Automatic Online Layer Separation for Vessel Enhancement in X-ray Angiograms for Percutaneous Coronary Interventions[†,††]

Hua Ma[a], Ayla Hoogendoorn[b], Evelyn Regar[c,d], Wiro J. Niessen[a,e], Theo van Walsum[a]

[a]*Biomedical Imaging Group Rotterdam, Erasmus MC, University Medical Center Rotterdam, The Netherlands*
[b]*Department of Cardiology, Biomedical Engineering, Erasmus MC, University Medical Center Rotterdam, The Netherlands*
[c]*Department of Cardiology, Erasmus MC, University Medical Center Rotterdam, The Netherlands*
[d]*Heart Center, University Hospital Zurich, Switzerland*
[e]*Quantitative Imaging Group, Faculty of Applied Sciences, Delft University of Technology, The Netherlands*

Abstract

Percutaneous coronary intervention is a minimally invasive procedure that is usually performed under image guidance using X-ray angiograms in which coronary arteries are opacified with contrast agent. In X-ray images, 3D objects are projected on a 2D plane, generating semi-transparent layers that overlap each other. The overlapping of structures makes robust automatic information processing of the X-ray images, such as vessel extraction which is highly relevant to support smart image guidance, challenging. In this paper, we propose an automatic online layer separation approach that robustly separates interventional X-ray angiograms into three layers: a breathing layer, a quasi-static layer and a vessel layer that contains information of coronary arteries and medical instruments. The method uses morphological closing and an online robust PCA algorithm to separate the three layers. The proposed layer separation method ran fast and was demonstrated to significantly improve the vessel visibility in clinical X-ray images and showed better performance than other related online or prospective approaches. The potential of the proposed approach was demonstrated by enhancing contrast of vessels in X-ray images with low vessel contrast, which would facilitate the use of reduced amount of contrast agent to prevent contrast-induced side effects.

Keywords:
Online Robust PCA, Layer Separation, Vessel Enhancement, X-ray Angiograms

1. Introduction

1.1. Motivation

Percutaneous coronary intervention (PCI) is a minimally invasive procedure for patients with advanced coronary artery disease. In this procedure, a stent pre-mounted on a delivery catheter is advanced over a guide-wire and through a guiding catheter at the site of narrowing in a patient's coronary arteries. Once the lesion site is reached, the delivery balloon is inflated and the stent is deployed against the coronary wall, assuring optimal patency of the artery. As there is no direct eyesight on the target area, these procedures are commonly performed under image guidance using X-ray angiography (XA), where coronary arteries are visualized with X-ray contrast agent. During the intervention, clinicians use XA images to navigate catheters and guidewires inside the patients.

As XA images contain useful information on anatomy and instrument position, many works have been published on extracting relevant information to improve the image guidance for cardiac interven-

[†]This manuscript has been accepted for publication in *Medical Image Analysis*. For the published journal article, please refer to https://doi.org/10.1016/j.media.2017.04.011

tions. For example, Panayiotou et al. (2014) have developed a retrospective motion gating technique of interventional X-ray images through vessel extraction. Also using the information of vessels, pre/intra-operative information fusion between CT angiography and XA have been reported in Baka et al. (2013) and Rivest-Henault et al. (2012). Apart from vessels in XA, there is interest to track structures such as the lungs, catheters and guidewires. Shechter et al. (2005) have used the position of diaphragm as an indicator of respiratory phase and constructed a patient specific coronary motion model based on that. In Baka et al. (2015), the position of guiding catheter tip has been related to the combination of respiratory and cardiac motion.

Since X-ray images are projections of 3D structures on a 2D plane, the image content can be interpreted as a composition of several opaque or semi-transparent structures, which have different appearances and motion patterns. The overlapping nature of the structures makes automatic analysis of XA challenging. Separating the structures from each other enables visualizing and analyzing different structures independently, which would, therefore, potentially facilitate the information processing of XA. For example, vessel extraction using Hessian-based filtering method in XA is often hampered by non-vascular structures, such as guiding catheters, diaphragm border and vertebral body edges, because of their tubular or curvi-linear appearance in XA. Separating non-vascular structures would improve the visibility of vessels and promote automatic vessel extraction that would ultimately facilitate the image guidance during interventions.

In the context of this work, we interpret the process of separating those structures in XA images as a separation of a set of additive 2D layer images which add up to the original image, and each of them has different structures. The purpose of this work is to develop and evaluate a fast method that can run prospectively for the effective and efficient separation of the structures on different layers for XA sequences. Following the terminology from earlier works (in Section 1.2), we adopt the term "layer separation" to refer to the separation of structures and putting them in different layers.

1.2. Related works

Existing methods for layer separation for X-ray fluoroscopic sequences can be categorized into two approaches: motion-based and motion-free .

Motion-based layer separation methods treat each frame of an X-ray fluoroscopic sequence as the outcome of the motion of each layer. Hence, the key part of obtaining the layers in these methods is estimating the motion of every layer. Various assumptions on the type of motion have been proposed. For instance, Close et al. (2001) have estimated translation, rotation and scaling for each layer in a region of interest. The layers are computed by transforming each frame with the estimated motion and averaging the transformed frames. This method computed a total of four layers for a sequence. Zhu et al. (2009) have proposed a two-layer separation scheme. They have developed a Bayesian framework that combines dense motion estimation, uncertainty propagation and statistical fusion to achieve layer separation. In a three-layer separation approach proposed by Zhang et al. (2009), a multi-scale framework has been developed based on different motion patterns for the static background, lung and vessels. In this work, a dense motion field of each layer has been constructed using thin plate splines. Fischer et al. (2015) have further extended this method by introducing a regularization term for layers with a Bayesian model to aid layer separation. In particular, they have proposed to use a robust data term and edge-preserving regularization. In the work of Auvray et al. (2009), a joint layer segmentation and parametric motion estimation scheme has been proposed for transparent image sequences. Similarly, Preston et al. (2013) jointly estimated layers and their corresponding smooth deformation to model the non-smooth motion observed in a fluoroscopic sequence. A total variation based regularization was used to encourage sparsity of gradients within and across the layer images.

Unlike motion-based methods, motion-free approaches do not require estimating the motion of layers. Instead, they directly model the background layer or/and foreground (vessel) layer of an image sequence under certain hypotheses. One of the simplest ways of modeling the background of XA is computing the median of several frames in a sequence, and obtaining the foreground by subtract-

ing the median image from the original frames (see Baka et al. (2014)). This method worked well for the background that is entirely static, but generates artefacts when there are moving objects in the background, e.g. diaphragm in XA. A more advanced method has been proposed by Tang et al. (2012) in which they assumed that the vessel and the backgrounds generate independent signals that are mixed in a sequence, so that the vessel-background separation becomes a blind source separation problem that is commonly solved by independent component analysis (ICA, Hyvärinen et al. (2004)).

Apart from ICA, robust principal component analysis (RPCA) is also a common approach for source decomposition. One of the most popular RPCA methods, principal component pursuit (PCP, Candès et al. (2011)), splits a data matrix into a low-rank component and a sparse component. It has been used for background modeling or foreground detection for surveillance videos (Bouwmans and Zahzah (2014)). In the field of medical image analysis, it found applications in reconstruction (Otazo et al. (2015)) and motion correction (Hamy et al. (2014)) in dynamic MRI. On the topic of layer separation for X-ray images, Ma et al. (2015) have used morphological closing to remove breathing structures from the images and adopted RPCA to separate a quasi-static layer and a vessel layer from XA. This method could only be used in a retrospective setting, since it requires all frames of a sequence. Volpi et al. (2015) have developed a method that worked in a prospective setting. The method used vesselness filtering (Frangi et al. (1998)) and RPCA to separate a foreground that contains interventional devices. They have implemented the foreground separation by solving RPCA with a mini-batch of data: for each new coming mini-batch, the average of the low-rank component was estimated and used as the background for the next mini-batch. The limitation of this method is that the foreground separation of a mini-batch is delayed by the processing of the previous complete block of data.

Online robust PCA (OR-PCA) is an online extension of the original RPCA method, proposed by Feng et al. (2013). OR-PCA overcomes the limitation of RPCA-based methods by reformulating the nuclear norm in the RPCA formulation as an explicit low-rank factorization, so that it does not require to "see" the complete dataset or a mini-batch of data, but can process each single data sample one at a time. This setting enables online processing of streaming data. In Song et al. (2015), a closed-form solution for the subspace basis update in OR-PCA has been proposed and shown to achieve better performance in image alignment tasks. OR-PCA has been used in computer vision tasks, such as background subtraction (Javed et al. (2015)) and foreground detection (Javed et al. (2014)), but its application in the field of medical imaging has not been investigated yet.

1.3. Overview and contributions

In this work we extended the method in Ma et al. (2015) that only worked in a retrospective or "offline" setting. To this end, we developed and evaluated an automatic motion-free online layer separation method for X-ray angiograms. The method robustly separates the layer that contains vessels and catheter tip from a (quasi) static background, while ignoring large-scale motion such as diaphragm movement. Our contributions are:

(a) We integrated OR-PCA in the layer separation scheme, enabling online layer separation for XA, which is a key ingredient for its potential application in a clinical workflow.

(b) Inspired by the work of Mairal et al. (2010), we proposed and analyzed three ways to down-weight past information that is able to improve the layer separation performance using the original OR-PCA algorithm.

(c) We compared the proposed method with other related background-removal approaches and evaluated the results visually and quantitatively on real patient XA data.

(d) We investigated the potential of improving the contrast of vessels in a low-contrast scenario using the proposed method with synthetic low-contrast XA sequences and real sequences acquired in a pig experiment in which various contrast levels were used.

2. Method

2.1. Overview

The proposed method treats the intensity of an XA frame as the sum of three layers, i.e., a "breathing" layer, a quasi-static layer and a layer that contains vessels. The method consists of two main steps: first, large-scale breathing structures, e.g. diaphragm, are separated and removed from the original XA frame, and second, smaller moving structures, e.g. vessels and guiding catheters, are separated from a quasi-static background using online robust PCA (OR-PCA). Fig. 1 provides an overview of the complete method, details are described in the remainder of Section 2.

2.2. Separation of breathing structures

To prevent artefacts remaining in the vessel layer due to respiratory motion, the layer that contains large-scale breathing structures, such as diaphragm, is removed from the original XA images in the first step.

The layer of breathing structures was obtained by removing "small" objects from the original X-ray angiographic frame. Depending on the field of view, those objects could include vessels, guiding catheters, guide wires, stitches and vertebral bodies. Following the approach of Ma et al. (2015), as a preprocessing step, we applied a morphological closing operation to the XA image with a circular structuring element of 8.5 *mm* in diameter, in order to remove any tubular and curvilinear structures smaller than that size. An example of a resulting image is shown in Figure 2b, where the guiding catheter and vessels are removed and vertebral contours are blurred, while structures that are susceptible to breathing motion remain in the image (diaphragm and lung tissue are shown as the white area in the upper left part of the image). The resulting image is referred to as the "breathing layer" in this paper and was next subtracted from the original image to obtain the difference image (DI, Figure 2c) of the XA frame for further processing.

2.3. Separation of vessel layer via OR-PCA

In this section, we briefly review the formulation of the online robust PCA method proposed by Feng et al. (2013) and different subspace basis update strategies for solving the OR-PCA problem (Feng et al. (2013), Song et al. (2015)). Then we propose three different ways of coping with previous frames to improve on these methods.

2.3.1. Notation

Bold letters are used to denote vectors. With the difference image (DI) of an XA frame represented with a $k \times k$ matrix, we concatenated all pixels in this matrix to form a single column vector $\mathbf{z} \in \mathbb{R}^p$, where $p = k^2$ is the dimension of the observed sample. Likewise, we use $\mathbf{x} \in \mathbb{R}^p$ to denote the quasi-static background of the XA frame and $\mathbf{e} \in \mathbb{R}^p$ represents the foreground. Hence, $\mathbf{z} = \mathbf{x} + \mathbf{e}$. Let n denote the number of frames in a sequence, t be the index of the sample/time instance of a frame and r denote the intrinsic dimension of the subspace underlying $\{\mathbf{x}_i \,|\, i = 1, 2, \ldots n\}$.

Matrices are denoted by capital letters in the following sections. In particular, $Z \in \mathbb{R}^{p \times n}$ is the matrix of a complete sequence of difference images (DIs), where its column $\mathbf{z}_i$ represents the i-th DI. Likewise, X and E are the background and the foreground matrices with $\mathbf{x}_i$ and $\mathbf{e}_i$ the vector for the i-th background and the i-th foreground. For an arbitrary real matrix M, let $\|M\|_1 = \sum_{i,j} |M_{i,j}|$ denote the L_1-norm of M, $\|M\|_F$ denotes the Frobenius norm $\|M\|_F = \sqrt{\sum_{i,j} |M_{i,j}|^2}$, and $\|M\|_* = \sum_i \sigma_i(M)$ denotes the nuclear norm, i.e., the sum of its singular values. $Tr(M)$ denotes the trace of a matrix.

2.3.2. Online robust PCA

Robust PCA (RPCA) aims at estimating the subspace underlying the observed samples. Among many popular RPCA methods, Principal Component Pursuit (PCP, Candès et al. (2011)) has been proposed to solve the RPCA problem by approximating the data matrix as the sum of a low-rank matrix and a sparse matrix. The concepts of low-rank and sparsity have been implemented using the nuclear norm and the L_1-norm of matrix respectively. This formulation is suitable for the separation of the vessel layer from the DI of an XA frame, since the background has merely minor changes, which can be modeled as a low-rank matrix. In addition, the fact that vessels and guiding catheters take up only a small portion of

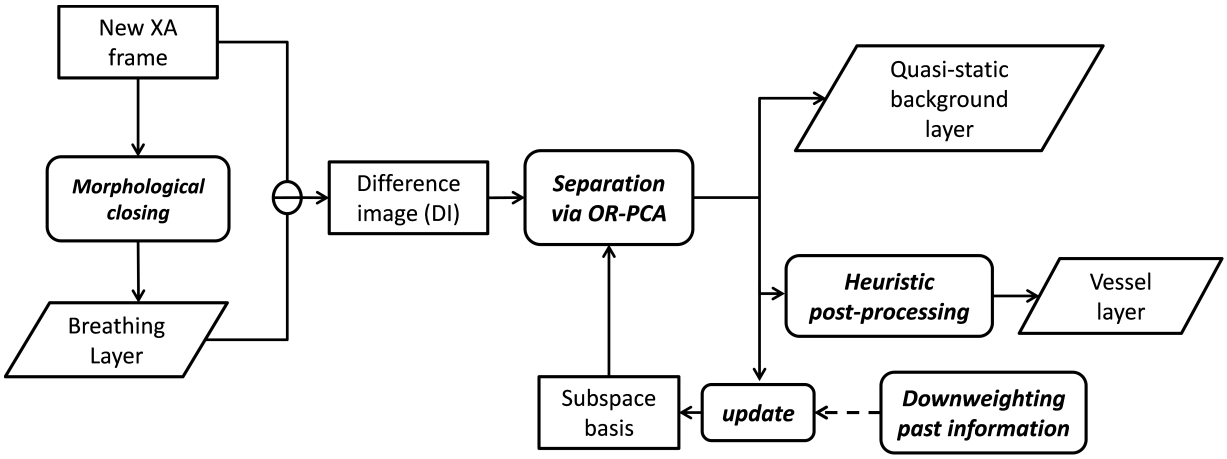

Figure 1: The overview of online layer separation for an XA frame.

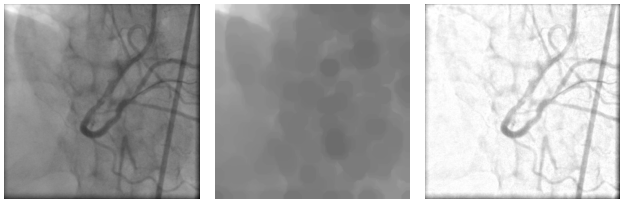

(a) (b) (c)

Figure 2: Morphological closing operation applied on an XA frame: (a) original frame, (b) image processed with morphological closing, (c) difference image (DI) (a-b).

the complete image content fits the requirement of sparsity.

2.3.2.1 The OR-PCA formulation

Different from the classical formulation in Candès et al. (2011), PCP can be reformulated as follows (Feng et al. (2013)):

$$\min_{X,E} \frac{1}{2}\|Z - X - E\|_F^2 + \lambda_1\|X\|_* + \lambda_2\|E\|_1 \qquad (1)$$

where λ_1 and λ_2 are regularization coefficients. Through minimizing the cost function (1) that contains the nuclear norm of the background X and the L_1-norm (sparsity) of the foreground E, the RPCA algorithm aims at obtaining the background (X) and foreground (E) that best approximate the XA sequence (Z). Because the nuclear norm couples all samples tightly, typical methods to solve Equation (1), such as Augmented Lagrangian Multiplier (ALM, Lin et al. (2011)), are often implemented in a batch manner, which limits its application in scenarios that deal with streaming data, e.g. X-ray cine angiography data during coronary interventions.

To overcome this problem, Feng et al. (2013) have proposed to use an equivalent form of the nuclear norm:

$$\|X\|_* = \inf_{L,R}\left\{\frac{1}{2}\|L\|_F^2 + \frac{1}{2}\|R\|_F^2 : X = LR^T\right\} \qquad (2)$$

where *inf* denotes the greatest lower bound of a subset of a partially ordered set, $L \in \mathbb{R}^{p\times r}$ is the basis of the low-dimensional subspace and $R \in \mathbb{R}^{n\times r}$ can be seen as the samples' coefficient with respect to the basis. Substituting Equation (2) into (1), the RPCA problem can be reformulated as (3):

$$\min_{L,R,E} \frac{1}{2}\|Z - LR^T - E\|_F^2 + \frac{\lambda_1}{2}(\|L\|_F^2 + \|R\|_F^2) + \lambda_2\|E\|_1 \qquad (3)$$

Following Feng et al. (2013), solving Equation (3) is equivalent to minimizing the following empirical cost function given a sequence Z consisting of n samples $[\mathbf{z}_1 \ldots \mathbf{z}_n]$:

$$f_n(L) \triangleq \frac{1}{n}\sum_{i=1}^{n} l(\mathbf{z}_i, L) + \frac{\lambda_1}{2n}\|L\|_F^2 \qquad (4)$$

where the loss function $l(\mathbf{z}, L)$ for each sample is defined as:

$$l(\mathbf{z}, L) \triangleq \min_{\mathbf{r},\mathbf{e}} \frac{1}{2}\|\mathbf{z}_i - L\mathbf{r} - \mathbf{e}\|_2^2 + \frac{\lambda_1}{2}\|\mathbf{r}\|_2^2 + \lambda_2\|\mathbf{e}\|_1 \qquad (5)$$

Note that Equation (4) enables the possibility of updating the basis L based on each individual sample. To handle streaming data in practice, in Feng et al. (2013), the estimation of basis L_t is obtained through minimizing the following surrogate function of (4) with respect to L for the t-th time instance:

$$g_t(L) \triangleq \frac{1}{t}\sum_{i=1}^{t}(\frac{1}{2}\|\mathbf{z}_i - L\mathbf{r}_i - \mathbf{e}_i\|_2^2 + \frac{\lambda_1}{2}\|\mathbf{r}_i\|_2^2 + \lambda_2\|\mathbf{e}_i\|_1) + \frac{\lambda_1}{2t}\|L\|_F^2 \quad (6)$$

Also observe that the loss function (5) optimizes $\mathbf{r}$ (the coefficient of $\mathbf{z}_i$ on the basis L) and $\mathbf{e}$ (the sparse component of $\mathbf{z}_i$) to minimize the cost given a fixed basis. Through an alternating optimization of $\mathbf{r}$, $\mathbf{e}$ and L, Equation (4) can be solved in an online manner. The complete stochastic optimization scheme for solving the OR-PCA problem is described in Algorithm 1.

Note that the right-hand side of Equation (7) in Algorithm 1 is equivalent to the loss function (5) for the t-th sample. To solve it, Feng et al. (2013) gives a closed-form solution to alternatively update $\mathbf{r}$ and $\mathbf{e}$ until a convergence criterion is met. The update of L_t in Equation (8) is discussed in the next section.

2.3.2.2 *Update the subspace basis L_t*

To minimize the function (6) with respect to L, note that the term $\frac{\lambda_1}{2}\|\mathbf{r}_i\|_2^2$ and $\lambda_2\|\mathbf{e}_i\|_1$ can be discarded, we then derived the following expression for L_t from (6):

$$L_t \triangleq \underset{L}{argmin}\frac{1}{2}Tr[L^T L\sum_{i=1}^{t}(\mathbf{r}_i\mathbf{r}_i^T + \frac{\lambda_1}{t}I)] - Tr(L^T\sum_{i=1}^{t}((\mathbf{z}_i - \mathbf{e}_i)\mathbf{r}_i^T)) \quad (9)$$

Using the two intermediate variables A_t and B_t that accumulate information of past frames, Equation (9) is equivalent to (8) in Algorithm 1. Equation (8) is then solved by the block-coordinate descent method, i.e., each column of the basis L is updated sequentially while fixing the other columns (see Algorithm 2).

Algorithm 1 Stochastic optimization for OR-PCA (Feng et al. (2013))

Require: $\{\mathbf{z}_1, \ldots, \mathbf{z}_T\}$ (sequentially revealed data samples), $\lambda_1, \lambda_2 \in \mathbb{R}$ (regularization parameters), $L_0 \in \mathbb{R}^{p\times r}$, $\mathbf{r}_0 \in \mathbb{R}^r$, $\mathbf{e}_0 \in \mathbb{R}^p$, $A_0 = \mathbf{0}^{r\times r}$, $B_0 = \mathbf{0}^{p\times r}$ (initial solution), T (number of samples).

1: **for** t = 1 **to** T **do**
2: Reveal the sample $\mathbf{z}_t$.
3: Given L_{t-1}, project the new sample:

$$\{\mathbf{r}_t, \mathbf{e}_t\} = \underset{\mathbf{r},\mathbf{e}}{argmin}\frac{1}{2}\|\mathbf{z}_t - L_{t-1}\mathbf{r} - \mathbf{e}\|_2^2 + \frac{\lambda_1}{2}\|\mathbf{r}\|_2^2 + \lambda_2\|\mathbf{e}\|_1 \quad (7)$$

4: $A_t \leftarrow A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T$, $B_t \leftarrow B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T$
5: Update the basis L_t

$$L_t \triangleq \underset{L}{argmin}\frac{1}{2}Tr[L^T L(A_t + \lambda_1 I)] - Tr(L^T B_t) \quad (8)$$

6: **end for**
7: **return** $X_T = L_T R_T^T$ (the low-rank matrix), E_T (the sparse matrix).

Algorithm 2 The basis update using block-coordinate descent (Feng et al. (2013))

Require: $L = [\mathbf{l}_1, \ldots, \mathbf{l}_r] \in \mathbb{R}^{p\times r}$, $A = [\mathbf{a}_1, \ldots, \mathbf{a}_r] \in \mathbb{R}^{r\times r}$, $B = [\mathbf{b}_1, \ldots, \mathbf{b}_r] \in \mathbb{R}^{p\times r}$. $\tilde{A} \leftarrow A + \lambda_1 I$.

1: **for** j = 1 **to** r **do**
2: Update the j-th column of L.

$$\mathbf{l}_j \leftarrow \frac{1}{\tilde{A}_{j,j}}(\mathbf{b}_j - L\tilde{\mathbf{a}}_j) + \mathbf{l}_j. \quad (10)$$

3: **end for**
4: **return** L.

Another way of solving Equation (8) is to derive a closed-form solution. Let the derivative of the right-hand side of (8) with respect to L be zero, we obtain

$$\frac{1}{2}L(A_t + \lambda_1 I)^T + \frac{1}{2}L(A_t + \lambda_1 I) - B_t = 0 \qquad (11)$$

where $A_t = A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T$, $B_t = B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T$. As $(A_t + \lambda_1 I)$ is symmetrical, a simple closed-form solution of (8) can be derived as

$$L_t = B_t(A_t + \lambda_1 I)^{-1}. \qquad (12)$$

This is equivalent to the form given in Song et al. (2015).

2.3.3. Downweighting the past information

The previous solutions for the subspace basis update treat all samples equally, which works well for scenarios where samples are independently drawn. For stream video data, however, adjacent frames have higher correlation than "distant" frames. Thus, it may be possible to improve the basis update by treating past frames with different weights, giving close-by frames higher impact to the result than the distant frames. Inspired by the work of Mairal et al. (2010) who has reported several possibilities to handle past data in an online dictionary learning problem, we propose three approaches to downweight past information for the OR-PCA algorithm. In Algorithm 1, as A_t and B_t contain information of past frames, variations can be made to replace the following equation set on line 4 in Algorithm 1:

$$\begin{aligned} A_t &\leftarrow A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T \\ B_t &\leftarrow B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T \end{aligned} \qquad (13)$$

A logical choice is to apply an exponential decay (ED) to "forget" past information as in (14):

$$\begin{aligned} A_t &\leftarrow (1-\epsilon)A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T \\ B_t &\leftarrow (1-\epsilon)B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T \end{aligned} \qquad (14)$$

where ϵ is the decay rate and $0 < \epsilon < 1$. So for the t-th time instance, the weight for the i-th sample is $(1-\epsilon)^{t-i}$.

Similar to Mairal et al. (2010), as a second option we consider supra-linear decay (SLD) approach:

$$\begin{aligned} A_t &\leftarrow \left(1 - \tfrac{1}{t}\right)^{\rho} A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T \\ B_t &\leftarrow \left(1 - \tfrac{1}{t}\right)^{\rho} B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T \end{aligned} \qquad (15)$$

where ρ is a tunable decay parameter and $\rho \geqslant 0$. At the t-th time instance, the weight for the i-th sample becomes $\left(\frac{i}{t}\right)^{\rho}$. Note that: when $\rho = 0$, (15) turns into (13); when $\rho = 1$, (15) degrades to a linear decay.

Apart from ED and SLD that scale the past data, it is also an option to focus only on adjacent frames in a fixed-size window, so that frames within the sliding window are treated equally, whereas the frames outside the window from the earlier times are not considered for the basis update, as follows:

$$\begin{cases} \begin{aligned} A_t &\leftarrow \mathbf{r}_t\mathbf{r}_t^T \\ B_t &\leftarrow (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T \end{aligned}, & t_0 = 1 \\ \begin{aligned} A_t &\leftarrow A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T \\ B_t &\leftarrow B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T \end{aligned}, & t_0 > 1 \text{ and } t \leqslant t_0 \\ \begin{aligned} A_t &\leftarrow A_{t-1} + \mathbf{r}_t\mathbf{r}_t^T - \mathbf{r}_{t-t_0}\mathbf{r}_{t-t_0}^T \\ B_t &\leftarrow B_{t-1} + (\mathbf{z}_t - \mathbf{e}_t)\mathbf{r}_t^T - (\mathbf{z}_{t-t_0} - \mathbf{e}_{t-t_0})\mathbf{r}_{t-t_0}^T \end{aligned}, & \text{else} \end{cases} \qquad (16)$$

where t_0 is the window size (number of frames within the window). This approach is referred to as "sliding-window (SW)".

2.4. Summary

The proposed online layer separation method consists of the following steps, as shown in Figure 1.

(a) **Breathing layer separation** When a new XA frame is obtained, the breathing layer is firstly extracted by applying morphological closing on that frame, as described in Section 2.2. Subsequently, the breathing layer is subtracted from the original frame to obtain the DI.

(b) **Quasi-static layer and vessel layer separation** Transform the DI from a matrix to vector by concatenating each column of the matrix one after another. This vector is then separated into two components by the OR-PCA method, as described in Section 2.3. The sparse component is reshaped to form the vessel layer, the other component is constructed as the quasi-static layer.

Finally, as pixels belonging to contrast agent always have negative value in the vessel layer, pixels with positive value in the vessel layer are heuristically set to zero to suppress artefacts.

3. Experiments

3.1. Image data

In this work, we used three types of data for evaluation: clinical X-ray angiograms, synthetic low-contrast XA and X-ray angiographic data of pigs with variations in contrast concentration.

3.1.1. Clinical X-ray angiographic data

Imaging data from clinical routine that were anonymized were used for our experiments. The data were acquired under standard clinical protocol from the Department of Cardiology at Erasmus MC in Rotterdam, the Netherlands. The 42 XA sequences are from 21 patients who underwent a PCI procedure and were acquired with Siemens AXIOM-Artis biplane system. The frame rate of all sequences is 15 frames per second (fps). The number of frames per sequence varies from 46 to 244. All 42 XA sequences have in total 4886 frames. 22 sequences have 512×512 pixels, 12 have 600×600 pixels, 2 have 776×776 and 6 have 1024×1024. Their corresponding pixel sizes are 0.216×0.216 or 0.279×0.279, 0.184×0.184, 0.184×0.184 and $0.139 \times 0.139\, mm^2$, respectively. In all sequences, inflow and wash-out of contrast agent can be observed.

3.1.2. Synthetic low-contrast XA

The synthetic image data was used to simulate the condition that a reduced amount (50%) of contrast agent is administered, for the purpose of testing our online layer separation method on low-contrast XA. To create these synthetic XA sequences from the real ones, we used the off-line layer separation method in Ma et al. (2015). The idea is that the real clinical XA sequence was firstly separated into three layers. The intensity of the vessel layer was then halved and added back to the other two layers to generate a new XA sequence that has half the amount of intensity compared to the original one, as shown in Equation (17):

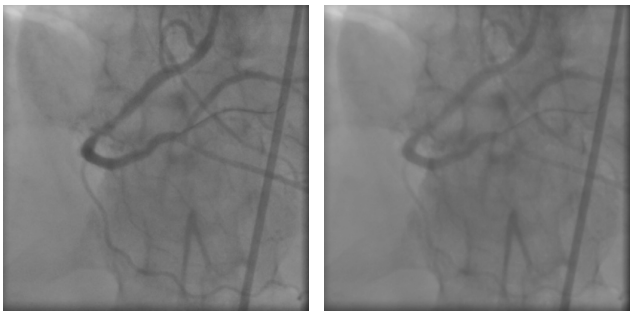

(a) Clinical XA (b) Synthetic XA

Figure 3: An example frame of real clinical XA sequences and synthetic low-contrast XA sequences: (a) the real image, (b) the synthetic XA frame with 50% vessel contrast.

$$I_{synthetic} = \alpha\, I^*_{vessel} + I^*_{static} + I^*_{breathing} \qquad (17)$$

where $I_{synthetic}$ denotes the synthetic XA sequence, I^*_{vessel}, I^*_{static} and $I^*_{breathing}$ are the vessel layer, quasi-static layer and breathing layer separated using the method in Ma et al. (2015), respectively, and $\alpha = 0.5$. The synthetic sequence has the same number of frames, same image size and resolution as its original in the clinical dataset. An example of a synthetic low-contrast XA is shown in Figure 3b. Note that the vessels have less contrast to the background than the original image in Figure 3a. We created a low-contrast XA sequence from each clinical XA described in Section 3.1.1, which results in 42 synthetic XA sequences in total.

3.1.3. X-ray angiograms of pigs

Additionally, *in vivo* XA data were acquired during a pig experiment performed at the Erasmus MC in Rotterdam, the Netherlands. 4 XA sequences with different contrast concentration levels were obtained from 1 FBM (familiar-hypercholesteremia Bretonchelles Meishan) pig which underwent a catheterization procedure after 14 months of high-fat diet. The XA sequences were acquired using a Siemens AXIOM-Artis monoplane system. The frame rate of all sequences is 15 frames per second. The number of frames per sequence varies from 48 to 79. The 4 XA sequences have in total 238 frames. All sequences have 776×776 pixels corresponding to a pixel size of $0.184 \times 0.184\, mm^2$. In all images, the inflow of contrast agent can be observed. The

XA images were made during a manual injection of iso-osmolar X-ray contrast medium (Visipaque 320, GE Healthcare, Buckinghamshire, U.K), delivered through the guide catheter. The full-contrast images were acquired with a contrast concentration of 320 g/mL. For the 25%, 50% and 75% contrast concentration images, the contrast agent was diluted accordingly with a 0,9% sodium-chloride solution (saline). Prior to image acquisition, the guide catheter was flushed with the right concentration of the contrast agent.

In practice, the full-contrast sequence had lower visual contrast than the images with 75% contrast. This might be due to incomplete flushing of the guiding catheter so that the contrast agent from the previous injection dilutes the current contrast agent.

3.2. Experiment 1: Parameter tuning for OR-PCA

OR-PCA has three parameters: the intrinsic rank of the subspace basis r and the regularization parameters λ_1 and λ_2. In Feng et al. (2013) and Song et al. (2015), both λ_1 and λ_2 were set to $1/\sqrt{p}$, where p is the dimension of data. This value had been proposed by Candès et al. (2011) as a general rule of thumb, but it can be slightly adjusted to achieve the best possible result. Javed et al. (2014), for example, have empirically selected different values for λ_1 and λ_2 instead of $1/\sqrt{p}$. Unlike the rule for choosing λ_1 and λ_2, the choice for r depends more on specific applications.

In order to find the optimal parameter setting for the layer separation application on the clinical XA data, we used the following way to quantify the outcome of layer separation with a certain set of parameters.

3.2.1. The definition of foreground and background

We firstly defined the "foreground" and the "background" for the objective of optimization in Section 3.2.2. It is worth noticing that the foreground and the background here are merely defined for computing the vessel contrast and thus should not be confused with the foreground and background's definition coming from the layer separation scheme described in the previous sections.

We used masks to define the foreground and background. A 1 *mm* wide area around manually-labeled vessel centerlines was considered as the foreground (shown as the dark area in the mask in Figure 4a). This area falls entirely within the vessel, and thus is a good representative of pixels belonging to vessels. For background, we adopted two different masks for measuring "global" and "local" contrast. The first one highlights all pixels outside a 4 *mm* wide area around the vessel centerlines (the white area in the mask in Figure 4b). This mask can quantify the effect of the removal of diaphragm, guiding catheters, etc. and can be used in a global measurement of contrast. The local background is defined as a 3 *mm* wide neighborhood area around the dark area in the global mask (the white area in the mask in Figure 4c).

For each clinical XA sequence, we randomly selected 8-15 frames for mask generation and contrast evaluation. The number of selected frames depends on the sequence length. As the vessel contrast is of main interest in this paper and in practice, only the frames with contrast agent were selected. This way we also avoided choosing non-contrast frames from the beginning of a sequence where the online algorithm has not converged yet. In total, 444 frames were chosen from 42 sequences.

We also created the masks for the four pig XA sequences. From each pig XA data, we randomly chose 8-12 frames. In total, 38 frames were chosen for the mask creation. These masks are only used for evaluation of the contrast level in pig data, not for parameter optimization.

3.2.2. The objective for parameter optimization

Metrics that have previously been used to measure vessel visibility include contrast-to-noise ratio (CNR), as the work in Ma et al. (2015), and the Jeffries-Matusita distance (JMD) (Zhu et al. (2009), Zhang et al. (2009)). These metrics evaluate the contrast of pixels from two groups, e.g. foreground and background. However, when tuning parameters of OR-PCA using either of these two measures as the objective for optimization, the optimal parameters are those that yield a very small standard deviation of the background, thus an almost constant background, and a flawed separation of vessel layer that loses much intensity of the vessel pixels. These would result in a large CNR or JMD, but do not lead to good foreground and background separation.

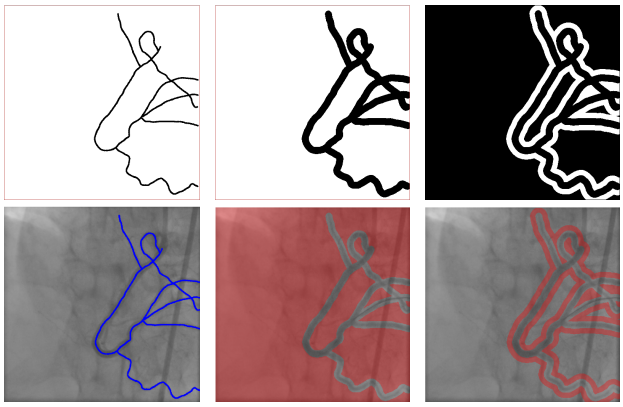

(a) (b) (c)

Figure 4: Examples of masks that are used for defining foreground and background in assessing the contrast of vessels. The first row shows the masks, the second row exhibits the overlay of masks on the corresponding original XA frame indicated by colors. (a) the foreground (blue), (b) the global background (red), (c) the local background (red).

To avoid the problem, the objective for OR-PCA parameter optimization should also consider the information loss in the vessel layer in addition to the vessel contrast. In this work, we integrated the difference between the original XA image and the sum of its three layers separated by OR-PCA method in the objective, such that losing too much information in the vessel layer would result in a large difference between the original XA and the sum of layer images. With this consideration, we used a corrected CNR (cCNR) as the objective to optimize the OR-PCA parameters:

$$cCNR = \frac{|\mu_F - \mu_B|}{\sqrt{\sigma_B^2 + MSE^V}} \tag{18}$$

where μ_F and μ_B are the mean of the pixel intensity value in the foreground and the background that were defined in Section 3.2.1, σ_B is the standard deviation of the pixel intensity in the background. MSE^V, the mean square error in the vessel area, which serves as a penalty term in Equation (18) to prevent too much information loss in the vessel layer, is defined as follows:

$$MSE^V = \frac{\sum_{x,y}(I^V_{original}(x,y) - I^V_{3-layer}(x,y))^2}{|I^V_{original}|} \tag{19}$$

where $I^V_{(\star)}$ denotes the operation that takes only pixels in the vessel area (defined by the dark region in the mask in Figure 4b) into consideration for image $I_{(\star)}$. The reason to focus only on the vessel area is that the information loss of vessel pixels only occurs in this region. $I_{original}(x,y)$ and $I_{3-layer}(x,y)$ are the pixel values of the position (x,y) in the original XA and the 3-layer sum image:

$$I_{3-layer} = I_{vessel} + I_{static} + I_{breathing} \tag{20}$$

where I_{vessel}, I_{static} and $I_{breathing}$ are the vessel layer, quasi-static layer and breathing layer, respectively. $|I^V_{original}|$ denotes the number of pixels in the vessel area in the frame.

According to (19), MSE^V indicates how well the original image can be reconstructed from the layer separation result. An undesirable reconstruction with pixel intensity loss in the vessel layer would result in a large MSE^V and, further, a small cCNR.

The cCNR for a complete sequence is defined as the average cCNR over all selected XA frames from the sequence. Global and local cCNR are computed respectively using the masks in Figure 4b and 4c.

3.2.3. Parameter optimization

The parameters of OR-PCA for both subspace basis update methods were optimized by exhaustively searching the optimal parameter set that maximizes the previously defined objective cCNR within a discrete set of parameters. First, cCNR was computed for every possible parameter combination $(\lambda_1, \lambda_2, r)$ within the parameter range for each clinical XA sequence. Then, the optimal parameters were obtained by searching for the parameter set that maximizes the average local cCNR over the 42 XA sequences. This optimization was performed for the two different basis update methods in Section 2.3.2.2 respectively.

The range for the intrinsic rank r was chosen as the integers in [2, 20]. The regularization parameter λ_1 and λ_2 were set to the same value as in Feng et al. (2013) and Song et al. (2015) both set $\lambda_{1,2}$ to $1/\sqrt{p}$. To search for the optimal $\lambda_{1,2}$, we explored the values in $[0.1/\sqrt{p},\ 10/\sqrt{p}]$ with a search step $0.1/\sqrt{p}$.

3.3. Experiment 2: Downweighting the past data in OR-PCA, Influence of the parameters

Once the optimal parameter settings of OR-PCA had been obtained from the previous experiment, we used this setting and study the influence of history parameters that were introduced in Section 2.3.3 for downweighting the past data on the performance of OR-PCA. The search spaces for exponential decay (ED), supra-linear decay (SLD) and sliding-window (SW) are:

$$\epsilon \in \{0.01, 0.3, 0.6, 0.9, 0.99, 0.9968, 0.999, 0.9997, 1-10^{-4}, 1-10^{-5}, 1-10^{-6}, 1-10^{-7}\},$$

$$\rho \in \{0.5, 1, 1.5, 2, 2.5, 3, 6, 9, ..., 36, 39, 42\},$$

$$t_0 \in [1, 15] \text{ and } t_0 \in \mathbb{N}.$$

In the experiments, we combined these three approaches with the two methods that update the subspace basis L that were mentioned in Section 2.3.2.2, i.e., block-coordinate descent (BCD) and the closed-form solution (CF). In this paper, for example, the OR-PCA method using exponential decay to downweight past data and using closed-form solution to update L is referred to as OR-PCA (ED+CF), or (ED+CF) as a short form.

As the OR-PCA parameters $\lambda_{1,2}$ and r tuned with cCNR assure a reliable layer separation, and tuning the parameters of ED, SLD and SW does not create the previously-mentioned undesirable layer separation, therefore, we evaluated the results for this experiment using the direct measure of vessel contrast in an image CNR, which were used in Ma et al. (2015), with the same masks from Section 3.2.1. In the evaluation, RPCA was used as a reference for the comparison purpose. Its regularization parameter λ had been optimized the same way as in Section 3.2.3 and was set to the optimal value $1.5/\sqrt{p}$ from the search space $[0.1/\sqrt{p},\ 10/\sqrt{p}]$ with a search step $0.1/\sqrt{p}$. The experiments were carried out with the 42 clinical XA sequences.

3.4. Experiment 3: Comparison with other methods

We compared the proposed approaches to several other related methods that can be used for prospective or online layer separation. The off-line method with the batch version of robust PCA in Ma et al. (2015) was used as a "benchmark" to show how close prospective or online methods can achieve to the performance of the off-line layer separation. The same way in Section 3.3, the regularization parameter λ of RPCA was set to $1.5/\sqrt{p}$ to achieve optimal performance. This method is referred to as **RPCA**. The following methods were tested in the experiment.

(a) **Median-subtraction** In Baka et al. (2014), static background has been suppressed by subtracting the median of the first 10 frames from each frame in the sequence. This method is referred to as **MS**.

(b) **Morphological-closing + median-subtraction** This advanced version of median subtraction method removes the breathing layer via morphological closing and then subtracting the median of the first 10 frames. This method is referred to as **MC+MS**.

(c) **Robust PCA with a sliding window** As mentioned in Section 1.2, Volpi et al. (2015) and Brosig et al. (2015) solved RPCA within a sliding window that consists of a few frames to enable prospective foreground separation. We adopted this idea of solving RPCA for our experiment. Different from their methods using Frangi filtering to preprocess images, to adapt to our application, we applied morphological closing to remove the breathing layer and then separate the other two layers from the difference images by solving RPCA with a sliding window. We used two sets of parameters for this method. The first one was used by Volpi et al. (2015) and Brosig et al. (2015): the window size was set to 4 and the regularization parameter λ was set to $1.5/\sqrt{p}$[1]. This one is referred to as **RPCA (SW)**. The second set of parameters was optimized the same way as in Section 3.2.3. The window size was set to 7 from the search space [2, 10] and the regularization parameter λ was set to $0.5/\sqrt{p}$ from the search space $[0.1/\sqrt{p},\ 10/\sqrt{p}]$ with a search step $0.1/\sqrt{p}$. This one is referred to as **RPCA (SW)***.

[1] λ was fixed to 3×10^{-3} in their papers. To adapt to different image sizes, we set λ to $1.5/\sqrt{p}$. This value is equivalent to 3×10^{-3} for images of size 512×512.

These methods were compared with the OR-PCA approaches using CNR as the evaluation metric with the same masks from Section 3.2.1. The CNR of OR-PCA approaches were computed in a leave-one-out (LOO) manner. In each LOO loop, firstly, the OR-PCA parameters $\lambda_{1,2}$ and r were optimized on the training sequences using the method in Section 3.2. Next, using the trained OR-PCA parameters, the history parameters were trained on the same set of sequences with local CNR using the approach in Section 3.3. Lastly, the optimal parameters obtained from the previous two steps were applied to compute CNR for the left-out data. The overall CNR was then computed as the average CNR over all LOO sessions. The LOO experiment was carried out with the 42 clinical XA sequences.

In addition to CNR as an evaluation metric of vessel contrast, to gain insight into how accurate the layers obtained with each method can reconstruct the original XA image, the reconstruction error was evaluated. It was computed as follows:

$$E_{recon} = \frac{\sum_{x,y} |I_{original}(x,y) - I_{3-layer}(x,y)|}{\sum_{x,y} I_{original}(x,y)} \quad (21)$$

where E_{recon} denotes the reconstruction error, $I_{original}$ and $I_{3-layer}$ are defined in the same way as Equation (19). The layer separation parameters used for computing E_{recon} were the same that were obtained during each LOO loop in the last paragraph. For each clinical XA sequence, the reconstruction error was computed for the frames that had been selected for mask generation in Section 3.2.1. The average error over all selected frames from a sequence was used as the empirical reconstruction error for this sequence. This error indicates the relative absolute difference between the pixel intensity of the reconstruction image and that of the original image with respect to the pixel intensity of the original image.

3.5. Experiment 4: Vessel enhancement in low-contrast XA

One possible application of layer separation is vessel enhancement. This can be achieved through enhancing the vessel layer and adding it back to the original image. To demonstrate this concept, we conducted experiments to enhance vessels in low-contrast XA using the online layer separation approaches.

The data we used are synthetic human XA data and real XA data acquired from pigs, as introduced in Section 3.1.2 and 3.1.3. We first separated the three layers, then enhanced the vessel layer by multiplying it by an enhancement factor $\beta > 0$. Finally, the vessel-enhanced image $I_{enhanced}$ equals the enhanced vessel layer plus the original XA image, as shown in Equation (22):

$$I_{enhanced} = \beta I_{vessel} + I_{original} \quad (22)$$

The results were evaluated using CNR. The layer separation method we used in this Section is OR-PCA (SW+CF).

For synthetic XA data, we used the parameters obtained from the leave-one-out evaluation in Section 3.4 for each synthetic data. For the pig XA data, the parameter set $(\lambda_{1,2}, r, t_0) = (2.1, 5, 3)$ was used.

3.6. Implementation

All algorithms were implemented in MATLAB (The MathWorks, Inc.). In particular, the computation time of layer separation was recorded in MATLAB 2014a on an Intel Core i7-4800MQ 2.70 GHz computer with 16 GB RAM running Windows 7.

4. Results

4.1. Optimal parameters for OR-PCA

The parameters of OR-PCA optimized over the whole XA dataset for the two different basis update methods are shown in Table 1. Here both λ_1 and λ_2 are set to the same value. Comparing the two methods, the $\lambda_{1,2}$ have similar values, while the intrinsic ranks r of the subspace are very different.

Table 1: The optimal parameter settings of OR-PCA for different subspace basis update methods. p is the dimension of the data, i.e. the number of pixels in a frame.

Basis Update Method	$\lambda_{1,2}$	r
Block-coordinate Descent (BCD)	$2.3/\sqrt{p}$	14
Closed-form Solution (CF)	$2.1/\sqrt{p}$	5

Using block-coordinate descent with the optimal parameter setting, an example of online layer separation of an XA sequence (512×512, 55 frames) is shown in Fig. 5. Note that the layer separation result for the first frame shows strong artefacts (e.g. the vertebral shape in the vessel layer) due to random initialization of the subspace basis. As time proceeds, the layer separation improves quickly. The 10^{th} frame already has a good layer separation.

4.2. Influence of the history parameters

With the optimal parameter setting of OR-PCA, we quantitatively assessed how the history parameters mentioned in Section 2.3.3 and 3.3 influence the layer separation performance. Depending on the image content, the influence of the history parameters on each individual sequence may vary from sequence to sequence (see Figure 1, Figure 2 and Figure 3 in Supplementary Material). The average measures over the whole dataset are shown in Figure 6 where the CNR values are normalized by dividing CNR by the CNR value obtained from the RPCA method (so RPCA has a constant value 1).

Compared to the original OR-PCA method (OR-PCA (BCD)) and OR-PCA (CF), the history parameters ϵ, ρ and t_0 all resulted in overall higher average local and global CNR. This improvement was more prominent in the case of global CNR compared to local CNR. It is also worth noticing that the combination of the history parameters with the CF method generally performed better than the (history parameter + BCD) option, which could be especially seen in the sliding-window case.

In addition to the overall comparisons of the three approaches, each of them presented a certain trend of CNR as the history parameter changed. For exponential decay (Figure 6a and 6b), as $-log(1-\epsilon)$ increases, the CNR values firstly increased fast; once they reached an optimal value when $-log(1-\epsilon)$ is between 0.5 and 2.5 (except for OR-PCA (ED+CF)), the CNRs dropped down slowly to reach a constant. In the case of supra-linear decay (Figure 6c and 6d), the CNRs increased as ρ increases, but did not change much when ρ is larger than 15. For the sliding-window approach (Figure 6e and 6f), although in general the CNRs dropped when the window size becomes larger, OR-PCA (SW+BCD) reached its optimum when t_0 equals 2, whereas OR-PCA (SW+CF) had an optimal t_0 between 3 and 5. Finally, as $-log(1-\epsilon)$ and ρ kept increasing (decreasing the weights of all past frames), the CNR curves of ED or SLD converged to where the CNR curves of SW started (only preserve information of the new frame).

4.3. Comparison with other methods

The optimal parameter sets for the methods based on OR-PCA that were obtained during the leave-one-out evaluation are listed in Table 2. In general, the methods that use closed-form solution (CF) had a smaller r than methods with BCD, but needed more information from the past data (lower ϵ and ρ, higher t_0) to achieve the best performance.

Table 2 also shows the counts of the optimal parameter sets that each LOO session generated for each method. For methods without history parameters (BCD and CF), the optimal parameter set with the largest count were identical to the results in Section 4.1. Most of the methods with a history parameter had a dominant optimal parameter set from LOO, except for the method (SW+CF) where the two most dominant optimal parameter sets had almost equal counts.

Table 3 lists the average CNR values for the original XA sequences and the vessel layers obtained with each method. Compared to the original XA, all methods achieved a substantial improvement on the CNR values in the vessel layer. Compared to method MS, MC+MS, RPCA (SW) and RPCA (SW)*, the methods that use OR-PCA (from BCD to SW+CF in Table 3) had higher CNR. The CNR values of the two types of methods (using or not using OR-PCA) were statistically significantly different with the two-sided Wilcoxon signed-rank test (see Table 1 in Supplementary Materials). The methods that downweight the past data (from ED+BCD to SW+CF) were able to improve the vessel contrast over the methods without history parameters (BCD and CF). The improvement was statistically significant (see Table 2 in Supplementary Materials). Among all the methods that are based on OR-PCA, ED+CF, SLD+CF and SW+CF showed similar or better average local CNR than RPCA, although without statistical significance. The performance of all OR-PCA based methods was

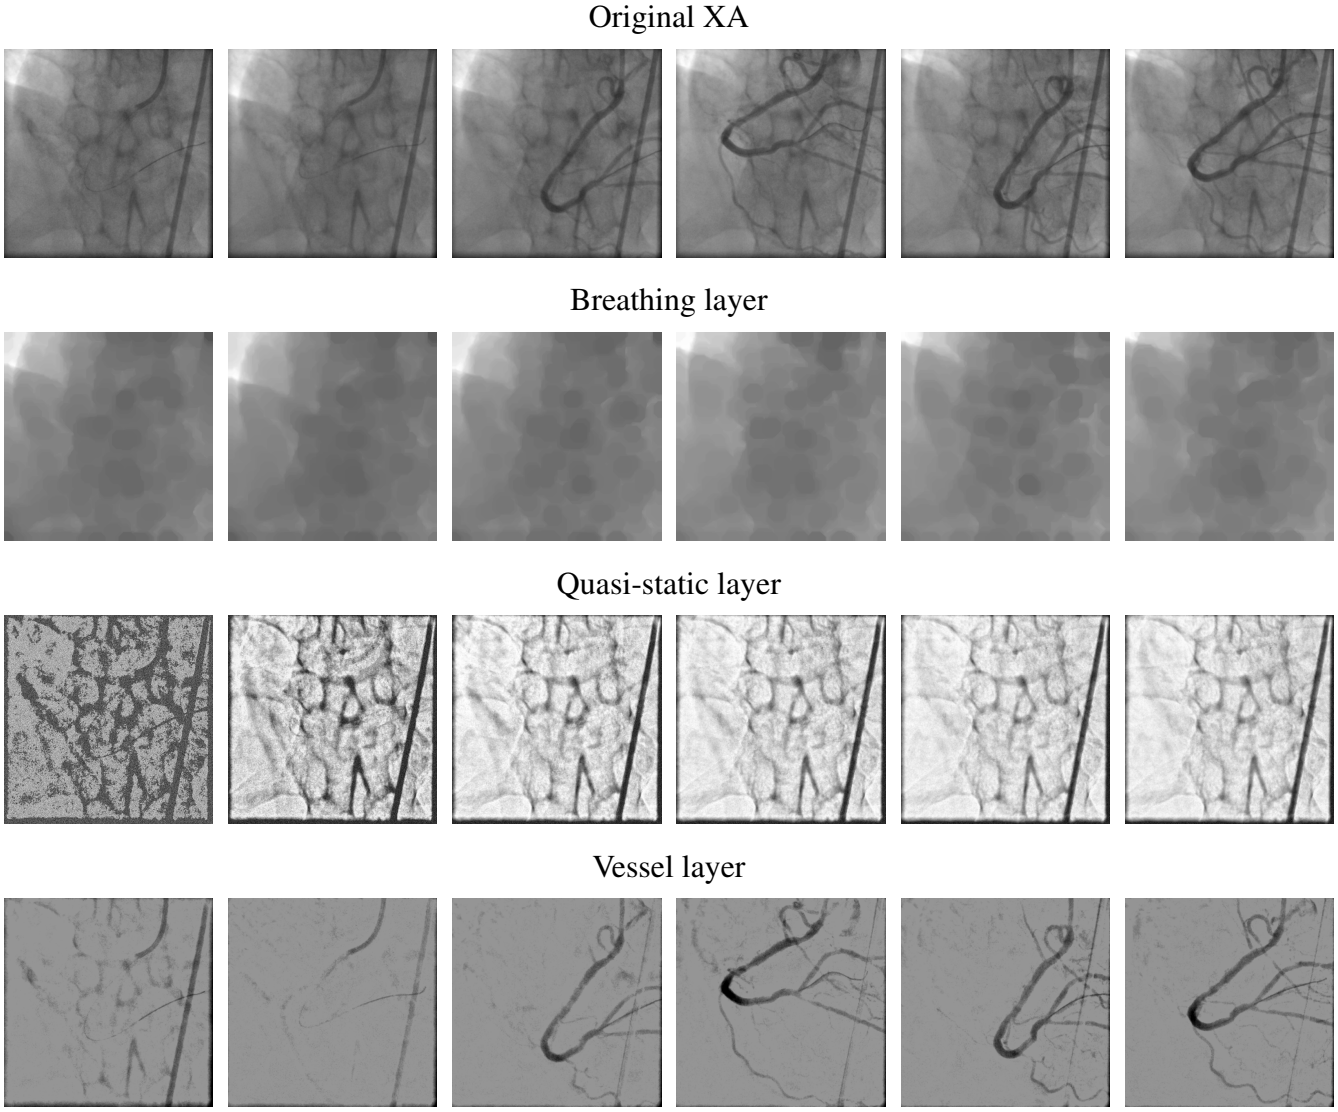

Figure 5: An example of online layer separation of an XA sequence using OR-PCA with the optimal parameter settings as listed in Table 1. The subspace basis update method used here is block-coordinate descent (BCD). Row 1-4 show the original frames, the breathing layer, the quasi-static layer and the vessel layer, respectively. Column 1-6 are 6 frames taken from the sequence in a chronological order and their layer separation outcomes. The frame ordinals from left to right are 1^{st}, 5^{th}, 10^{th}, 15^{th}, 20^{th}, 25^{th}.

closer to the off-line benchmark RPCA than those not using OR-PCA.

Table 3 also shows the average reconstruction error between the original image and the three-layer sum image for each method. For MS, MC+MS and the RPCA (SW) methods, the foreground (vessel layer) were obtained via subtraction of the background. therefore these methods, by definition, have a reconstruction error of zero. The methods based on OR-PCA made minor reconstruction errors (less than 3% of the average pixel intensity of the original images). The history downweighting techniques reduced the reconstruction errors of BCD and CF for about 17% and 26% respectively. These errors are about three times larger than the one made by RPCA.

The comparison between different methods is illustrated in Figure 7, where the CNR values were normalized in the same way as in Figure 6. Similarly, the methods that use OR-PCA outperformed the other methods on both local and global CNR. The improvement that results from downweighting history was more substantial in global CNR than local CNR. For the methods that are based on OR-PCA, those that use the closed-form solution achieved slightly better normalized CNR values than the ones with BCD.

Figure 8 presents examples of results of five representative prospective or online layer separation methods and original images on four XA sequences: MC+MS, RPCA (SW), OR-PCA (CF), OR-PCA (SW+CF) and RPCA. All methods improved the visibility of vessels in the image, but MC+MS generated strong artefacts near the guiding catheters. RPCA (SW) shows slightly better results than MC+MS, but still presents some motion artefacts near the guiding catheters. OR-PCA (CF) achieved good layer separation, while OR-PCA (SW+CF) was able to produce "cleaner" background (row 1 column 4 and 5, the left part of the images) and still maintained the vessel information (column 4 and 5). The vessel layer separated using OR-PCA (SW+CF) had more similar appearances to the ones produced by RPCA (column 5 and 6) than other methods.

4.4. Vessel enhancement in low-contrast XA

The results of vessel-enhancement on synthetic low-contrast XA and real pig XA data are shown in this section.

4.4.1. Synthetic low-contrast XA

Table 4 shows the average CNR values for vessel enhancement in synthetic XA sequences. With enhancement factor $\beta = 1$, the vessel-enhanced XA showed better local and global CNR than the synthetic XA with statistical significance. Compared to the original XA in Table 3, the CNR values of the vessel-enhanced XA was slightly lower, but these CNR values could be improved with a larger enhancement factor β.

A few examples of vessel-enhancement on synthetic XA data are shown in Figure 9. Compared to the original images (the first row), the synthetic XA (the second row) had poorer vessel contrast. The proposed layer separation method (SW+CF) was still able to extracted the vessel layer (the third row), while maintaining a majority of the information, and enhance the vessel contrast (the last row) to the visually similar level of the original images.

4.4.2. Real XA of pigs

We show the CNR values for vessel enhancement experiment with pig XA data in Table 5. In the table, since the full-contrast sequence showed lower visual contrast (see Section 3.1.3), the four sequences were sorted by their local CNR values in an ascending order. With enhancement factor $\beta = 2$, the vessel layers and the enhanced sequences showed an improvement on local and global CNR. This improvement increased as the local CNR of the sequence became higher.

Similar observation could be found in Figure 10, where representative frames from each pig XA sequence are shown. For example, the proposed method was able to increase the vessel contrast in the image of 25% contrast to the similar level as in the image of 50% contrast (see Figure 10i and 10b). The vessel contrast in the enhanced image of 50% contrast (Figure 10j) had better contrast than the image of full contrast (Figure 10c) and 75% contrast (Figure 10d). It is also observed that, in Figure 10 from left to right, the false positive enhancement of non-vessel structures (e.g. the dark spots in the right part of the images) decreased as the visual contrast increased.

4.5. Computation time

The computation time of layer separation for each frame is shown in Figure 11. In this figure, the box plots of per-frame processing time for each method that is based on OR-PCA is illustrated. The processing times of these methods ranged from 0.15 to 1.60 seconds per frame. The methods that use a closed-form solution to update L were approximately two times faster than the ones that use block-coordinate descent. The methods which use ED and SLD to treat past information were slightly faster than their counterparts that do not weight past data, while the methods with SW needed slightly longer time to process one frame than their corresponding methods BCD and CF. The "outliers" shown as red marks in Figure 11 are from images of larger size.

The average computation time per image size is shown in Table 6. Generally, the table shows that images of larger size needed longer processing time per frame. On XA images of common size in clinics (512×512, 600×600 and 776×776), it is possible to achieve a processing rate of 3-6 frames per second (fps) with the proposed methods on our hardware.

Compared to MS and MC+MS, OR-PCA based methods showed much better layer separation performance (see Table 3 and Figure 7), while the processing time of the fastest one (ED+CF) is only about 40% slower than MC+MS.

The processing time for RPCA (SW) is also shown in Table 6. In RPCA (SW), the foreground separation of a frame is delayed by the processing of its previous data block, therefore, we computed the average per frame processing time of all blocks in a sequence. Compared to RPCA (SW), the methods that use OR-PCA perform faster with a factor up to 3.

5. Discussion and conclusion

We have presented a fast automatic online method to robustly separate cardiac interventional X-ray angiograms into three image layers: a breathing layer, a quasi-static layer and finally a vessel layer that contains information of moving thin structures, such as coronary arteries. The method relied on morphological closing and online robust PCA and we investigated different possibilities for downweighting information from previous frames for further improving layer separation. The parameters of OR-PCA were optimized on 42 clinical XA sequences. In addition, a pilot study was performed on synthetic XA sequences and pig data to show the potential of the proposed method for vessel enhancement in XA.

The integration of OR-PCA algorithm into layer separation enables online processing XA images from the beginning of the sequence. The mechanism behind this is that OR-PCA only needs to be "fed" one frame each time, but is able to update the subspace basis of the low-rank component based on the new information. This is an important difference from the method in Ma et al. (2015) which worked "off-line" and needed the complete sequence as input. The proposed approach is also different from the method in Volpi et al. (2015), where the online implementation needed several frames to solve RPCA in a mini-batch manner, and hence, resulting in a delay in processing for each following mini-batch. Although the layer separation results of our proposed method might suffer from random initialization at the beginning, the algorithm converges fast and obtain reasonable layer separation after a few frames (see an example in Figure 5).

In addition to the advantage of online processing, the methods that use OR-PCA show good performance on layer separation. It significantly improves the vessel visibility of the original XA images with minor reconstruction errors: the background structures were removed and the vessel contrast was visually and quantitatively enhanced. Compared to those methods that model a total static background, e.g. MS and MC+MS, the approaches that are based on OR-PCA are superior because they are able to model a dynamic scene, therefore can adapt to small dynamic changes in the background. The method that separates layers by solving RPCA with a mini-batch of data suffers from motion artefacts that remain around vessels and catheters. This might be because it uses the frames in the previous block to infer the background for the current block of images, which might fall behind the background change. Furthermore, it uses the same background for all images in the block to compute their vessel layers, which ignores the possible small background change within the block. OR-PCA updates L frame by frame to keep up the background change and has a "customized" background for each individual frame.

It is also worth noticing the advantage of removing breathing structures prior to the separation of the other two layers with OR-PCA. If the original XA image is directly "fed" to the OR-PCA method, the breathing structures should stay in the same layer with the static structures in order to obtain a reasonable vessel layer, because OR-PCA decomposes a source image into only two components. Due to breathing motion, this background layer will contain more variation than a layer that only contains quasi-static structures, which might require a much higher r parameter of OR-PCA to allow reasonable convergence of the algorithm. However, it is often inevitable to observe strong breathing motion artefacts in the output vessel layer. Figure 12 provides an example that illustrates the cases without removing breathing structures prior to the OR-PCA computation. In Figure 12c, the image still contains a large amount of noise and a static dark band on the left. In Figure 12d, OR-PCA has a better convergence with a higher r, less noise and no dark band is observed, but a stronger artefact of diaphragm remains.

The parameters of OR-PCA used in this work were optimized based on our image data, instead of being assigned the "rule of thumb" value $1/\sqrt{p}$ as in the works of Song et al. (2015) and Javed et al. (2014). The optimum of $\lambda_{1,2}$ for method BCD and CF are similar and the values are close to $1/\sqrt{p}$. The optimal intrinsic rank r for BCD and CF are different. A possible explanation is that CF computes the quasi-static subspace basis L in one step, whereas BCD updates L column-wisely and thus needs more variations to achieve the same accuracy as CF.

The performance of layer separation using OR-PCA can be improved by downweighting the past frames. In this work, we have proposed three different ways: exponential decay (ED) and supra-linear decay (SLD) are methods to scale all past data, and the sliding-window approach (SW) only preserve the information of the few most recent frames, which could be interpreted as "binary" scaling. The results on vessel visibility and reconstruction error showed that all three ways improve the overall layer separation by giving recent frames higher weight than earlier frames. This suggest that not all past information is necessary for best inferring the current status in the scenario of online learning. A possible explanation of this finding is that using the downweighting techniques in the online learning algorithm promotes faster convergence (Mairal et al. (2010)). Although the improvement might vary between sequences, depending on specific image content, most of them present an improvement with history parameters (see supplementary material).

The optimized history parameters show that only the most recent 2-5 frames are needed to update the subspace basis L, which seems too "few". The reasons might be two-fold. First, since L is the subspace basis of the quasi-static layer which does not contain much variation, it should not need information from a large number of frames to update L. Second, for the case of BCD, note that in Equation (10), the update of L still partly relies on its previous version, not solely on A and B. This means that every version of L can inherit information from its previous version and thus is a compact representation of all past information, but updating L needs only the very recent frames.

The combination of the subspace basis update methods and the past information downweighting techniques yields 8 different OR-PCA variants, while in practice one might choose one of them for a layer separation task. In terms of the performance on improving vessel visibility, (SW+CF) might be the best choice, as it gives the highest CNR value. On the other hand, if speed is of great concern, (ED+CF) is a good option, since it runs the fastest among the eight and the performance on CNR is not much worse than (SW+CF). In addition, the implementation of (ED+CF) is easier than (SW+CF), in that it does not explicitly store a few past values of A and B, but does the scaling implicitly.

In terms of computational efficiency, the methods that use OR-PCA run fast. For 512×512 frames, OR-PCA with CF was able to achieve a 5-6 fps processing rate on standard PC; for 1024×1024 frames, the proposed methods could reach at highest about 2 fps. This is faster than RPCA (SW) based approaches either from the result of our experiment (see Table 6) or the literature, e.g. Volpi et al. (2015) who reported about 1 fps for frame size in the range of 824×1024 to 1024×1024, and Brosig et al. (2015) who achieved 3 fps for 512×512 images. According to Feng et al. (2013), the computation complexity for batch RPCA is $O(np^2)$ and for OR-PCA is $O(pr^2)$. Since $p \gg r$ and normally also $p \gg r^2$, OR-PCA runs much faster than RPCA. We also notice that OR-PCA with CF ran faster than BCD, this is because the computation complexity of both methods is $O(pr^2)$, and in our experiments, the r for CF was smaller than the r for BCD. The methods with history downweighting schemes are faster than BCD and CF. This might be due to the faster convergence of OR-PCA when downweighting the past information. Finally, it is worth noticing that the timing reported in this paper were based on a MATLAB implementation that ran on a single CPU core. A parallelized version of the method may achieve real-time processing rate (about 15 fps) for clinical applications.

One of the potential direct clinical applications of the proposed method is to enhance vessels in X-ray images with low vessel contrast, which suffer from poor diagnostic quality. X-ray contrast agent used for angiography may have side effects including allergic reaction which can be life-threatening, and nephrotoxicity (contrast-induced nephropathy, CIN)

which can result in chronic renal failure with all its sequelae (Andreucci et al. (2014), Tepel et al. (2006)). Thus, it is clinically relevant to limit the use of X-ray contrast agent during interventions. The method we proposed in this paper provides a possibility to use X-ray contrast of lower concentrations. In the experiments, we have evaluated whether it is possible to achieve better vessel visibility on low-contrast images by enhancing vessels using the methods based on OR-PCA. We have used synthetic 50%-contrast XA images and 4 real pig XA sequences with different contrast concentrations for the test. The results showed a good improvement on the vessel visibility on both kinds of images, implying a potential application for coronary interventions.

The proposed layer separation methods are based on some assumptions. First, the morphological closing operation with a circular structuring element of 8.5 *mm* in diameter worked well for small vessels, such as coronary arteries. However, for other kinds of interventions that operate on large vessels, such as aorta or pulmonary arteries, the structuring element of the proposed size is not large enough. In those cases, one might consider using a larger structuring element for morphological closing and adjusting the parameters of OR-PCA for a reasonable layer separation. Another important assumption underlying the methods that use OR-PCA is that there is dynamic change in the foreground, and it detects the dynamic change. This assumption holds true most of the time because coronary arteries always move together with heartbeat. However, in the case that the guiding catheter tip segment moves together with heartbeat, the proposed layer separation method cannot separate this moving catheter segment from vessels. The method also requires a certain amount of contrast agent, i.e. the signal of the vessels should not be too weak. As the methods based on OR-PCA assumes a sparse foreground, when the contrast of vessel is not strong enough, the proposed methods might enhance noise or detect other non-vessel structures in the foreground, as can be seen in Figure 10.

In the future, it is of great interest to investigate the potential of the proposed method. One important direction would be to evaluate the clinical potential, e.g. how the proposed layer separation method would work under different contrast concentration levels on a larger dataset. One could also think of improving the visibility of instruments, such as catheters or guidewires for other cardiac applications. From a methodological point of view, it might be interesting to unify the steps in the proposed method into one optimization problem, for example, incorporating the heuristic post-processing into the OR-PCA algorithm with an additional non-positive constraint on the sparse component.

In conclusion, we have presented a fast automatic online layer separation method for robust vessel enhancement in X-ray angiograms. The method separated an XA frame into three layers: a breathing layer, a quasi-static layer and a vessel layer. We proposed three ways to improve the layer separation outcome by downweighting the past frames. The proposed method significantly improved the vessel visibility and outperformed other related prospective or online layer separation approaches. The method does not need much computation time, making it potentially applicable for clinical practice without the necessity of using advanced hardware, opening the way for relevant clinical applications, such as improving the vessel visibility under conditions of low contrast concentrations, so as to allow a reduced amount of contrast agent usage to prevent contrast-induced side effects.

Acknowledgement

This work was supported by Technology Foundation STW, IMAGIC project under the iMIT program [grant number 12703].

References

Andreucci, M., Solomon, R., Tasanarong, A., 2014. Side effects of radiographic contrast media: pathogenesis, risk factors, and prevention. BioMed Research International 2014.

Auvray, V., Bouthemy, P., Liénard, J., 2009. Jointmotion estimation and layer segmentation in transparent image sequenc: application to noise reduction in X-ray image sequences. EURASIP Journal on Advances in Signal Processing 2009, 19.

Baka, N., Lelieveldt, B., Schultz, C., Niessen, W., van Walsum, T., 2015. Respiratory motion estimation in X-ray angiography for improved guidance during coronary interventions. Physics in Medicine and Biology 60, 3617.

Baka, N., Metz, C., Schultz, C., Neefjes, L., van Geuns, R.J., Lelieveldt, B.P., Niessen, W.J., van Walsum, T., de Bruijne, M., 2013. Statistical coronary motion models for 2D+ t/3D registration of X-ray coronary angiography and CTA. Medical Image Analysis 17, 698–709.

Baka, N., Metz, C., Schultz, C.J., van Geuns, R.J., Niessen, W.J., van Walsum, T., 2014. Oriented Gaussian mixture models for nonrigid 2D/3D coronary artery registration. Medical Imaging, IEEE Transactions on 33, 1023–1034.

Bouwmans, T., Zahzah, E.H., 2014. Robust PCA via principal component pursuit: a review for a comparative evaluation in video surveillance. Computer Vision and Image Understanding 122, 22–34.

Brosig, R., Hariharan, S.G., Volpi, D., Kowarschik, M., Carlier, S., Navab, N., Demirci, S., 2015. Implicit background subtraction for cardiac digital angiography, in: Joint MICCAI workshops on Computing and Visualisation for Intravascular Imaging and Computer-Assisted Stenting.

Candès, E.J., Li, X., Ma, Y., Wright, J., 2011. Robust principal component analysis? Journal of the ACM 58, 11.

Close, R.A., Abbey, C.K., Morioka, C.A., Whiting, J.S., 2001. Accuracy assessment of layer decomposition using simulated angiographic image sequences. Medical Imaging, IEEE Transactions on 20, 990–998.

Feng, J., Xu, H., Yan, S., 2013. Online robust PCA via stochastic optimization, in: Advances in Neural Information Processing Systems, pp. 404–412.

Fischer, P., Pohl, T., Köhler, T., Maier, A., Hornegger, J., 2015. A robust probabilistic model for motion layer separation in X-ray fluoroscopy, in: Information Processing in Medical Imaging, Springer. pp. 288–299.

Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A., 1998. Multiscale vessel enhancement filtering, in: Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 130–137.

Hamy, V., Dikaios, N., Punwani, S., Melbourne, A., Latifoltojar, A., Makanyanga, J., Chouhan, M., Helbren, E., Menys, A., Taylor, S., et al., 2014. Respiratory motion correction in dynamic MRI using robust data decomposition registration–application to DCE-MRI. Medical Image Analysis 18, 301–313.

Hyvärinen, A., Karhunen, J., Oja, E., 2004. Independent component analysis. volume 46. John Wiley & Sons.

Javed, S., Oh, S.H., Sobral, A., Bouwmans, T., Jung, S.K., 2014. OR-PCA with MRF for robust foreground detection in highly dynamic backgrounds, in: Asian Conference on Computer Vision. Springer, pp. 284–299.

Javed, S., Sobral, A., Bouwmans, T., Jung, S.K., 2015. OR-PCA with dynamic feature selection for robust background subtraction, in: Proceedings of the 30th Annual ACM Symposium on Applied Computing, ACM. pp. 86–91.

Lin, Z., Liu, R., Su, Z., 2011. Linearized alternating direction method with adaptive penalty for low-rank representation, in: Advances in Neural Information Processing Systems, pp. 612–620.

Ma, H., Dibildox, G., Banerjee, J., Niessen, W., Schultz, C., Regar, E., van Walsum, T., 2015. Layer separation for vessel enhancement in interventional X-ray angiograms using morphological filtering and robust PCA, in: Augmented Environments for Computer-Assisted Interventions. Springer, pp. 104–113.

Mairal, J., Bach, F., Ponce, J., Sapiro, G., 2010. Online learning for matrix factorization and sparse coding. The Journal of Machine Learning Research 11, 19–60.

Otazo, R., Candès, E., Sodickson, D.K., 2015. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Magnetic Resonance in Medicine 73, 1125–1136.

Panayiotou, M., King, A.P., Housden, R.J., Ma, Y., Cooklin, M., O'Neill, M., Gill, J., Rinaldi, C.A., Rhode, K.S., 2014. A statistical method for retrospective cardiac and respiratory motion gating of interventional cardiac X-ray images. Medical Physics 41, 071901.

Preston, J.S., Rottman, C., Cheryauka, A., Anderton, L., Whitaker, R.T., Joshi, S.C., 2013. Multi-layer deformation estimation for fluoroscopic imaging., in: Information Processing in Medical Imaging, pp. 123–134.

Rivest-Henault, D., Sundar, H., Cheriet, M., 2012. Nonrigid 2D/3D registration of coronary artery models with live fluoroscopy for guidance of cardiac interventions. Medical Imaging, IEEE Transactions on 31, 1557–1572.

Shechter, G., Shechter, B., Resar, J.R., Beyar, R., 2005. Prospective motion correction of X-ray images for coronary interventions. Medical Imaging, IEEE Transactions on 24, 441–450.

Song, W., Zhu, J., Li, Y., Chen, C., 2015. Image alignment by online robust PCA via stochastic gradient descent. IEEE Transactions on Circuits and Systems for Video Technology pp (99).

Tang, S., Wang, Y., Chen, Y.W., 2012. Application of ICA to X-ray coronary digital subtraction angiography. Neurocomputing 79, 168–172.

Tepel, M., Aspelin, P., Lameire, N., 2006. Contrast-induced nephropathy a clinical and evidence-based approach. Circulation 113, 1799–1806.

Volpi, D., Sarhan, M.H., Ghotbi, R., Navab, N., Mateus, D., Demirci, S., 2015. Online tracking of interventional devices for endovascular aortic repair. International Journal of Computer Assisted Radiology and Surgery 10, 773–781.

Zhang, W., Ling, H., Prummer, S., Zhou, K.S., Ostermeier, M., Comaniciu, D., 2009. Coronary tree extraction using motion layer separation, in: Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 116–123.

Zhu, Y., Prummer, S., Wang, P., Chen, T., Comaniciu, D., Ostermeier, M., 2009. Dynamic layer separation for coronary DSA and enhancement in fluoroscopic sequences. Medical Image Computing and Computer-Assisted Intervention , 877–884.

Table 2: Counts of the optimal parameter sets obtained during the leave-one-out evaluation for each OR-PCA method. As 42 sequences are used, that yields 42 leave-one-out sessions and 42 optimal parameter sets in total.

Method	$\lambda_{1,2}$ $(1/\sqrt{p})$	r	History Parameter	Counts
BCD	2.3	14		39
	2.4	17		2
	3.3	19		1
CF	2.1	5		31
	2.3	7		11
ED+BCD	2.3	14	$\epsilon = 0.9$	39
	2.4	17	$\epsilon = 0.9$	2
	3.3	19	$\epsilon = 0.9$	1
ED+CF	2.1	5	$\epsilon = 0.6$	31
	2.3	7	$\epsilon = 0.6$	8
	2.3	7	$\epsilon = 0.3$	3
SLD+BCD	2.3	14	$\rho = 42$	39
	2.4	17	$\rho = 18$	2
	3.3	19	$\rho = 9$	1
SLD+CF	2.1	5	$\rho = 36$	29
	2.3	7	$\rho = 36$	11
	2.1	5	$\rho = 15$	2
SW+BCD	2.3	14	$t_0 = 2$	38
	2.4	17	$t_0 = 2$	2
	2.3	14	$t_0 = 1$	1
	3.3	19	$t_0 = 2$	1
SW+CF	2.1	5	$t_0 = 3$	17
	2.1	5	$t_0 = 5$	14
	2.3	7	$t_0 = 5$	11

Table 3: The average local CNR, global CNR and reconstruction error E_{recon} (mean value ± standard deviation) for the original XA and all methods.

Method	Local CNR	Global CNR	E_{recon}
Original XA	0.991 ± 0.330	0.507 ± 0.305	
MS	1.811 ± 0.490	2.211 ± 0.648	0
MC+MS	2.396 ± 0.635	3.210 ± 0.869	0
RPCA (SW)	2.134 ± 0.511	2.278 ± 0.676	0
RPCA (SW)*	2.170 ± 0.578	1.881 ± 0.623	0
(BCD)	3.010 ± 1.065	3.453 ± 1.223	0.026 ± 0.005
(CF)	3.010 ± 1.088	3.509 ± 1.204	0.030 ± 0.007
(ED+BCD)	3.209 ± 1.192	4.257 ± 1.316	0.021 ± 0.006
(ED+CF)	3.226 ± 1.263	4.422 ± 1.466	0.023 ± 0.007
(SLD+BCD)	3.169 ± 1.172	4.085 ± 1.282	0.022 ± 0.006
(SLD+CF)	3.246 ± 1.224	4.500 ± 1.393	0.022 ± 0.006
(SW+BCD)	3.170 ± 1.161	4.150 ± 1.295	0.022 ± 0.006
(SW+CF)	3.281 ± 1.290	4.602 ± 1.465	0.022 ± 0.007
RPCA	3.227 ± 1.301	5.176 ± 2.004	0.007 ± 0.002

Table 4: The average local and global CNR (mean value ± standard deviation) for the synthetic XA data, the vessel layers separated from the synthetic data using OR-PCA (SW+CF), and the vessel-enhanced XA sequences ($\beta = 1$). The two-sided Wilcoxon sign-rank test indicates statistically significantly difference in local and global CNR between the synthetic data and the vessel-enhanced images ($p < 0.01$).

Image types	Local CNR	Global CNR
Synthetic XA	0.592 ± 0.236	0.338 ± 0.245
Vessel Layer	3.170 ± 1.290	4.048 ± 1.592
Enhanced	0.875 ± 0.312	0.452 ± 0.285

Table 5: The local and global CNR of the 4 XA sequences obtained from pigs, the separated vessel layer and their vessel-enhanced sequence ($\beta = 2$). The 4 sequences are sorted by their local CNR values in an ascending order.

Contrast Concentration	25%		50%		100%		75%	
Metric (CNR)	local	global	local	global	local	global	local	global
Original	0.307	0.164	0.490	0.340	0.523	0.395	0.690	0.569
Vessel Layer	2.546	0.682	5.156	0.830	5.175	2.759	6.835	5.473
Enhanced	0.488	0.262	0.954	0.562	1.041	0.729	1.528	1.147

Table 6: The average processing time (seconds/frame) for each layer separation method on XA sequences of different frame size. The bold numbers indicate processing time that the fastest OR-PCA based method needs for a certain frame size.

Method	512×512	600×600	776×776	1024×1024
BCD	0.36	0.51	0.69	1.14
CF	0.20	0.30	0.36	0.58
ED+BCD	0.30	0.41	0.63	1.02
ED+CF	**0.16**	**0.21**	**0.34**	**0.57**
SLD+BCD	0.30	0.42	0.63	1.02
SLD+CF	0.17	0.22	**0.34**	**0.57**
SW+BCD	0.36	0.50	0.75	1.27
SW+CF	0.22	0.31	0.46	0.77
RPCA (SW)	0.37	0.53	0.85	1.50
RPCA (SW)*	0.46	0.68	1.10	1.97
MS	0.02	0.03	0.05	0.11
MC+MS	0.11	0.15	0.25	0.45

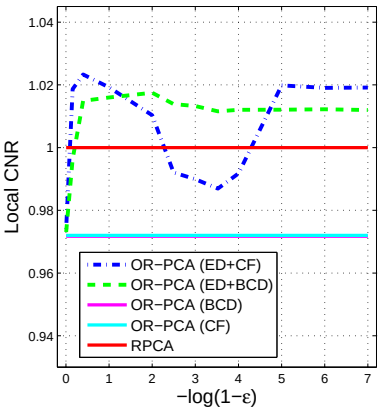

(a) Exponential decay (local CNR)

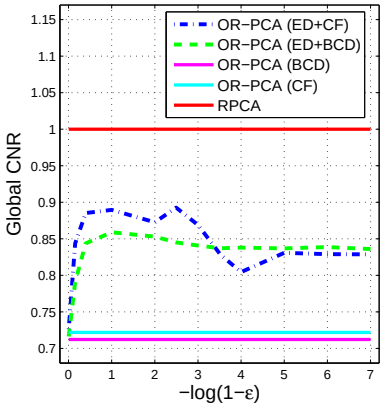

(b) Exponential decay (global CNR)

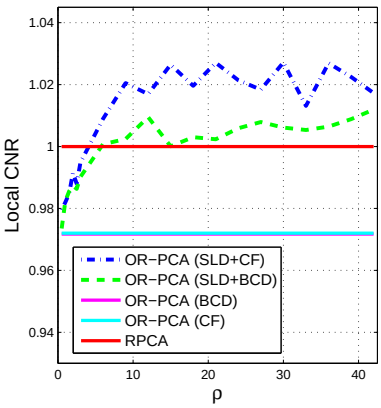

(c) Supra-linear decay (local CNR)

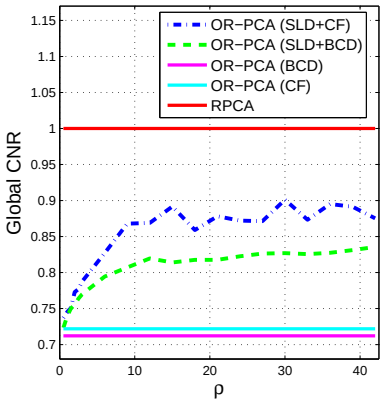

(d) Supra-linear decay (global CNR)

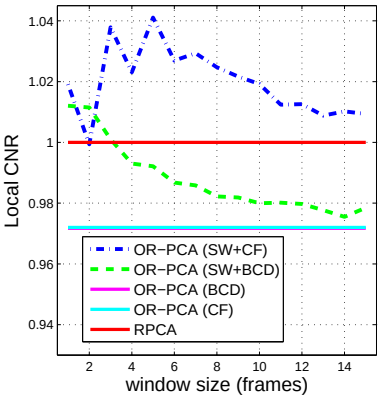

(e) Sliding-window (local CNR)

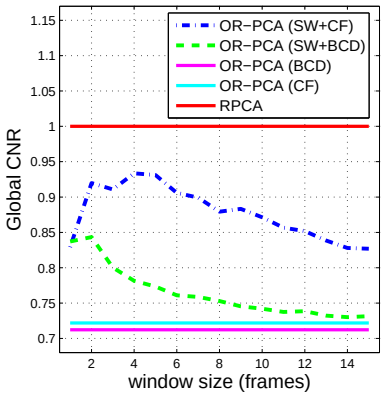

(f) Sliding-window (global CNR)

Figure 6: The influence of history parameters ϵ, ρ and the sliding-window size t_0 on local and global CNR. All values are normalized using RPCA method as the reference. In (a), (c) and (e), the local CNR for OR-PCA (BCD) and OR-PCA (CF) are very close that the two lines almost overlap.

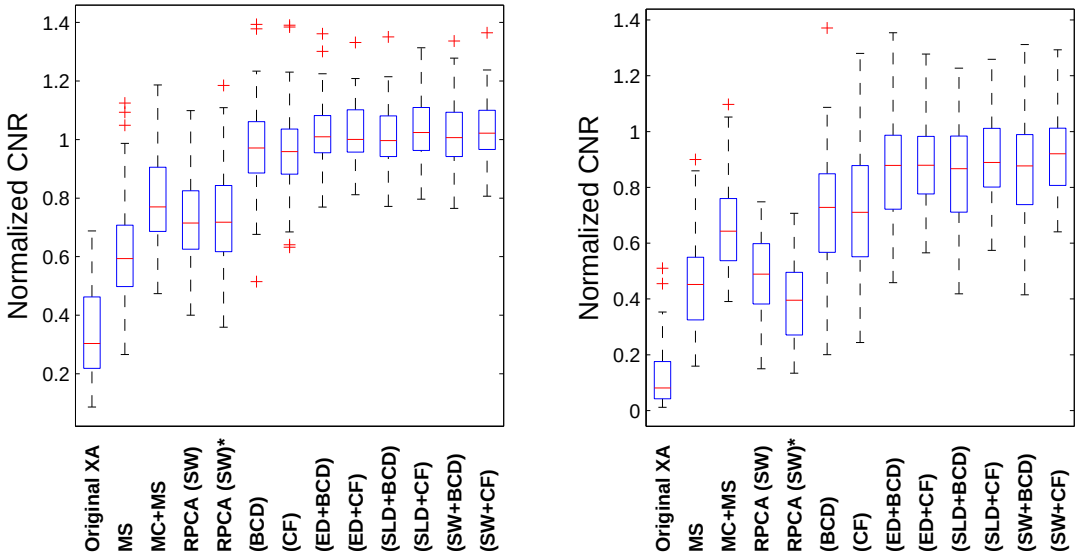

(a) Normalized local CNR

(b) Normalized global CNR

Figure 7: The boxplot that compares various methods on their performance of layer separation. In these diagrams, the CNR values of the mentioned methods are normalized by dividing their CNR by the CNR obtained with the RPCA method.

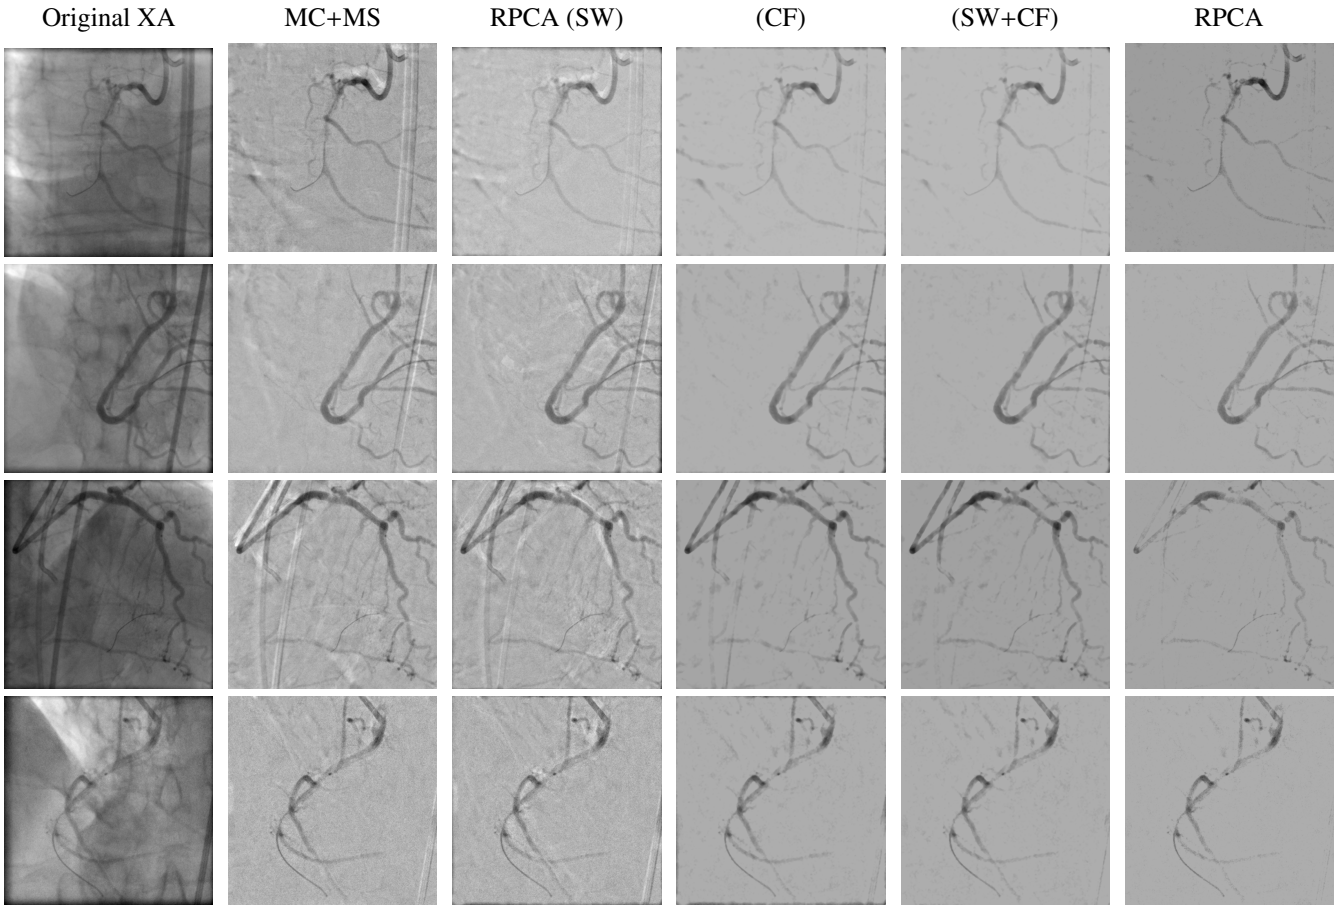

Figure 8: Comparison of five different layer separation methods on four example XA sequences. One representative frame is selected from each sequence to visualize the results. Row 1-4 show four sequences. Column 1 is original XA, column 2-6 are the separated foreground (vessel layer) obtained from MC+MS, RPCA (SW), OR-PCA (CF), OR-PCA (SW+CF) and RPCA.

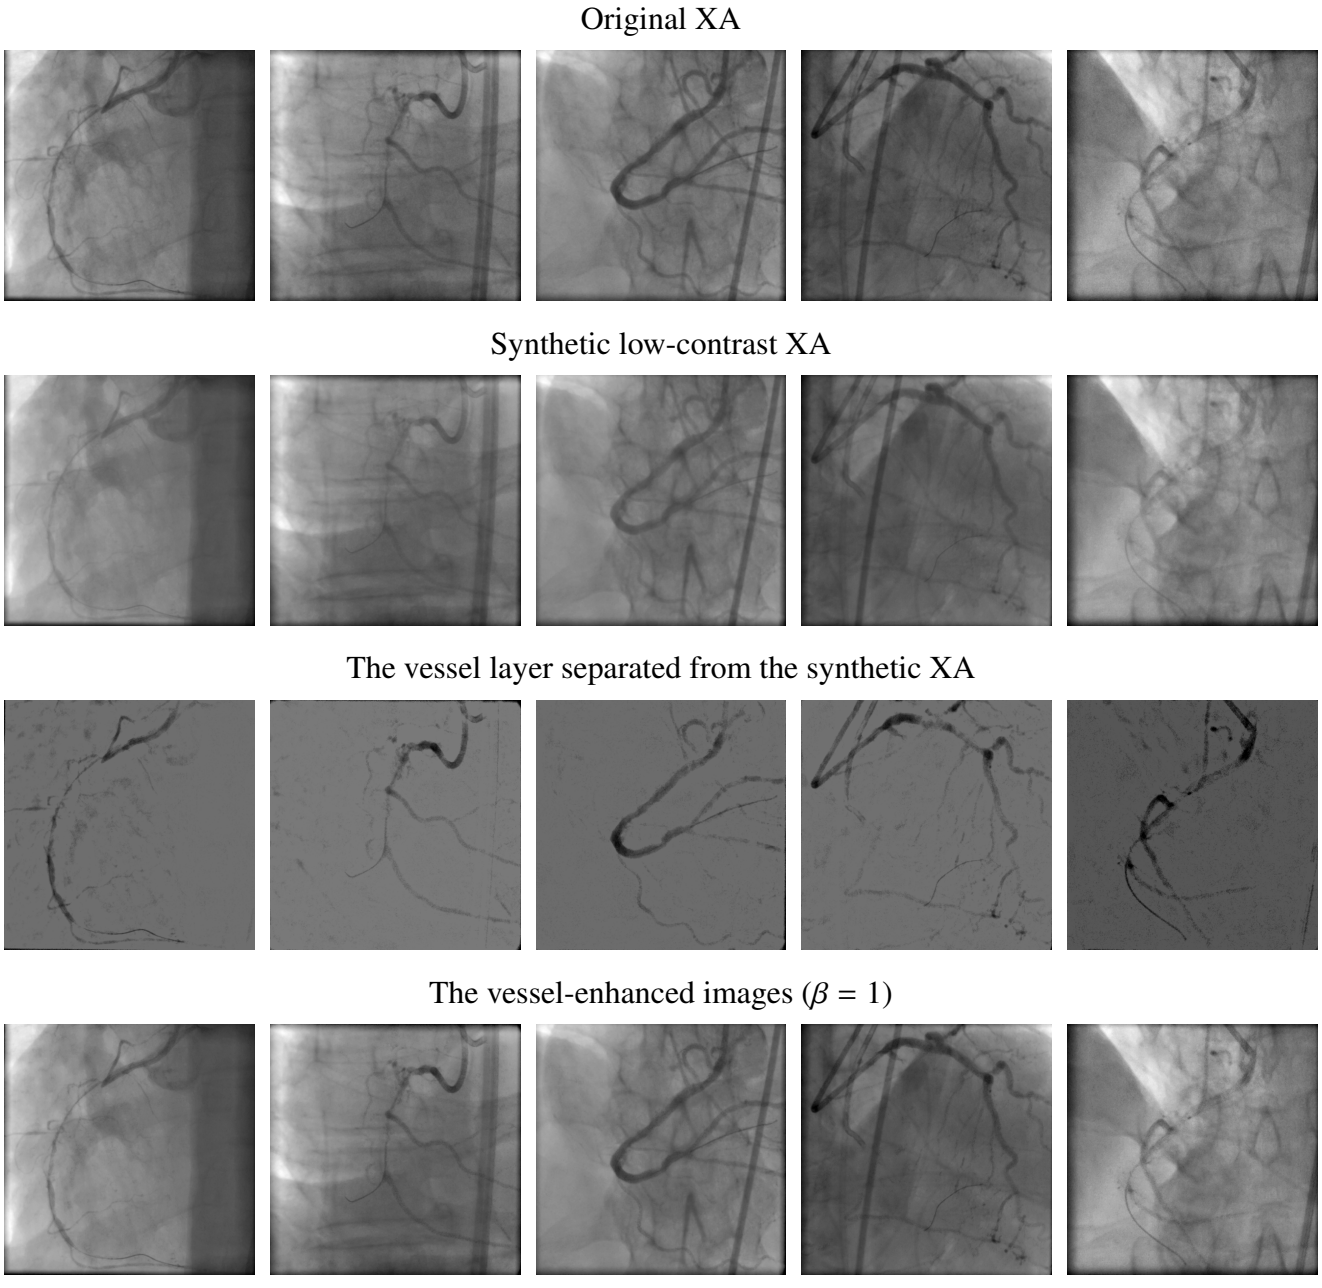

Figure 9: Five examples of vessel enhancement on synthetic low-contrast XA images (Column 1-5).

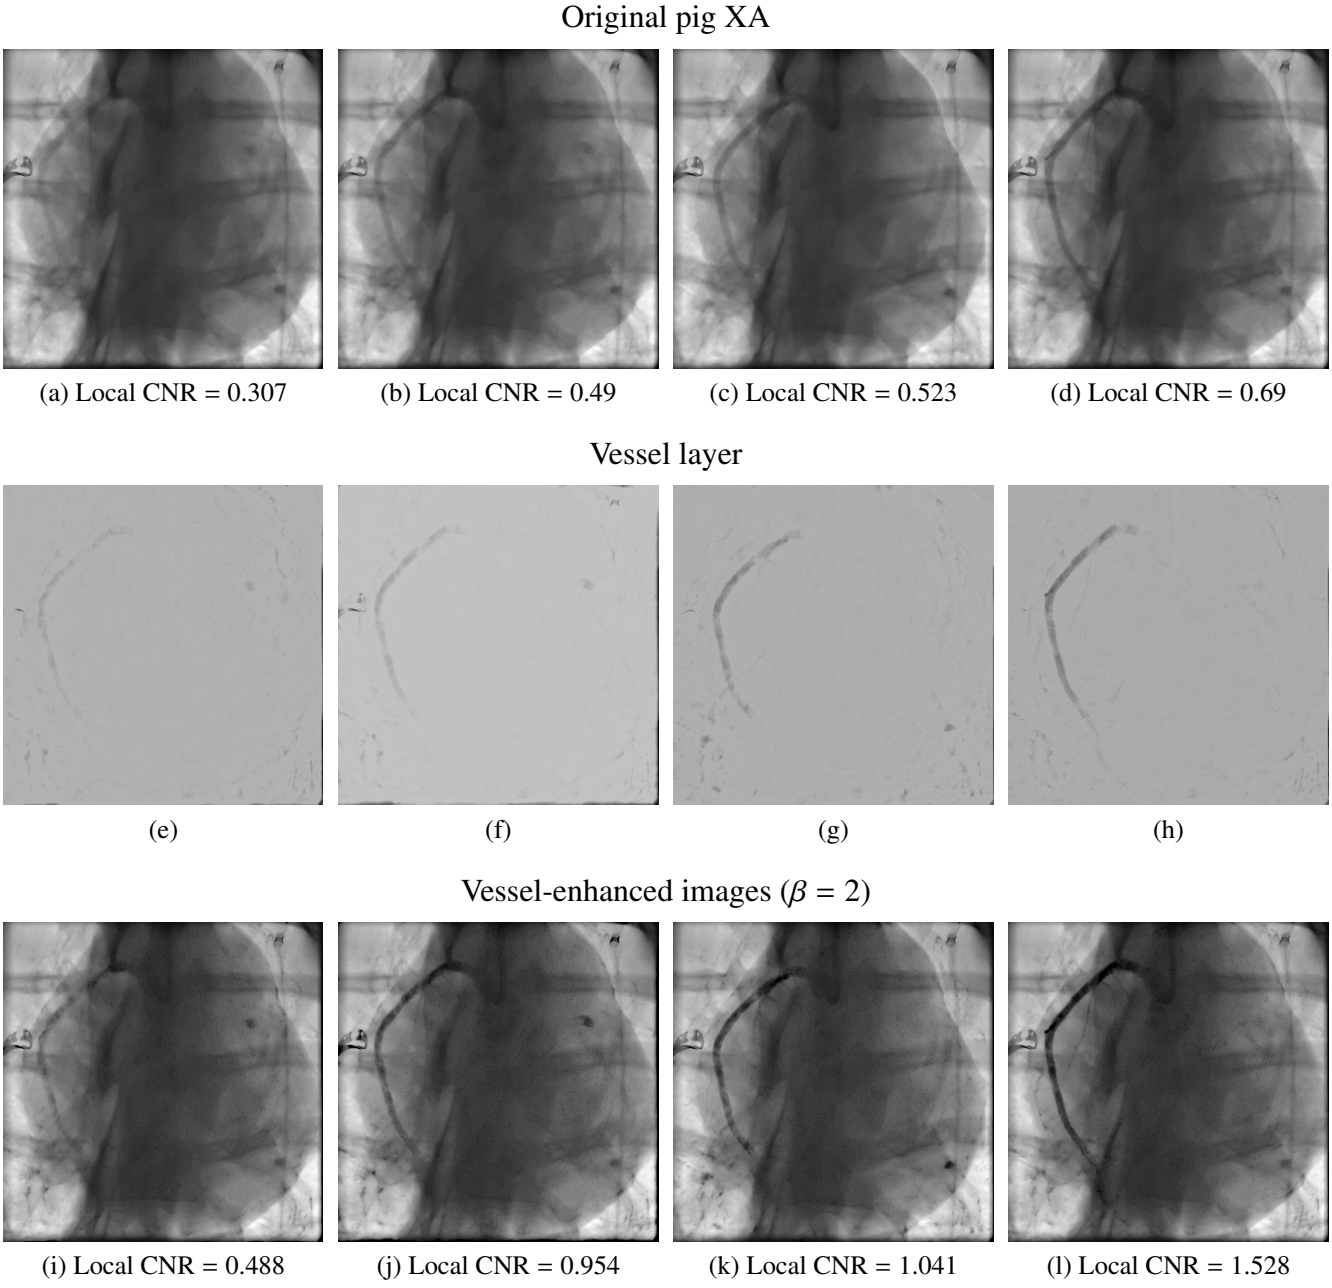

(a) Local CNR = 0.307 (b) Local CNR = 0.49 (c) Local CNR = 0.523 (d) Local CNR = 0.69

(e) (f) (g) (h)

(i) Local CNR = 0.488 (j) Local CNR = 0.954 (k) Local CNR = 1.041 (l) Local CNR = 1.528

Figure 10: Vessel enhancement on pig XA sequences with different level of contrast agent. From left to right, the contrast concentration used for the sequence are 25%, 50%, 100% and 75%, respectively, while the local CNR increases. Row 1: the original pig XA sequences. Row 2: the separated vessel layer. Row 3: the vessel-enhanced images (β = 2).

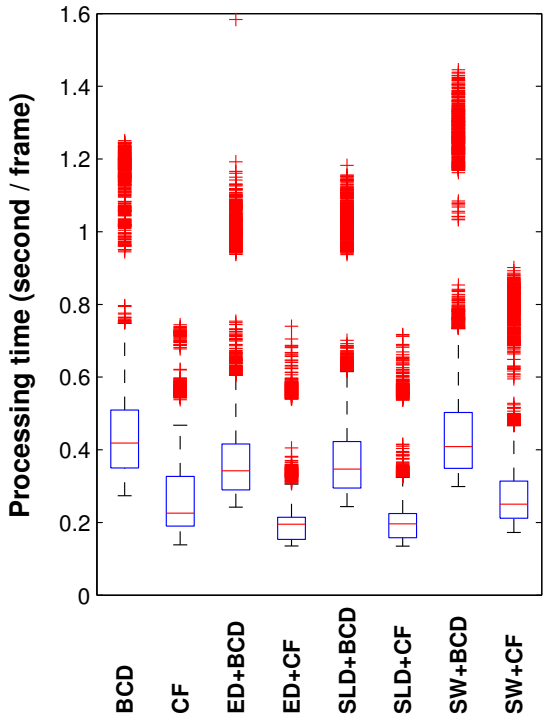

Figure 11: The processing time that each method that uses OR-PCA needs for layer separation. The per frame processing time (second) of every single frame in the whole dataset is shown as box plots.

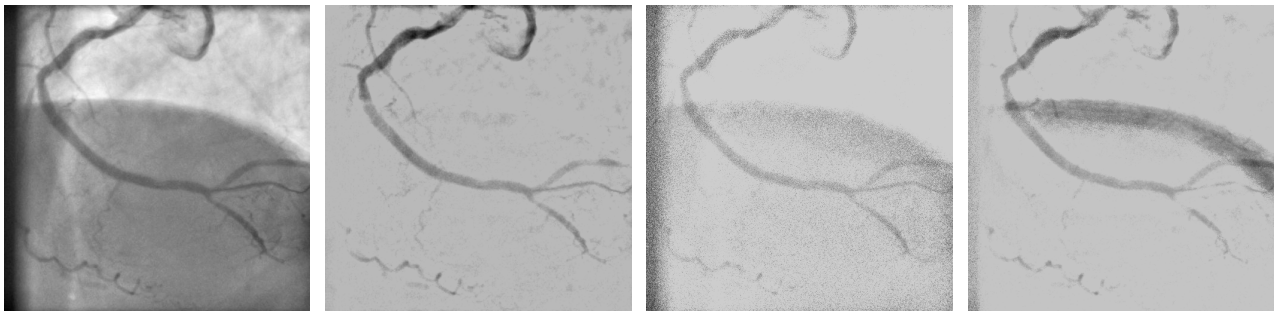

(a) (b) (c) (d)

Figure 12: An example that compares the cases without removing breathing layer before the OR-PCA operation to the one resulted from the proposed method: (a) original image; (b) the vessel layer obtained with OR-PCA (BCD) using the proposed method and parameters; (c) the vessel layer obtained without the separation of breathing layer using OR-PCA (BCD) with the same parameters as the case in (b); (d) the vessel layer obtained without removing breathing structures using a higher r value ($r = 50$) for OR-PCA (BCD). Strong artefacts due to breathing motion can be observed in (c) and (d).