Second pass metal artifact reduction for moving and static objects

Dirk Schäfer, William van der Sterren, Michael Grass

***Abstract*— Metal artefacts such as streaks and shading in C-arm cone beam CT (CBCT) scans are often *not* caused by "photon starvation", i.e. there is still information in the metal shadow. Instead, the artefacts are caused by inconsistent line-integral measurements between different projections (e.g. lateral and anterior-posterior). These inconsistencies have two main causes: First, inconsistent removal of scatter or other imaging deficiencies for rather big *static* high-contrast objects like screws and metal hip implants. Second, motion inconsistencies for rather small *moving* high-contrast objects like contrast filled vessels or catheters.**

A newly developed two-pass method based on the subtraction of the registered metal shadow that is filtered with a matched 2D filter is presented. Good results in the case of moving *and* static 'metal' objects are achieved.

I. Introduction

A variety of methods have been investigated to reduce metal artifacts originating from static objects in X-ray computed tomography. A comprehensive survey as well as a qualitative and quantitative comparison of 12 methods has been performed by Mouton et al. [4]. Simple linear 1D sinogram interpolation [2] is known to wipe out information in the proximity of metal objects but nevertheless proved to be quite effective when measuring the Normalized Root Mean Squared Error (NRMSE) and very efficient compared to more sophisticated methods [4]. Iterative multi-pass methods with more than one forward projection and two back-projections may improve with respect to image quality but increases computational burden beyond interventional use.

Motion is more likely to occur due to rather long scanning times in the interventional setting when compared to diagnostic CT, and we differentiate between moving and static metal objects. In case of rather small *moving* high-contrast objects like contrast filled vessels, the motion inconsistencies result in *high frequency* artifacts. In the case of *static* (non-moving) metal objects, *low frequency* metal streaks and shading spoil the soft tissue visualization, but high contrast bone structures in the proximity of metal are still visible in relevant level/window settings in the original image.

In this paper, we develop a unified second pass method for static and moving objects for interventional use.

II. Methods

The general workflow of second-pass metal artifact correction is as follows and also sketched in Fig. 2.

1) Perform a standard first pass reconstruction.
2) Segment the first-pass volume using a single threshold (approx. 1500 HU) to select only metal voxels.
3) Forward-project the metal voxels to determine the position of the metal shadow in each (line-integral) projection.
4) Correction of the metal shadow: registration, background estimation, filtering, subtraction.
5) Perform a reconstruction using the corrected projections (second pass).
6) Re-insert the segmented metal voxels from the first pass into the second pass volume.

DS and MG are with Philips Research Hamburg, Röntgenstraße 24–26, 22335 Hamburg, Germany. WvdSt is with Philips Healthcare, Veenpluis 6, 5684 PC Best, The Netherlands. corresponding author: dirk.schaefer@philips.com

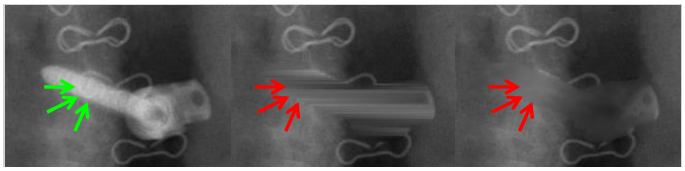

Fig. 1. High contrast structure clearly visible in metal shadow of screw (green arrows, left), linear 1D interpolated projection (middle), linear 2D (right). The interpolation erases the structures in the metal shadow.

A. Motion estimate by registration of metal shadow

The measured projection data is denoted $g(u, v, \lambda)$, at pixel u, v and source position λ. The corresponding forward projection of the segmented metal volume gives the metal shadow $M_0(u, v, \lambda)$.

For every pixel $g(u, v, \lambda)$, where $M_0 > 0$ a patch of the given size $p_u \times p_v$ is generated. A search window of the size $s_u \times s_v$ is created around the corresponding pixel. The normalized cross-correlation (NCC) between the metal shadow patch and the original projection is calculated at each possible location in the search window. The search starts in the middle without a shift of the patch and gradually increases the shift towards the outer borders. For the best match or the first one where NCC> 0.9, the metal shadow patch is added on the resulting registered shadow M_1 at the corresponding positions. Finally, every pixel in M_1 is normalized with the number of added patches.

B. Background estimation

The metal shadow has to be matched to the metal structure within the projection. Therefore, a background estimate is

DOI: 10.12059/Fully3D.2017-11-3202033

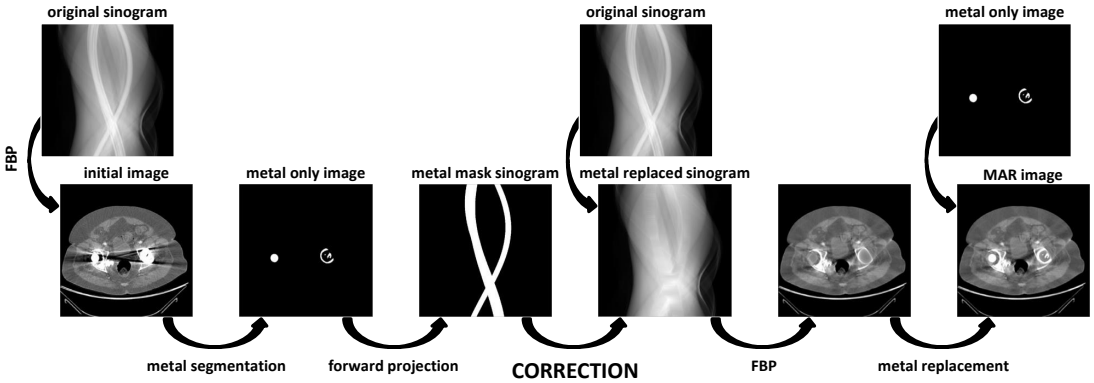

Fig. 2. Schematic algorithm of second pass MAR.

needed to calculate a matched filter for the motion compensated metal mask.

The metal shadow pixels are interpolated using a 2D linear approach. A pixel is defined as border if any of its 8 neighboring pixels $(u \pm 1, v \pm 1)$ is segmented. For each segmented pixel (u_i, v_i) a 2D perimeter $\mathcal{U}(\sigma) = \{(u,v) \in \mathbb{N}^2 |\ |u-u_i| = \sigma \ \vee \ |v-v_i| = \sigma \ \}$ with iteratively increasing radius σ is evaluated, and all border pixel p_i on this perimeter are collected for interpolation. If the ratio $N_{\mathrm{Border}}(\sigma)/N_{\mathrm{Perimeter}}(\sigma)$ of collected border pixel N_{Border} to the current perimeter size $N_{\mathrm{Perimeter}}$ exceeds the threshold Φ_{Ipol}, the line integral value for the segmented pixel is interpolated by a weighted sum of the border pixel values:

$$g_B(u,v,\lambda) \quad = \quad \frac{1}{\mathcal{N}} \sum_{i=1}^{N_{\mathrm{Border}}} g(p_i) \cdot \frac{1}{\sigma(p_i) - \chi} \qquad (1)$$

$$\text{if} \quad \frac{N_{\mathrm{Border}}(\sigma)}{N_{\mathrm{Perimeter}}(\sigma)} > \Phi_{\mathrm{Ipol}} \qquad (2)$$

$$\text{with:} \quad \mathcal{N} = \sum_{i=1}^{N_{\mathrm{Border}}} \frac{1}{\sigma(p_i) - \chi},$$

where $\sigma(p_i)$ provides the radius of the perimeter to which the border pixel p_i belongs. χ is an additional weighting factor to control the relative weight of close and distant border pixels. The default values are $\Phi_{\mathrm{Ipol}} = 0.5$ and $\chi = 0$. An example of this background signal estimated by linear 2D interpolation is shown in Fig. 1.

C. Matched 2D filtered subtraction

Subtraction of the metal shadow from the original projection to maintain overlaid information in the line-integral projection may only be successful if the forward projected metal shadow line-integrals are consistent with the measured line-integral data. There are several effects violating this condition:

1) The object moved outside or to the edge of the metal shadow. Typically, in this case the initial 3D segmentation of the object also suffers from blurring. The compensation for the motion may be handled by a registration of the metal shadow (see Sec. II-A)
2) After registering the metal shadow to the measured data, typically the shape corresponds quite well but the quantitative line-integral values do not match. Therefore, the forward projected metal shadow has to be adjusted to the measured data before subtraction.

This adjustment is achieved by applying a matched 2D filter to the registered metal shadow M_1. Afterwards, the registered and filtered metal patch is subtracted from the measured line-integral data. Let $\mathbf{g} \in g(u,v,\lambda)$ denote a 2D patch of the measured data and $\mathbf{b} \in g_B(u,v,\lambda)$ the corresponding patch of the estimated background. The registered metal patch $\mathbf{r} \in M_1(u,v,\lambda)$ is filtered according to:

$$||(\mathbf{g}-\mathbf{b}) - \mathbf{A}\ \mathbf{r}||^2 \ - \ \lambda\ ||\mathbf{A}||_F^2 \stackrel{!}{=} \min \qquad (3)$$

where $\mathbf{A}$ is the matrix formulation of a 2D convolution filter. F denotes the Frobenius matrix norm and λ is a scalar to adjust the regularization strength. The padding pixels needed for the convolution are added to the patch size, i.e. for the example in Fig. 3, the 5×5-filter is optimized on a $(7+2) \times (7+2)$ patch.

The filter coefficients are computed by setting their partial derivatives to zero. The computation can be done analytically and involves a matrix inversion. The derivation is omitted here for the sake of brevity. To avoid strong spatial fluctuations of the filter coefficients and to capture the general imaging deficiencies, a regularization term $\lambda\ ||\mathbf{A}||_F^2$ is added. This 2D filtered subtraction method is illustrated in Fig. 3. The filter size should be smaller than the used patch sizes to allow an adaptation to the main blurring features in this region of the projection. The general behavior of the filter can be seen in the middle image of Row C of Fig. 3 (i.e. $\lambda = 0.01$): The filter coefficients have high values along the diagonal from top left to bottom right following the main direction of the catheter in this patch. The filter coefficients in the bottom left and top right corner are negative leading to a sharpening of the blurred metal shadow.

D. Linear 1D with frequency-split

For comparison, results are also presented with the frequency-split method [3], which is shortly presented here.

The standard initial reconstruction f^{orig} is segmented by thresholding to get a 3D metal volume f^{metal}. Forward projection, linear interpolation of the metal shadow and reconstruction of the corrected sinogram yields the second pass

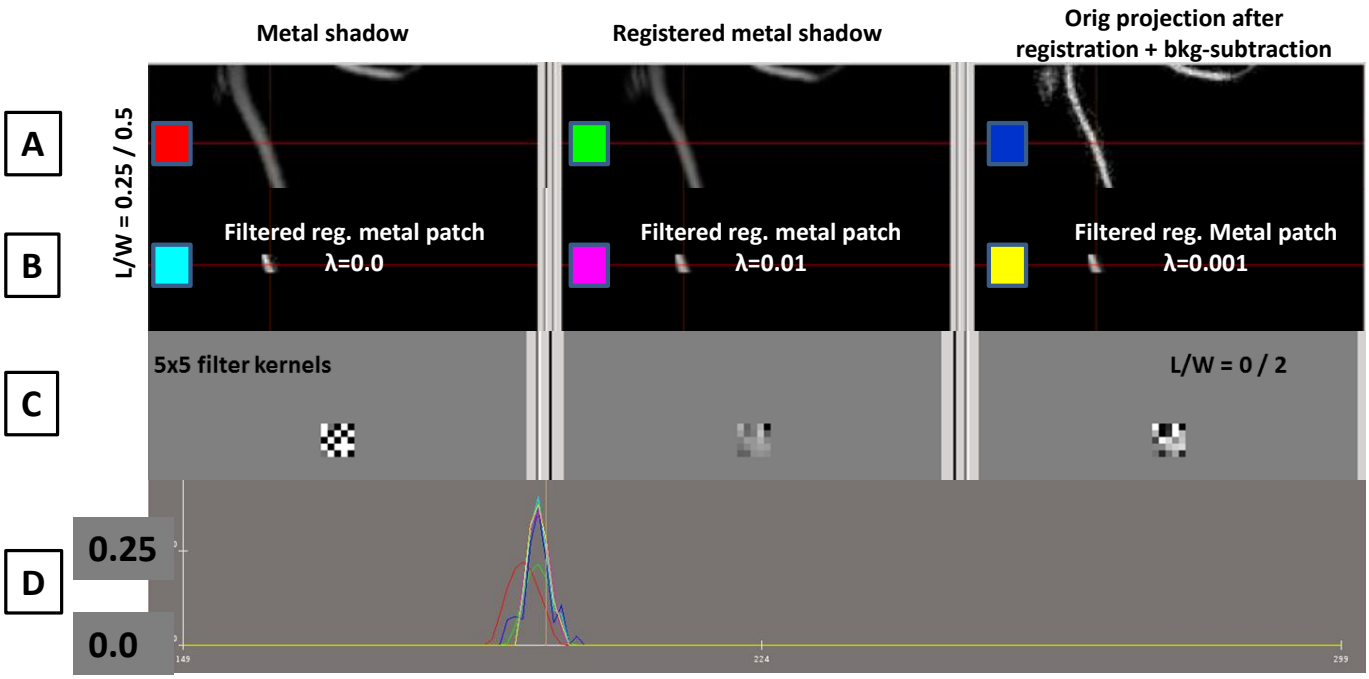

Fig. 3. Illustration of the 2D filtered subtraction method:

Row A: Metal shadow M_0, registered metal shadow M_1 and original projection after background subtraction g_B.

Row B: Filtered 7×7 registered metal patches with different regularization values λ using the 5×5-filters displayed in Row C.

Row C: Matched 5×5-filters. Not to scale! Magnified for better visualization.

Row D: Cross-sectional plots through the images shown in Row A+B. The colors of the plots correspond to the colors in Row A+B.

volume f^{lin1D}. The high frequencies from f^{orig} are computed via subtraction of a low-pass filtered version of the volume:

$$f_{hp}^{orig} = f^{orig} - f^{orig} * G(\sigma_f), \tag{4}$$

where $G(\sigma_f)$ is a 3D Gaussian filter with standard deviation σ_f. The low-pass and high-pass filtered versions of the second pass volume are computed accordingly, using the same σ_f:

$$\begin{aligned} f_{lp}^{lin1D} &= f^{lin1D} * G(\sigma_f) \\ f_{hp}^{lin1D} &= f^{lin1D} - f_{lp}^{lin1D}. \end{aligned}$$

Spatial weights w are computed from the segmented metal volume with very large Gaussian filters:

$$w = G(\sigma_{w2}) * \Theta\left(G(\sigma_{w1}) * f^{metal}\right), \tag{5}$$

where Θ maps all smoothed metal voxel to one and the background to zero. The subsequent convolution with $\sigma_{w2} < \sigma_{w1}$ guarantees a smooth transition of the weighting function from one to zero. The final frequency-split volume is the computed according to:

$$f^{fs} = f_{lp}^{lin1D} + w\, f_{hp}^{orig} + (1 - w)\, f_{hp}^{lin1D}. \tag{6}$$

III. Results

The matched filter method has been developed, tested and evaluated on a data-base consisting of 12 data sets (5 liver angiographies, 3 needle ablations, 1 spine data set, 2 abdominal scans and 1 hip data set). All scans are acquired with XperCT protocols using a Philips Allura Xper FD20 system (Philips Healthcare, Best, The Netherlands).

Results are shown in Fig. 4 for a spine case with 570 projections and a liver case with 312 projections. For comparison, results are also shown for the standard linear 1D interpolation method (lin1D) [2] and the lin1D-method combined with the frequency-split method (lin1d + fs) [3].

All these three methods use the same 3D metal segmentation with a single threshold of 1500 HU. The lin1D method has been combined with the frequency split method using $\sigma_f = 1$ voxel, $\sigma_{w1} = 15$ voxels and $\sigma_{w2} = 7.5$ voxels.

For the 2D registration 7×7-patches with a search window of 7×7 pixels have been used. The perimeter parameter for the lin2D method was set to $\Phi_{\text{Ipol}} = 0.5$. The 2D filter patches are set to a size of 5×5.

The original reconstruction is shown also in a wider display window to allow better visual assessment.

The frequency split method is motivated by the observation

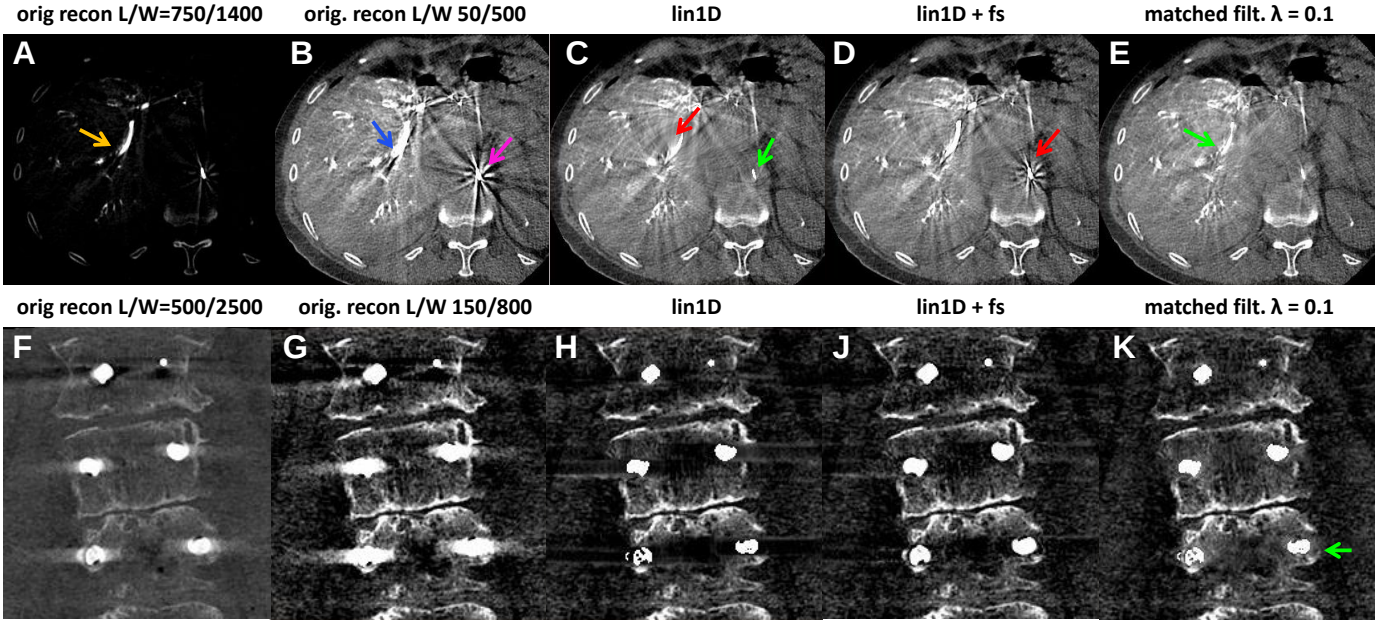

Fig. 4. Trans-axial slice of a liver angiography (top row) and coronal slice of a spine scan with metal screws. Original reconstruction in different windows (column 1-2) and three different MAR methods (column 2-5).

that in case of static (non-moving), "metal" objects, *low frequency* metal streaks and shading spoil the soft tissue visualization, but high contrast structures in the proximity of metal are still visible in other level/window settings in the original image (see orange arrow in Fig. 4A).

In the liver case, this kind of 'static' metal-artefact caused by the high concentration of iodine agent in the liver artery (see blue arrow Fig. 4B) occurs in the same trans-axial slice as 'motion' metal-artefacts (see magenta arrow Fig. 4B) caused by the catheter in the aorta, which is moving due to the pulsatile blood flow.

The lin1D-method is able to remove the streaks around the small catheter (green arrow Fig. 4C), but wipes out all information in the proximity of the large hepatic artery (red arrow Fig. 4C). Adding the frequency-split to the lin1D-method re-introduces the natural texture in the neighborhood of the hepatic artery, but also re-introduces the high-frequency 'motion' metal-artefacts around the catheter (red arrow Fig. 4D). The proposed matched filter method nicely recovers the small vessel next to the hepatic artery (green arrow Fig. 4E) ***and*** removes the streaks around the catheter.

Good performance is also obtained in the 'static'-only case of the spine scan. The faint bone structure close to the screw can be recovered for the matched filter method (green arrow Fig. 4K).

IV. Summary and Conclusion

The standard two-pass lin1D interpolation method wipes out structures close to larger metal objects and produces artificially flat regions. In the case of rather small, slightly moving high-contrast objects (catheters, vessels, needles) the 'metal artefacts' also consist of *high frequency* streaks close to the metal voxels. The standard lin1D interpolation method effectively removes most streaks, but the frequency split method re-introduces the high frequency artefacts.

The presented two-pass method based on the subtraction of the registered metal shadow that is filtered with a matched 2D filter yields good results in the case of moving *and* static 'metal' objects. This is exemplarily shown for a liver angiography, where the subtraction method recovers the small vessels next to the large 'metal' object and at the same time the streak artifacts from the moving catheter are strongly reduced. More complex sinogram interpolation methods could be investigated, e.g. linear 3D interpolation [5] or interpolation based on the structure tensor of the sinogram [1]. However, the complexity increases significantly, and a parallel independent processing of the projections is not possible anymore.

Acknowledgment

The authors gratefully acknowledge data acquisition of the liver XperCT by Prof Geert Maleux, University Hospitals Leuven and the spine XperCT by the imaging research team at Cincinnati Children's Hospital.

References

[1] M. Bertram, J. Wiegert, D. Schäfer, T. Aach, G. Rose, Directional view interpolation for compensation of sparse angular sampling in cone-beam CT, IEEE Trans Med Imaging, Vol. 28(7), 1011-1022, 2009.

[2] W. A. Kalender, R. Hebel, and J. Ebersberger, Reduction of CT artifacts caused by metallic implants., Radiology, Vol. 164, no. 2, 576, 1987.

[3] E. Meyer, R. Raupach, M. Lell, B. Schmidt M. Kachelriess, Frequency split metal artifact reduction (FSMAR) in computed tomography, Med. Phys. 39 (4), 1904-1916, 2012.

[4] A. Mouton, N. Megherbi, K. van Slambrouck, J. Nuyts, T. Breckon, An experimental survey of metal artefact reduction in computed tomography, Journal of X-Ray Science and Technology, 2013, Vol. 21(2), 193-226, 2013.

[5] D. Prell, W. Kalender, Y. Kyriakou, Development, implementation and evaluation of a dedicated metal artefact reduction method for interventional flat-detector CT, The British Journal of Radiology, 83, 1052–1062, 2010.