.g. the tip of the catheter, balloon, stent) have to be displayed with the highest spatial and temporal resolution and fidelity so that their size, relative position and temporal changes can be estimated accurately, while irrelevant background structures (bones, muscles, etc.) should be suppressed. Despite developments in vascular imaging (e.g. MR, PET, US), X-ray [1,2] remains an important diagnostic and therapeutic tool. In X-ray angiography, a radio-opaque contrast agent (also called contrast medium, indicator or bolus) is usually injected into the vessels of interest by means of a catheter. The contrast agent travels through the vessels, making them visible in the X-ray, and is eventually washed out by the blood stream. An X-ray sequence capturing the passage of the agent through a given vessel region gives information about the vessel morphology and the dynamics of the blood flow, but, in general, not all vessels of interest are opacified equally and simultaneously in any frame.

Current methods computing morphological information from a XA image sequence tend to utilize only a fraction of the spatio-temporal information available. The use of temporal information (i.e. the evolution of the grey-level values in time) is generally limited to simple averaging and subtraction operations on selected frames.

The prevalent technique in clinical practice is *Digital Subtraction Angiography* (DSA), surveyed by Meijering et al. in [3]; here, we summarize the points relevant for our work. In DSA, a *mask frame* is acquired before the appearance of the contrast agent, then subtracted from subsequent frames containing opacified vessels (*live* or

* Corresponding author. Tel.: +36 20 4351999.
E-mail addresses: lasso@topcat.iit.bme.hu (A. Lassó), e.trucco@hw.ac.uk (E. Trucco).

0895-6111/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.compmedimag.2005.02.002

contrast frames) to enhance the vasculature. A basic technique to compose a single frame containing all the contrasted vessels is *maximum opacification*, which collects the lowest grey-level (if the agent darkens the vessels) in the sequence at each pixel position. Unfortunately, results include several artifacts in addition to vessel pixels, due, among others, to noise, variations in the power of the X-ray source, and patient motion. Sensor noise and artifacts that can be modeled as random noise can be somewhat reduced by *mask averaging*, i.e. using as mask frame the average of several frames taken before the contrast agent appears. In general, the suboptimal performance of DSA techniques for determining vessel morphology depends crucially on the fact that not all vessels are opacified simultaneously, which calls for a proper temporal analysis of the grey-level dynamics. This is the focus of our work.

Temporal bandpass filtering (TBF) methods have been proposed as an alternative to DSA including temporal dynamics [4–9], with the aim of achieving high-quality images with optimal contrast-to-noise ratio and better dose efficiency.

Contrast-to-noise ratio (CNR) is defined as

$$\mathrm{CNR} = \frac{\bar{V} - \bar{N}}{\sqrt{\sigma_{\mathrm{V}}^2 + \sigma_{\mathrm{N}}^2}}, \tag{1}$$

where $\bar{V}$, $\bar{N}$ and σ_{V}, σ_{N} are, respectively, the means and standard deviations of intensities at vessel and background (non-vessel) pixels.

The temporal dynamics is captured by the *indicator dilution curve* (IDC), a function describing the temporal evolution of the grey levels at a given pixel in a contrasted vessel. As an example, Fig. 1 shows the positions of four vessel pixels in three different frames of an image sequence, and Fig. 2 the respective IDCs.

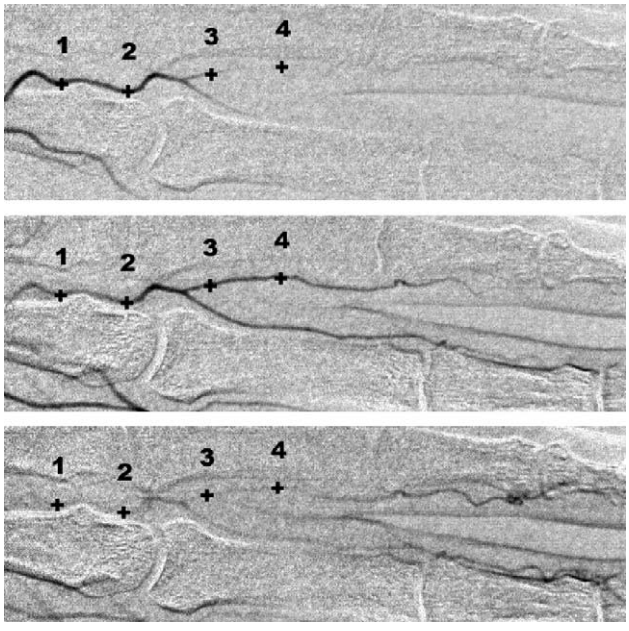

Fig. 1. Three frames of a XA image sequence. Note that different parts of the vessel tree are opacified in each frame.

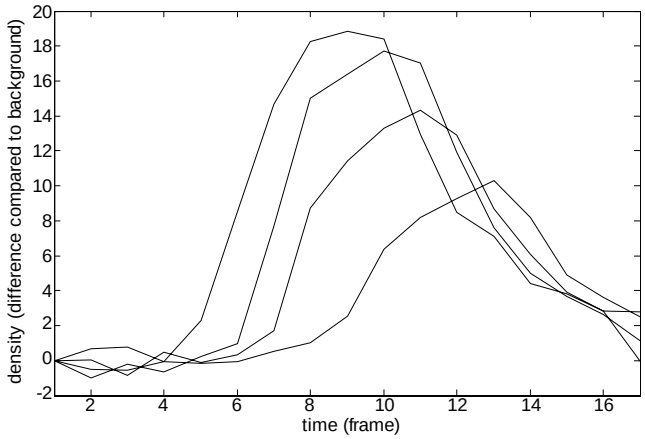

Fig. 2. Grey-level evolution (IDC curves) for the 4 pixels in Fig. 1. The position and shape of the curves change smoothly along the segment. The average number of samples per curve is 10. Logarithmic processing was applied, and background subtracted.

TBF methods model the IDC through recursive filters [4] or matched filters [5]. Reported filters are designed manually from observations in example sequences. The benefit of modeling temporal dynamics are apparent: [3,5] reports an improvement in contrast-to-noise ratio by a factor of 2–3 at a parity of X-ray dose, and dose reduction by a factor 4–6 at a parity of contrast-to-noise ratio with respect to simple subtraction.

However, the crucial problem for temporal methods is how to model the evolution of grey levels with sufficient generality. Several factors make this evolution difficult to capture in filters or models constructed manually. For instance, the IDC is different in different parts of the body; it depends on vessel anatomy (diameter, wall elasticity, etc.) and on the distance from the point of injection of the contrast agent; and it is influenced by noise, variations of the X-ray source, and patient motion. Kruger et al. [7] considered small, user-defined windows around vessel pixels, instead of single pixels, to reduce the effect of noise. To handle the blood flow differences between vessel segments, Kruger and Liu [5] recommended determining the IDC in several windows and applying the acquired model to vessel segments in the neighborhood of each window. However, they did not give guidance about how to implement such a system.

Kump et al. [6] used approximate matched filtering (AMF) consisting of a correlation with a kernel that approximates the optimal matched filter kernel, followed by a maximum opacity operation. They reported that the AMF nearly achieved the performance of the optimal matched filter on simulated images. However, the maximum opacity operation makes this method sensitive to noise and motion present in clinical images, and user interaction is needed to mark vessel points for obtaining filter kernel estimates.

The complexity and generality of the temporal evolution of grey levels, so hard to capture by a hand-crafted model, can be, however, *learnt automatically* from examples. We propose a novel method for enhancing vessel visibility based on statistical learning, a technique which has proven capable of tackling very complex modeling problems in computer vision [14–16]. The method relies on *Support Vector Machines* (SVM), which have been successfully applied to a wide range of classification and regression problems thanks to their good generalization performance, efficiency and robustness [10–13]. We use a SVM first to learn the IDC at vessel pixels from training sequences, then to classify pixels as vessel or not vessel in the same sequence, or in other sequences that were acquired with the same examination protocol. We also show that, using a linear SVM, the result of the classification is equivalent to matched filtering (but notice that the SVM *learns* the optimal filter, while manual filter design must *estimate* it). Although we do not address the problem of motion compensation directly, our method reduces artifacts due to motion or inaccurate registration as they result in IDC significantly different from those of vessel pixels. Our method offers therefore a principled and efficient solution to the problem of obtaining the IDC without user interaction.

The rest of this paper is organized as follows. A brief summary of XA imaging and SVMs is given in Section 2. Section 3 introduces our method for vessel enhancement. We present results obtained on simulated and clinical XA image sequences in Section 4. Section 5 closes the paper with concluding remarks and possible future research topics.

2. Theoretical background

2.1. Constructing an optimal matched filter for vessel enhancement

In this section, we recall briefly the contrast-to-noise ratio formula for a matched filter applied to vasculature enhancement, following [5]. The relation between matched filtering and the proposed SVM temporal filtering (STF) method is discussed in Section 3.1.

We assume a logarithmic conversion in the imaging process so that the grey-level of a pixel is proportional to the logarithm of the incident X-ray energy flux at the corresponding detector elements. Therefore, the concentration of the contrast agent in a vessel is proportional to the grey levels of the vessel's pixel [7]. Given this equivalence between concentration and pixel values, we can express the IDC in intensity values, as done so far.

Kruger and Liu [5] constructed a matched filter of the form

$$d(m,n) = \sum_{t=1}^{N} k_t f(m,n,t),$$

where $f(m,n,t)$ is the grey-level of the pixel at (m,n) position in the tth frame of the sequence consisting of N frames, k_t is the filter's coefficient and $d(m,n)$ is the resulting filtered image. The filter itself was designed by hand from the grey levels measured at a vessel pixel, as follows.

Assuming white Gaussian noise with constant standard deviation, and the normalizing constraint

$$\sum_{j=1}^{N} k_j = 0, \tag{2}$$

which forces to zero the output of non-vessel, stationary background pixels, it can be shown [5] that the contrast-to-noise ratio (filter output to noise) is maximized when

$$k_t = s(t) - \frac{1}{N}\sum_{j=1}^{N} s(j), \tag{3}$$

where $s(t)$ is the value of the IDC (or the grey-level measured at a vessel pixel) at time t. Thus, the filter has the same shape as the IDC, vertically shifted so that the sum of the coefficients is zero.

2.2. Support vector machines

In this section, we give a concise account of *support vector machines* (SVM) performing *binary* classification. The reader is referred to [13–16] for details on the theory and application of SVM.

An SVM projects data (sample points) in a high-dimensional space and determines a discriminating hyperplane in that space, i.e. sample points belonging to the same class lie on the same side of the hyperplane. SVMs minimize the *structural risk*, i.e. the probability of misclassifying previously unseen data generated by a fixed but unknown probability distribution [13]. This contrasts with other learning methods minimizing the *empirical risk*, i.e. the probability of misclassifying data from the training set. Section 3.1 shows how we use SVM for enhancement of vessels based on grey-level changes over time using not only the sign, but also the absolute value of the output of the SVM (the distance from the optimal separating hyperplane).

2.2.1. Linearly separable data set

For a *linearly* separable data set, $S = (\mathbf{x}_i, y_i)_{i=1}^{l}$, $\mathbf{x}_i \in R^N$, with class labels $y_i \in \{-1,1\}$, the goal of binary classification is to find the hyperplane that divides the set so that points that belong to the same class lie all on the same side of the hyperplane. The separating hyperplane, which minimizes the structural risk, is the one having maximum distance from the training samples of both classes.

A hyperplane can be defined by the equation $\langle \mathbf{w},\mathbf{x}\rangle + b = 0$, where diamond brackets indicate inner product, and $\mathbf{w} \in R^N$, $b \in R$. By definition, for a linearly separable set there exist a $(\mathbf{w}, b)$ such that, for all samples

$$y_i(\langle \mathbf{w}, \mathbf{x}_i\rangle + b) \geq 1. \tag{4}$$

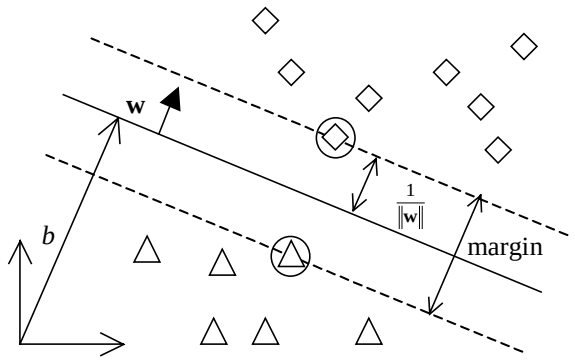

Fig. 3. Optimal separating hyperplane in R^2. Support vectors (circled) are the closest point to the hyperplane.

Choosing $(\mathbf{w}, b)$ so that the distance of the closest point is $1/\|\mathbf{w}\|$, the *optimally separating hyperplane* (OSH, also called *maximal margin hyperplane*, see Fig. 3) is obtained by maximizing $1/\|\mathbf{w}\|$, subject to inequality (4).

The Lagrangian formulation of this problem is

$$\min L(\mathbf{w}, b, \boldsymbol{\alpha}) = \min\left[\frac{1}{2}\langle \mathbf{w}, \mathbf{w}\rangle - \sum_{i=1}^{l} \alpha_i\{y_i(\langle \mathbf{x}_i, \mathbf{w}\rangle + b) - 1\}\right],$$

where $\alpha_1,\ldots,\alpha_l$ are the non-negative Lagrange multipliers associated with the inequality constraints (4) and the minimization is performed over $\mathbf{w}$, b and $\boldsymbol{\alpha}$.

The solution to the problem is equivalent to determining the saddle point $(\mathbf{w}^*, b^*, \boldsymbol{\alpha}^*)$ of the function L: where L has a minimum in terms of $\mathbf{w}$ and b, and a maximum in terms of $\boldsymbol{\alpha}$. This problem can be reformulated as a *dual problem* [13,16]: maximize the function

$$L_{\mathrm{d}}(\boldsymbol{\alpha}) = \sum_{i=1}^{l} \alpha_i - \sum_{i,j=1}^{N} \alpha_i \alpha_j y_i y_j \langle \mathbf{x}_i, \mathbf{x}_j\rangle,$$

subject to $\alpha_i \geq 0$, $i=1,\ldots,l$ and $\sum_{i=1}^{l} y_i \alpha_i = 0$.

From the solution $\boldsymbol{\alpha}^*$ the parameters of the optimal $(\mathbf{w}^*, b^*)$ hyperplane can expressed as

$$\mathbf{w}^* = \sum_{i=1}^{l} \alpha_i^* y_i y_j \mathbf{x}_i, \quad b^* = y_j - \langle \mathbf{w}^*, \mathbf{x}_j\rangle,$$

where $\mathbf{x}_j$ is any sample such that α_j^* is non-zero. It can be shown that α_j^* can be non-zero only for the $\mathbf{x}_j$ closest to the optimal separating hyperplane. These $\mathbf{x}_j$ are called *support vectors*. Thus, $\mathbf{w}^*$ becomes a linear combination of only a *small* fraction of the total number of sample points, and a new data point, $\mathbf{x}$, is classified by simply evaluating the sign of the expression $\langle \mathbf{w}^*, \mathbf{x}\rangle + b^*$ involving only the few support vectors.

2.2.2. Linearly non-separable data set

Linear separation is rarely possible in practical problems. *Soft-margin classification* generalizes hyperplane-based classification by accepting a minimum number of misclassified points. To achieve this, the inequality constraints (4) is modified by introducing additional *slack variables* $\xi_i \geq 0$, $i=1,\ldots,l$, so that, for all samples,

$$y_i(\langle \mathbf{w}, \mathbf{x}_i\rangle + b) \geq 1 - \xi_i. \tag{5}$$

The variables ξ_i are non-zero only for the $\mathbf{x}_i$ which do not satisfy (4). We obtain the *generalized optimal separating hyperplane* by solving the following optimization problem: minimize $\langle \mathbf{w}, \mathbf{w}\rangle + C\sum_{i=1}^{l} \xi_i$ subject to (5). The role of parameter C is to balance the number of misclassified points against the minimum distance between the separating hyperplane and the samples. Small C values allow more misclassified points (more errors) but achieve a larger minimum distance (more overall robustness).

It can be shown [13] that this soft-margin classification problem is equivalent to finding the maximal-margin hyperplane, with the additional constraint that all α_i are upper bounded by C, which intuitively limits the influence of outliers.

2.2.3. Non-linear SVM

In many real-world applications, the separating function to be learned cannot be expressed as a simple linear combination of the data vectors. In this case, a strategy commonly used in statistical learning is to transform the data into abstract 'features' via a non-linear mapping, i.e. to map data points, $\mathbf{x}$, from a given input space to points $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}),\ldots,\phi_N(\mathbf{x}))$ in a feature space.

The advantage of a suitable non-linear mapping is that the original non-linear problem could be solved in a higher-dimensional feature space by a linear classifier. A key problem is to determine the smallest set of features $\phi_1,\ldots,\phi_N$ capturing the essential information contained in the original data, keeping computational costs low and generalization performance high. In this sense, SVMs minimize both the error and the complexity of the classifier [14].

An important property of SVMs is that the data appears always in the form of inner products. Therefore, a so-called *kernel function* is defined for a specific feature mapping

$$K(\mathbf{x}, \mathbf{z}) = \langle \phi(\mathbf{x}), \phi(\mathbf{z})\rangle,$$

which calculates the inner product of the transformed data. Notice that the kernel function can be evaluated more efficiently than implementing the above formula literally [16]. Only symmetric, positive definite function satisfying certain conditions, called *Mercer's conditions* [14], are valid kernels.

In non-linear SVMs, the classification of a sample $\mathbf{x}$ is given by the sign of the following expression

$$\sum_{i=1}^{N_s} \alpha_i^* y_i K(\mathbf{s}_i, \mathbf{x}) + b^*, \tag{6}$$

where N_s is the number of support vectors determined in the learning phase, $\mathbf{s}_i$ is a support vector, y_i is its label, and α_i^* and b^* are values determined in the learning phase.

3. Vessel enhancement by temporal statistical learning

This section gives a detailed description of our method. First we show how a SVM can be used for vessel enhancement in XA images and how to choose the kernel and learning parameters. The sensitivity of the method to the parameters is discussed in Section 4.1. Then we describe how to obtain the necessary training set for the learning phase, and how we generated simulation data and implemented the SVM software.

3.1. SVMs as matched filtering for vessel enhancement

In this section, we show that the distance of a sample from the hyperplane that optimally separates sample vectors is equivalent to the output of an optimal matched filter for the XA image, in the sense that the contrast-to-noise ratio is maximized. We use the assumptions of Section 2.1, i.e. stationary background, and additive, uncorrelated, time-independent Gaussian noise.

Consider a discrete, finite function **x** holding the grey levels of a pixel throughout the sequence:

$$\mathbf{x} = (f(m,n,1), f(m,n,2), \ldots, f(m,n,N)). \tag{7}$$

This function can be written for a non-vessel pixel (stationary background) and for a vessel pixel, respectively, as

$$\mathbf{x}_{\mathrm{n}} = (p + \nu_{\mathrm{n},1}, p + \nu_{\mathrm{n},2}, \ldots, p + \nu_{\mathrm{n},N}), \tag{8}$$

$$\mathbf{x}_{\mathrm{v}} = (p + s_1 + \nu_{\mathrm{v},1}, p + s_2 + \nu_{\mathrm{v},2}, \ldots, p + s_N + \nu_{\mathrm{v},N}), \tag{9}$$

where p is the grey-level of the stationary background, s_i is the value of the IDC at time i, and $\nu_{\mathrm{n},i}$, $\nu_{\mathrm{v},i}$ are the random noise.

If there is no noise, and there is at least one non-zero IDC value, a separating hyperplane can be found (Fig. 4). The non-vessel points are located on the line $(1,1,\ldots,1)q$, and vessel points on the parallel line $(1,1,\ldots,1)r+(s_1, s_2,\ldots,s_N)$, $q,r \in R$. The optimal separating hyperplane is defined by the equation $\langle \mathbf{w},\mathbf{x}\rangle + b = 0$, $\mathbf{w} \in R^N$, $b \in R$, and the distance between the closest sample and the hyperplane is $1/\|\mathbf{w}\|$.

We can determine **w** and b with the help of normal vector **n** that is perpendicular to both the vessel and non-vessel lines,

$$\langle \mathbf{n}, (1,1,\ldots,1)q \rangle = 0, \tag{10}$$

and the length of which is the distance between the two lines. Since the non-vessel line crosses the origin, **n** is a point on the vessel line:

$$\mathbf{n} = (1,1,\ldots,1)r + (s_1, s_2, \ldots, s_N). \tag{11}$$

Substituting **n** from (11) into (10) gives

$$q \sum_{i=1}^{N} (r + s_i) = 0,$$

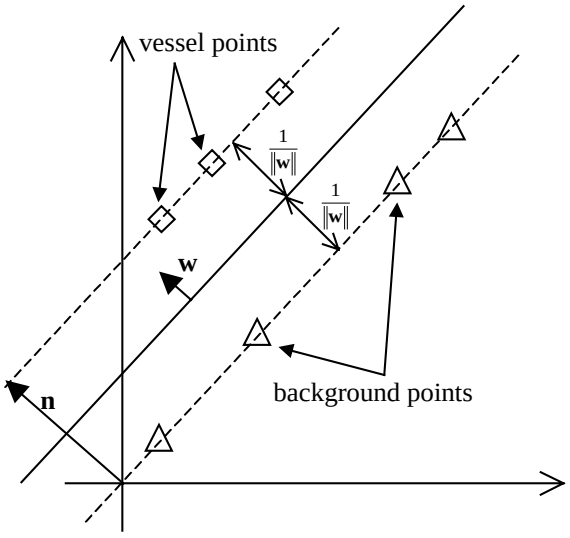

Fig. 4. Determining the OSH between stationary background and background with IDC added, without noise.

and therefore

$$r = -\frac{1}{N} \sum_{i=1}^{N} s_i = -\bar{s}.$$

The normal vector **n** can be expressed as $\mathbf{n} = (s_1 - \bar{s}, s_2 - \bar{s}, \ldots, s_N - \bar{s})$, and the hyperplane is:

$$\mathbf{w} = \frac{2}{\|\mathbf{n}\|^2}\mathbf{n}, \quad b = -1.$$

The output of the classifier, i.e. the signed distance from the optimal separating hyperplane for sample **x**, is therefore\

$$\begin{aligned}\langle \mathbf{w}, \mathbf{x}\rangle + b &= \frac{2}{\|\mathbf{n}\|^2}\langle \mathbf{n}, \mathbf{x}\rangle + b \\ &= \frac{2}{\sum_{i=1}^{N}(s_i - \bar{s})} \sum_{i=1}^{N} [x_i(s_i - \bar{s})] - 1 \\ &= a(s_1, \ldots, s_N) \sum_{i=1}^{N} [x_i(s_i - \bar{s})] - 1,\end{aligned}$$

where $a(s_1,\ldots,s_N)$ is a scaling constant independent of $\mathbf{x}_i$:

$$a(s_1, \ldots, s_N) = \frac{2}{\sum_{i=1}^{N}(s_i - \bar{s})}.$$

According to (2) and (3) the output of the optimal filter for sample **x** is

$$\sum_{i=1}^{N} \left[x_i \left(s(i) - \frac{1}{N} \sum_{i=1}^{N} s(i) \right) \right]$$

$$= \sum_{i=1}^{N} \left[x_i \left((p + s_i) - \frac{1}{N} \sum_{i=1}^{N} (p + s_i) \right) \right] = \sum_{i=1}^{N} [x_i(s_i - \bar{s})],$$

where $s(i)$ is a grey-level measured at a vessel pixel, including an unknown stationary background, and s_i is an

IDC value (without the background) as defined in (8) and (9).

Thus, the signed distance from the optimal separating hyperplane equals the normalized output of the optimal linear matched filter. The normalization is carried out by a scaling and offset giving $+1$ for vessel, -1 for background. If the samples are noisy, a linearly separating hyperplane may not exist, so we have to use the soft-margin classifier. Therefore, a SVM with a simple linear kernel is applicable to distinguish vessel and non-vessel pixels from the temporal evolution of their grey levels, and the signed distance from the optimal separating hyperplane is a meaningful metric for 'vesselness'.

In the general case, noise affects both grey levels and labels, as incorrect labeling of training samples may occur, and there are slightly different IDCs in different vessel segments. The effects of these outliers are controlled by the parameter C (Section 2.2.3): the SVM maximizes the global minimum contrast-to-noise ratio, where 'noise' includes data noise, label noise and the effect of IDC differences.

C values influence classification as follows. High C values maximize the margin in the presence of high data noise, but too high values tend to misclassify valid vessel points as outliers. Low C values improve the CNR in low-contrast pixels (where IDC values are low compared to the local background), but at the cost of a smaller improvement in high-contrast pixels. Finally, C gives an upper limit on the maximum weight of a sample in determining the OSH, therefore limiting the effect of label noise.

3.2. Choosing a kernel

Choosing a kernel, and therefore a mapping of the input samples to a high-dimensional feature space, is a fundamental decision that determines the complexity of the target function to be learned by the SVM. We can base this decision on prior knowledge about the problem, or evaluate commonly used kernels, relying on the good generalization properties of SVMs. In our case, the interpretation of the distance from the maximal margin hyperplane as an optimal vessel enhancement filter (Section 3.1) justifies the use of the simple linear kernel,

$$K_{\mathrm{L}}(\mathbf{x},\mathbf{z}) = \mathbf{x}\cdot\mathbf{z}. \tag{12}$$

Given the dynamics of the propagation of the contrast agent in the blood flow, one expects a simple unimodal IDC as shown in Fig. 2. However, the IDC of different vessel segments can differ significantly, depending on, among others, distance from the location where the contrast media is injected, anatomy of the vessel, blood velocity, etc. We notice that most such factors change smoothly in space. For example, Fig. 2 shows four IDCs measured, respectively, at four points along a same vessel; smooth variation with spatial location is evident. This suggests to introduce a *spatial localization constraint*, whereby only spatially close training pixels play a role in the classification of a sample. This leads to a CNR improvement, as the classification is not disturbed by distant, potentially different IDCs. An additional advantage is that the effect of IDC differences on the classification results is controlled by this constraint and the optimal C value can be chosen with more freedom.

Such localization could enhance the output of other filtering methods (see [4–9]), but at the price of significantly increased complexity; moreover, there is no proven way to do that. The locality constraint can be naturally built into the SVM framework by constructing an appropriate non-linear kernel. We combined linear and radial basis function (RBF) kernels to form a new kernel with the desired localization property

$$K_{\mathrm{RL}}(\mathbf{x},\mathbf{z}) = \exp\left(-\gamma\sqrt{(x_1-z_1)^2+(x_2-z_2)^2}\right)\sum_{i=1}^{N} x_{i+2}z_{i+2}, \tag{13}$$

where the first two elements of the sample vectors are the spatial coordinates, and subsequent elements the grey levels. Eq. (13) is a valid kernel function due to its construction from kernel functions [16]. The parameter γ controls the localization strength: the smaller its value, the more distant training vectors influence a specific point. For $\gamma = 0$, Eq. (13) becomes a simple linear kernel.

3.3. Kernel and learning parameters

The SVM determines the optimal separating hyperplane for the training set using the mapping into the feature space defined by the kernel. The result depends mainly on kernel type, its parameters, and the regularization parameter C (introduced in Section 2.2 for soft-margin classifiers).

The linear kernel has no parameters. The RBF-linear kernel has one parameter, γ, which controls the size of the neighborhood influencing the classification result. The optimal γ value depends primarily on how much the IDC shapes vary at different spatial positions. Choosing a small γ value results in a suboptimal CNR, because the SVM cannot learn different IDCs at different regions of the image. On the other hand, too high γ values mean too strict localization, i.e. only a few training samples determine the classification for each pixels; as noise makes the filter coefficients calculated from fewer samples less accurate, the result is a lower CNR. Since the blood-flow dynamics and image noise are similar in most XA images, a single value is expected to perform well for most of the images. Optionally, optimal values can be determined for specific examination protocols by fine-tuning the parameter on a number of images of the same type.

Expanding the comments on C in Section 3.1, we observe that C controls the effect of outliers due to misclassified or noise-corrupted training samples, and differences of positive sample vectors. The optimal value of this parameter depends on the scaling of the data. A low C value causes

most of the training vectors to be identified as support vectors, with all α_j^* multipliers equal to C. This prevents the SVM from identifying outliers properly ('overtraining'). Too high C values result in considering normal data points as outliers, preventing them from influencing the classification ('undertraining'). The algorithm works well, in general, with the following default C value (used by Joachims [11])

$$C = \frac{1}{\frac{1}{l}\sum_{i=1}^{l}\langle \mathbf{x}_i, \mathbf{x}_i\rangle}, \tag{14}$$

that is the inverse of the average squared norm of the training vectors (but notice that the learning problem can be reformulated so that a scale-dependent parameter v can be used instead of C [17]). We choose γ and C so to maximize the CNR measured on the resulting images.

Using a valid, positive-definite kernel function, the SVM is guaranteed to converge to a global optimal solution. The speed of convergence depends on the difficulty of the problem (number of training vectors, number of outliers, learning parameter value), the number of training samples, and the optimization algorithm adopted.

After the learning phase is completed, the SVM classifies the pixels of an input image by Eq. (5). The sign of the result gives the binary classification, and its absolute value a measure of 'vesselness', as explained in Section 3.1.

3.4. Obtaining training samples

The learning phase requires sets of examples of the IDC at both vessel pixels (positive samples) and non-vessel pixels (negative samples). We need therefore an automatic procedure picking such samples reliably.

There are several vessel-enhancement algorithms based on vessel tracking and other computer vision techniques [18]. The methods giving the best results seem to require user interaction or impose serious computational costs. Making use of the generalization properties of SVM and their good tolerance for misclassifications in the training set, we applied a simple automated method, inspired by [19], to locate vessel pixels that can be used as training points. The algorithm is conservative, i.e. designed to avoid false positives. Reliable candidates are taken as axial pixels (pixels along the vessel axis) as vessel borders are often blurred and difficult to locate. False positives are therefore limited in number. The algorithm is as follows.

1. Compute all the pixelwise, frame-to-frame differences (for all $j=1, 2,\ldots,(N-m)$ frames the pixels of frame j subtracted from frame $(j+m)$), threshold the results at level T, and combine them pixelwise with a logical OR. We chose $m=1$ s/frame rate (1 s time difference between the subtracted frames) and $T = \frac{1}{2}\overline{\max_i f(m_v, n_v, i) - \min_i f(m_v, n_v, i)}$ (threshold is half of the average difference between the highest and lowest grey-level of contrasted vessel pixels).
2. Apply morphological open using a 3×3 pixel disk-shaped structuring element, removing noise and smoothing contours.
3. Apply morphological skeletonization, keeping only central vessel pixels as the most reliable samples.
4. Randomly choose from the selected pixels to get the desired amount of vessel samples.

Non-vessel pixels are found by modifying the above algorithm to *avoid* points with large grey-level changes. However, since grey levels of non-vessel pixels are well modeled by Poisson noise [2], it is more efficient to measure the mean and standard deviation of the non-vessel pixel population and generate *synthetic* negative samples with these parameters. This method also guarantees no false-negative errors in the training set. In our experiments, using synthetic background pixels instead of real, manually selected one had little impact on classification results.

Using the linear kernel (12), the sample vector contains the grey levels of a pixel for all the frames, Eq. (7); with the combined RBF-linear kernel (13) the sample vector contains also the spatial coordinates (m,n):

$$\mathbf{x} = (m, n, f(m, n, 1), f(m, n, 2), \ldots, f(m, n, N)).$$

We chose the positions of non-vessel training pixels randomly, but found it better to use the same spatial distribution as the positive samples to avoid brightness variance caused by local imbalance in the number of positive and negative samples. The simplest way to achieve this is to use the spatial coordinates of the positive samples to generate a distribution of synthetic negative ones.

3.5. Synthetic data

In order to evaluate our method and compare it to other state-of-the-art methods in a controlled environment, we used synthetic (simulated) data generated by the model used in [6]. This model describes the transport of the contrast agent in an artery with the following differential equation:

$$\frac{\partial C}{\partial t} = D\frac{\partial^2 C}{\partial z^2} - v\frac{\partial C}{\partial z}, \quad z > 0. \tag{15}$$

This is a one-dimensional convective–dispersive model, where C is the average concentration of the contrast agent in the plane normal to the axial flow, depending on the axial position z and time t. The first term represents dispersion and the second term represents convection. The model parameters are D, the effective dispersion coefficient, and v, the average blood velocity. Noise is modeled by adding Poisson noise to the noise-free background.

To compare results, we adopted the model parameters used in [6], i.e. $D=84$ cm^2/s, $v=17$ cm/s, Poisson noise standard deviation $\lambda=6.0$, image frame rate $=4$ frames/s,

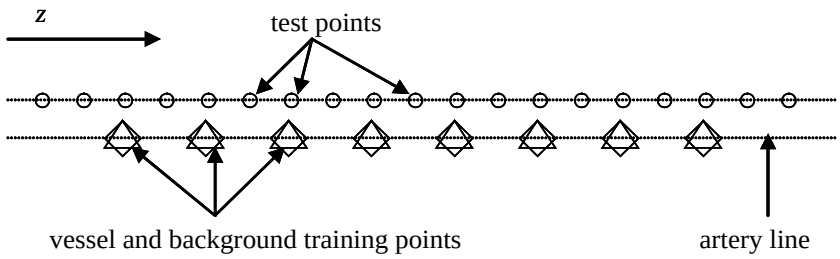

Fig. 5. Arrangement of vessel, background and test points for simulated data.

$z=0$–100 cm, $Q_0/Q=0.5$ (ratio of volume flow of the injected solution to the volume flow of blood), $\tau=3$ s (duration of the injection).

The training samples for the SVM are generated from this model as follows. Vessel points (positive samples) are evenly spaced along a line representing the artery, and the values are the grey values calculated from the model, Eq. (15), at the corresponding z positions. Background points (negative samples) are placed at the same spatial positions as vessel points and their grey values are Poisson random deviates. We generate the same amount of positive and negative samples, and these together form the training data for the SVM.

Note that, although we use a one-dimensional flow model, the filter output is slightly dependent on the distance of the test point from the artery line, due to the two-dimensional spatial localization constraint in the STF. We chose the test points to be on a line 0.2 cm far from the artery line (Fig. 5) because, typically, these test points can belong to vessel as well as background pixels in clinical images. To calculate the CNR we determine the STF output for vessel IDC and noise samples at these test points.

3.6. SVM implementation

Several SVM packages are available with full source code. We used SVMLight [11] and LIBSVM [20]. SVMLight was generally faster, but sometimes converged very slowly when used with simple linear kernels; LIBSVM was reliable, but usually slower. We used SVMLight for classification and non-linear training, and LIBSVM for linear training. We wrote a separate pre-processing program that extracts training and test vectors from the images and writes to a file that can be interpreted by the SVM programs. We wrote Matlab code to generate the random (non-vessel) training samples and to analyze the SVM output.

4. Results

This section summarizes test results of the proposed vessel (CNR) enhancement method on simulated and clinical image sequences.

4.1. Optimizing kernel and learning parameters

We choose the kernel and learning parameters of the SVM so to maximize the CNR defined by Eq. (1). To determine the optimal C learning parameter we calculated the CNRs with C ranging from 10^{-8} to 10 and kernel parameter γ from 10^{-5} to $10\,\mathrm{cm}^{-1}$ at a given artery position ($z=50$ cm). The result is shown in Fig. 6.

The highest CNR value is achieved at around $\gamma=0.15\,\mathrm{cm}^{-1}$, and the sensitivity of the CNR for using a sub-optimal γ value is primarily depending on the C learning parameter. The sensitivity is the smallest when C is very low, but in this case the SVM is overtrained (most training vectors are support vectors) and has poor generalization properties. Lack of generalization makes the results more sensitive to label noise and structured noise in the image (we show later the effect of structured noise on the CNR). Therefore, the C value determined by Eq. (14), which has a magnitude of 10^{-4}, is a good compromise between sensitivity for γ and overtraining.

A single, optimal γ value exists for maximizing the CNR at a given vessel point. Now we measure how this optimal γ changes along the artery line. Fig. 7 shows the dependency of CNR on γ at different vessel points. It can be seen that approximately the same γ value maximizes the CNR all over the artery line.

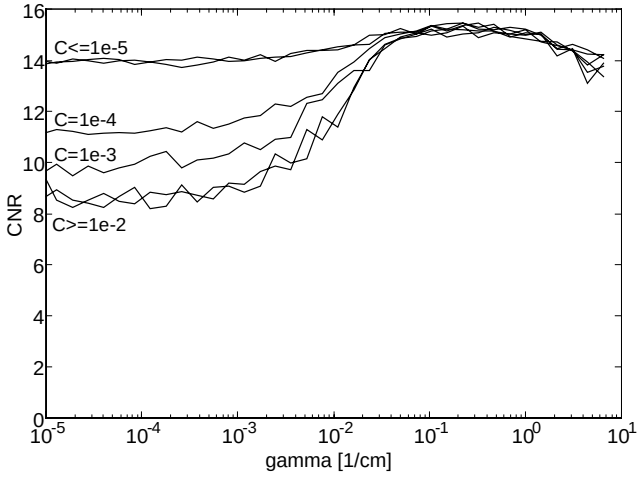

Fig. 6. Dependence of the CNR on γ and C, calculated from simulated data ($z=50$ cm).

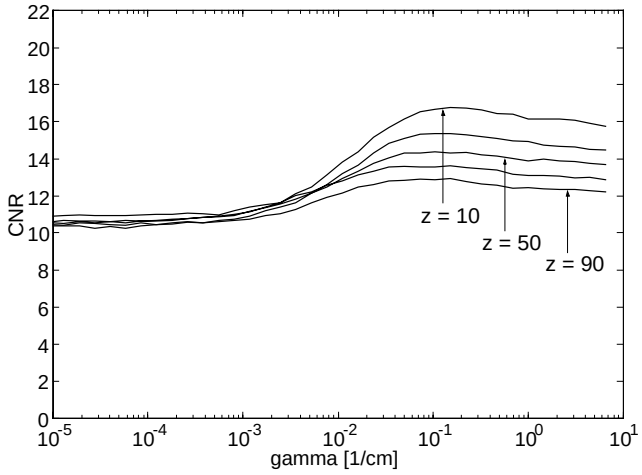

Fig. 7. Dependence of the CNR on γ and z, calculated from simulated data (C = default).

In summary, for an optimal CNR in the studied exam we shall use $\gamma = 0.15\ \mathrm{cm}^{-1}$ and the default C from Eq. (14).

4.2. Evaluation on simulated data

We implemented all the image stacking methods evaluated in [6] and compared them with the STF method. We refer the reader to [6] for details on the various methods. The resulting CNR along a vessel is shown in Fig. 8.

In accordance with [6], we found that recursive filtering (REC), maximum opacification (MO) and simple subtraction (REF) achieve the lowest CNRs. Approximate matched filtering (AMF) and SVM temporal filtering (STF) nearly achieved the CNR of the optimal (ideal) matched filter (MAT). As the optimal matched filter cannot be realized in practice, STF and AMF are the two most effective methods for image stacking, and STF performed slightly better than AMF (2–12% higher CNR) throughout the vessels in our sequences.

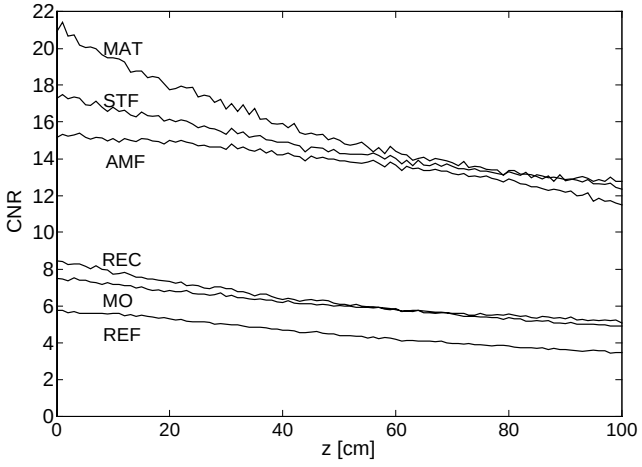

Fig. 8. CNR calculated from simulated data along the vessel line with six different stacking methods ($\gamma = 0.15\ \mathrm{cm}^{-1}$, C = default).

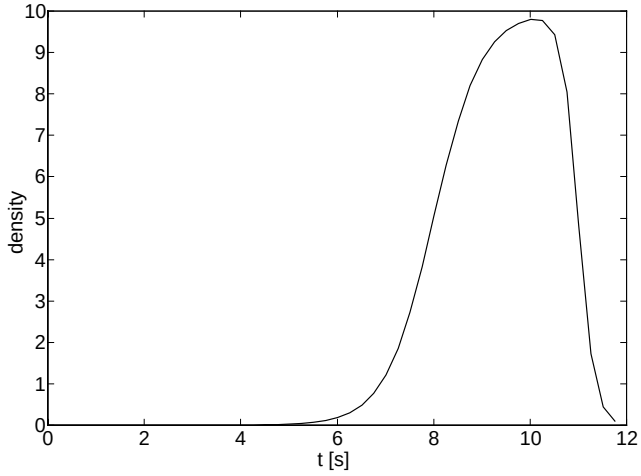

Fig. 9. Impulse that were represented the structured noise ($a_1 = 8$ s; $b_1 = 10/0.1$ s, $b_2 = 10/0.3$ s, $c = 10$; 10 g is about the half of the density of vessels).

The AMF output is the *maximum* of the cross-correlation of the grey-level evolution of a given pixel with the AMF kernel. This maximum operation makes AMF sensitive to structured noise, like motion artifacts. STF suppresses most of the artifacts because it enhances only those grey-level changes that spatio-temporally correlate with grey-level change caused by the contrast agent. We simulate motion artifact by adding an impulse function to the simulated background and vessel pixel densities:

$$x_{\mathrm{n},s} = \frac{c}{2}\left(\tanh\left(\frac{t-a_1}{b_1}\right) - \tanh\left(\frac{t-a_2}{b_2}\right)\right).$$

This is an asymmetrical impulse increasing sharply at about a_1, with maximum slope $1/b_1$, reaching its maximum value c, decreasing sharply at about a_2, with slope $1/b_2$, then vanishing. To construct a realistic signal, we determined the constants from observations of motion artifacts in clinical sequences. The resulting signal (Fig. 9) appears 6 s after injection and disappears a couple of seconds later; its density is about half of the density change caused by the contrast agent.

Fig. 10 compares the STF and AMF results with Poisson noise only and with Poisson and structured noise. The presence of structured noise does not significantly reduce the CNR of STF-enhanced vessels. However, the CNR drops by about 25% when we use the AMF method. Altogether STF achieves up to 45% better CNR compared to AMF in the presence of structured noise.

4.3. Evaluation on clinical images

In this section, we show that vessel morphology (size and position) is preserved by STF filtering, and contrast increased at both normal vessels and most clinically interesting places. Then we show the results of STF processing on clinical XA image sequences.

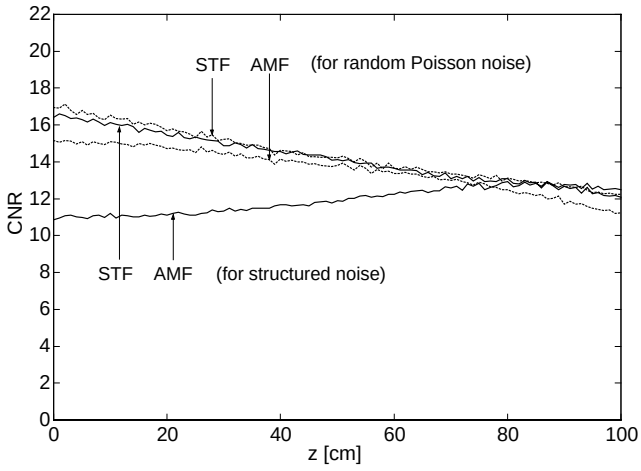

Fig. 10. Filter output when only Poisson random noise is present (dotted lines) and when also structured noise is present (solid line). STF $\gamma = 0.15\ \mathrm{cm}^{-1}$, C = default value, calculated from Eq. (14).

4.3.1. Classification of vessel border pixels

The IDC $\mathbf{x}_b$ at a vessel border pixel can be approximated by interpolating the values $\mathbf{x}_v$ of the closest pixel on the vessel axis, and $\mathbf{x}_n$ of the closest non-vessel pixel:

$$\mathbf{x}_b \approx \kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n. \tag{16}$$

Measurements on clinical images show that the approximation has the smallest error when $\kappa \approx 0.5$, as exemplified in Fig. 11.

Thus, border pixel grey values are between the vessel center and non-vessel pixel grey values: vessel border pixels are half-way between the vessel center and non-vessel pixels in the feature space.

Now we show that the position of a border pixel remains between the vessel center and non-vessel pixels in the classified image too. From Eq. (16)

$$\mathrm{SVM}(\mathbf{x}_b) \approx \mathrm{SVM}(\kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n),$$

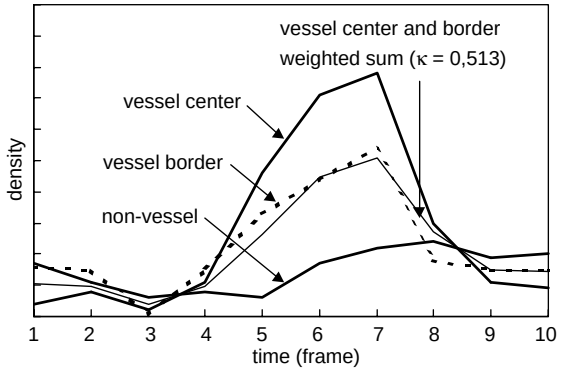

Fig. 11. IDC measured at vessel center, border and non-vessel pixel. The border IDC (thin solid line) can be constructed (dotted line) from the weighted sum of the vessel center and non-vessel pixel IDC.

and from Eq. (6)

$$\mathrm{SVM}(\mathbf{x}) = \sum_{i=1}^{N_s} \alpha_i^* y_i K(\mathbf{s}_i, \mathbf{x}) + b^*.$$

The STF operation with an invariant linear kernel (notice that the RBF-linear kernel is approximately invariant, see below) applied to Eq. (16) is

$$\begin{aligned}
&\mathrm{SVM}(\kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n) \\
&\quad = \sum_{i=1}^{N_s} \alpha_i^* y_i K(\mathbf{s}_i, \kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n) + b^* \qquad (17)\\
&\mathrm{SVM}(\kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n) \\
&\quad \approx \sum_{i=1}^{N_s} \alpha_i^* y_i \mathbf{s}_i \cdot (\kappa\mathbf{x}_v + (1-\kappa)\mathbf{x}_n) + b^* \\
&\quad = \left(\sum_{i=1}^{N_s} \alpha_i^* y_i \kappa \mathbf{s}_i \cdot \mathbf{x}_v + b^*\right) \\
&\qquad + \left(\sum_{i=1}^{N_s} \alpha_i^* y_i (1-\kappa)\mathbf{s}_i \cdot \mathbf{x}_n + b^*\right) \\
&\quad = \mathrm{SVM}(\kappa\mathbf{x}_v) + \mathrm{SVM}((1-\kappa)\mathbf{x}_n), \qquad (18)
\end{aligned}$$

which proves that the border pixel remains between the vessel center and non-vessel pixels in the classified image.

The approximation of Eqs. (17) and (18) is exact for the linear SVM kernel. In the case of RBF-linear kernel, the smaller γ, the more accurate the approximation, which imposes an upper limit on γ.

As border pixels are situated between the vessel center and non-vessel pixels both in feature space and in the classification results, the spatial position of vessel borders are not changed by the filtering.

As an example, Fig. 12 shows how the grey values are changing along the cross-section of a vessel (the vessel

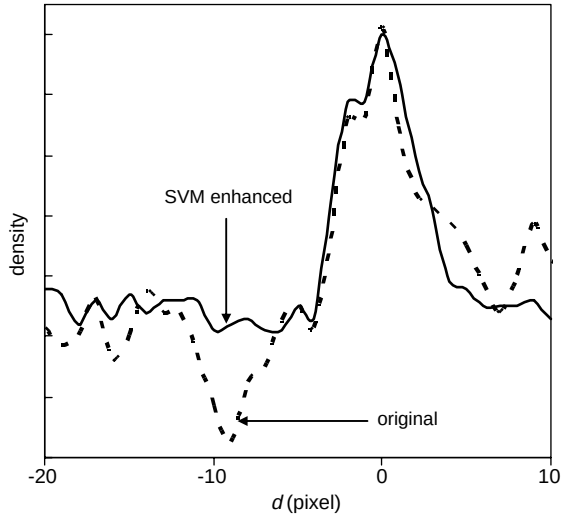

Fig. 12. Grey values along a line perpendicular to a vessel point. Vessel border position (around $x = -4$ and $x = 4$) is not changed by the filtering, but variation in non-vessel pixel grey values are reduced.

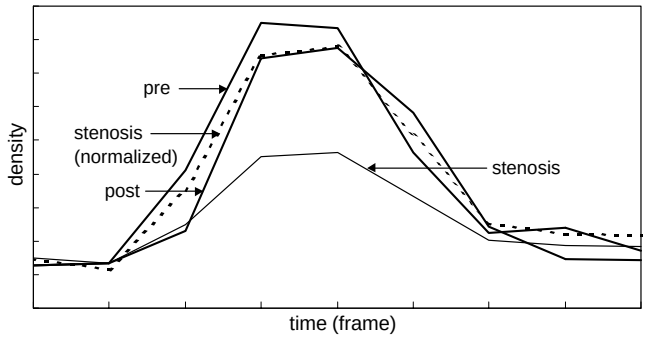

Fig. 13. Sample of IDC at a pixel before the stenosis (*pre*), in the middle of the stenosis (*stenosis*) and after the stenosis along the artery line (*post*). The IDC changes smoothly along the vessel line, and the normalized IDC at the stenosis is between the pre- and post-IDC.

centerline is at $d=0$). The vessel border position is not changed by the filtering (remains at around $d=-4, 4$), but variation in non-vessel pixel grey values are reduced.

4.3.2. Clinically interesting regions

Due to the spatial and temporal continuity of the blood flow, the IDC at a stenosis is similar to the IDC before and after the obstructed vessel segment (Fig. 13). Thus, the IDC at a stenosis can be estimated with a weighted sum of pre- and post-stenosis IDC.

The continuity of the contrast agent spreading in the blood flow is not guaranteed in large-volume broadenings, such as huge aneurysms. In these places, the STF filtering may not improve the visibility of the opacified vasculature.

4.3.3. Evaluation on clinical images

We present the results of enhancement of three clinical XA image sequences from different parts of the body, containing different types of motion and blood-flow dynamics. All of them were acquired by a General Electric AdvantX XA system. The *abdomen* sequence (Fig. 15) is composed of 16 frames showing the abdominal aorta. The two other sequences are peripheral: *arm* (Fig. 14) is a 11-frame sequence of an arm with respiratory and intestinal motion; *hand* is a 16-frame sequence of a hand with no significant motion. We did not perform motion compensation to test the effectiveness of suppression of motion artifacts as well as of noise and other artifacts.

Figs. 14–16 show comparative results after processing the images with subtraction and STF ($\gamma=0.15\,\mathrm{cm}^{-1}$,

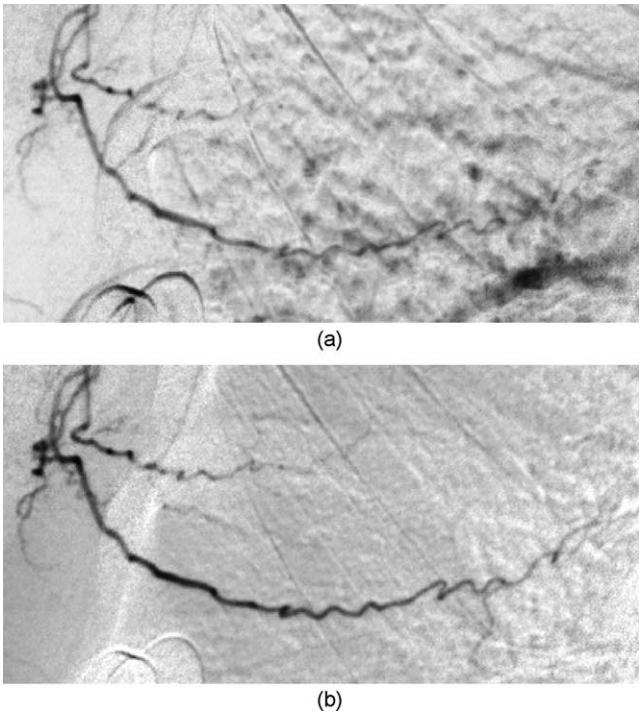

Fig. 14. Comparison of improved vessel visibility in the *arm* image using subtraction (a) and STF (b). STF suppressed the moving background structures of ribs and soft tissues significantly better than the subtraction method. (a) Subtraction (averaging frames 1–3 for mask and maximum opacification in frames 4–11); (b) STF.

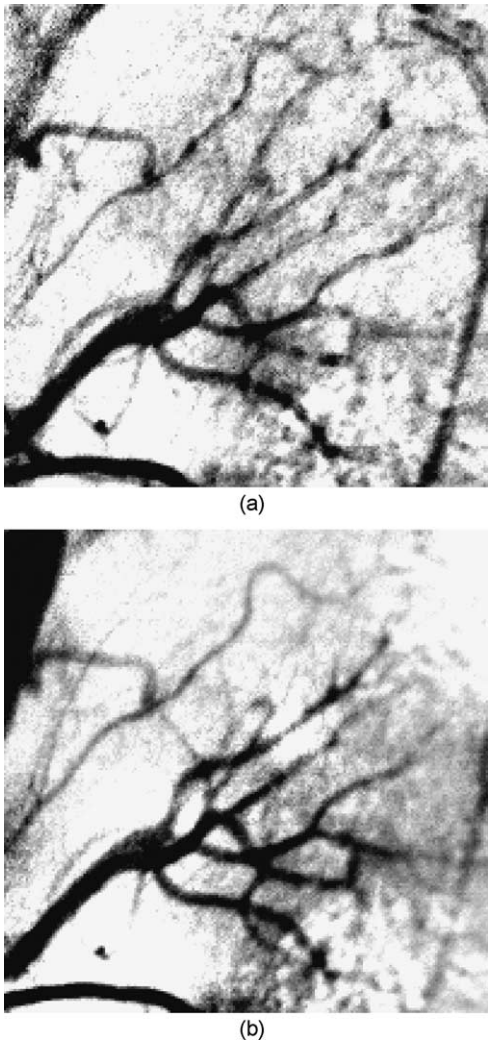

Fig. 15. Comparison of improved vessel visibility in the *abdomen* image using subtraction and STF with RBF-linear kernel: (a) subtraction (averaging frames 3–5 for mask and maximum opacification of frames 6–12); (b) STF.

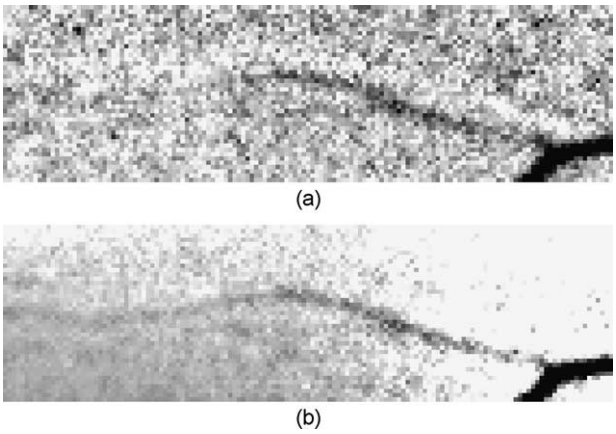

Fig. 16. Comparison of improved vessel visibility in the *hand* image using subtraction and STF with RBF-linear kernel: (a) subtraction (averaging frames 1–3 for mask and maximum opacification of frames 4–16); (b) STF.

C = default). STF gives good results, even when image registration and subtraction cannot fully eliminate complex motion artifacts (e.g. Fig. 16 shows some suppression of moving ribs and soft tissues using subtraction and STF processing). STF also reduces artifacts resulting from inaccurate registration, as the motion causing such artifacts differs from the motion of the blood flow. Therefore, the STF vessel enhancement can be used to further enhance vessel visibility in a motion-compensated image sequence.

5. Conclusions and future work

This paper presented a method, STF, using Support Vector Machines for vessel enhancement in XA sequences. Like others, the method exploits a temporal model of the IDC at vessel pixels, but, unlike any other method known to us, it *learns* the model from real sequences. In this sense, the method can be regarded as matched linear filtering optimizing the CNR, in which learning takes care automatically of modeling difficulties. It also provides a unified framework for incorporating constraints and prior knowledge.

Results show that STF can learn IDCs at vessel pixels which can be successfully used to ameliorate the CNR. A simple, automatic training point selection algorithm is sufficient to select pixels providing positive and negative (vessel and non-vessel) training IDCs. The improvement in vessel visibility brought about by STF is comparable with that of the theoretically optimal matched filtering.

Results with simulated data show that STF filtering gives superior vessel visibility compared to subtraction, and in the presence of structured noise significantly better results than AMF. With clinical sequences, STF performed very well where the commonly used subtraction techniques with image registration do not give good results, e.g. complex motion or motion of soft tissues; and effectively reduced artifacts introduced by inaccurate registration.

Future work is needed on learning different IDCs, incorporating spatial vessel enhancement methods into the STF framework, and extending it to tackle currently unsolved registration problems.

6. Summary

In this paper we present a vessel enhancement method, SVM temporal filtering (STF), for X-ray angiographic (XA) images based on statistical learning. We use a Support Vector Machine (SVM) to learn the characteristics of the evolution of grey levels in time caused by the contrast media injection.

We briefly review currently known the image stacking methods and the basics of SVMs and we describe how the SVM can be used for vessel enhancement. First we show that a simple linear SVM applied to vessel enhancement can be regarded as a matched linear filter optimizing the contrast-to-noise ratio in XA images. Then we propose a modified non-linear kernel function—combining a radial basis function kernel with a linear kernel—for the SVM leading to a better enhancement with noisy, varying grey-level dynamics at vessel pixels.

We give a simple method for automatically obtain training samples from clinical images. Positive training samples are randomly selected from the detected vessel centerline points. A simple method is used for extracting vessel centerline points: we compose a single image from those pixels where pixelwise, frame-to-frame differences are above a given threshold then apply morphological open and skeletonization operations. Negative samples are generated synthetically from Poisson noise.

We determine the optimal kernel and learning parameters of the SVM from simulated data. Then we compare the performance of subtraction and other image stacking methods (matched filtering, approximate matched filtering, maximum opacification, recursive temporal filtering) to the STF on simulated data and clinical images. One key advantage over the well-known matched filters is that an optimal filter is learnt from image sequences, not estimated at design stage. Results on clinical XA images show that learning-based enhancement achieves better results compared to the tested image stacking methods.

Acknowledgements

Thanks go to Yves Trousset for making the images available and for his input on the XA application; to László Vajta and Robert Heidsieck for making possible the first author's visit to Heriot Watt University, where this work was carried out; and to Francesca Odone for helpful discussions on Support Vector Machines.

References

[1] Cowen AR. Digital X-ray imaging. IEE Colloquium Med Imaging: Image Process Anal 1992;8/1–8/3.

[2] Krestel E. Imaging systems for medical diagnostics. New York: Wiley; 1990.

[3] Meijering EHW, Niessen WJ, Viergever MA. Retrospective motion correction in digital subtraction angiography: a review. IEEE Trans Med Imaging 1999;18(1):2–21.

[4] Kruger RA, Liu P, Bateman W, Nelson JA. Time domain filtering using computerized fluoroscopy—intravenous angiography applications. SPIE 1981;314:319–26.

[5] Kruger RA, Liu Y. Digital angiography using a matched filter. IEEE Trans Med Imaging 1982;1(1):16–21.

[6] Kump KS, Saidel GM, Wilson DL. Comparison of algorithms for combining X-ray angiography images. IEEE Trans Med Imaging 2001;20(8):742–50.

[7] Kruger RA, Mistretta CA, Houk TL, Kubal W, Riederer SJ, Ergun DL, et al. Computerized fluoroscopy techniques for intravenous study of cardiac chamber dynamics. Invest Radiol 1979;14(4): 279–87.

[8] Liu P, Kruger RA, Nelson JA, Miller FJ, Osborn AG, Wojtowycz M. Digital angiography: matched filtration versus mask-mode subtraction. Radiology 1985;154(1):217–20.

[9] Riederer SJ, Enzmann DR, Hall AL, Pelc NJ, Djang WT. The application of matched filtering to X-ray exposure reduction in digital subtraction angiography: clinical results. Radiology 1983;146: 349–54.

[10] Pontil M, Verri A. Support vector machines for 3D object recognition. IEEE Trans PAMI 1998;20(6):637–46.

[11] Joachims T. Making large-scale SVM learning practical. In: Schölkopf B, Burges C, Smola A, editors. Advances in kernel methods—support vector learning. Cambridge (MA): MIT Press; 1999. Available: http://www-ai.cs.uni-dortmund.de/DOKUMENTE/joachims_99a.pdf.

[12] Furey T, Cristianini N, Duffy N, Bednarski D, Schummer M, Haussler D. Support vector machine classification and validation of cancer tissue samples using microarray expression data. Bioinformatics 2000;16:906–14.

[13] Vapnik V. The nature of statistical learning theory. New York: Springer; 1995.

[14] Vapnik V. Statistical learning theory. New York: Wiley; 1998.

[15] Burges C. A tutorial on support vector machines for pattern recognition. Data Mining Knowl Discov 1998;2(2):121–67.

[16] Cristianini N, Shawe-Taylor J. An introduction to support vector machines and other kernel-based learning methods. Cambridge, UK: Cambridge University Press; 2000.

[17] Scholkopf B, Smola A, Williamson R, Bartlett P. New support vector algorithms. Neural Comput 2000;12:1207–45.

[18] Poli R, Valli G. An algorithm for real-time vessel enhancement and enhancement. Comput Methods Programs Biomed 1996;52:1–22.

[19] Santos AC, Furie SS, Gutierrez MA. Estimation of coronary blood flow by contrast propagation using simulated X-ray angiography. IEEE Comput Cardiol 1999;26:379–82.

[20] Chang CC, Lin CJ. LIBSVM: a library for support vector machines 2001. Available: http://www.csie.ntu.edu.tw/~cjlin/libsvm.

András Lassó received MSc degree in Electrical Engineering in 2000 and currently he is a PhD candidate at Budapest University of Technology and Economics, Hungary. Since 2000 he is working for GE Healthcare, developing medical applications for digital X-ray angiography systems. His research interest is medical image enhancement and segmentation and X-ray image synthesis.

Emanuele Trucco is a Reader (Associate Professor) in the School of Engineering and Physical Sciences, Heriot Watt University, Edinburgh (UK). He graduated in 1984 and obtained his PhD in 1990, both from the University of Genoa, in Italy. He has been an active researcher in computer vision and image processing since 1984, publishing more than 150 refereed papers, co-authoring two books, and organizing several international events in various countries. He is an Honorary Editor of the IEE Proceedings on Vision, Image and Signal Processing, and an Editor of the IEEE Transactions on Systems, Man and Cybernetics Part C. His interests span a variety of applications, including manufacturing, robotics, communications and medical image analysis.