

# DIGITAL CRADLE REMOVAL IN X-RAY IMAGES OF ART PAINTINGS

Rujie Yin<sup>1</sup>, David Dunson<sup>1</sup>, Bruno Cornelis<sup>2</sup>, Bill Brown<sup>3</sup>, Noelle Ocon<sup>3</sup>, Ingrid Daubechies<sup>1</sup>

1: Duke University, USA; 2: Free University Brussels (VUB), Belgium; 3: North Carolina Museum of Art, USA

## ABSTRACT

We introduce an algorithm that removes the deleterious effect of cradling on X-ray images of paintings on wooden panels. The algorithm consists of a three stage procedure. Firstly, the cradled regions are located automatically. The second step consists of separating the X-ray image into a textural and image component. In the last step the algorithm learns to distinguish between the texture caused by the wooden cradle and the texture belonging to the original painted wooden panel. The results obtained with our method are compared with those obtained manually by best current practice.

**Index Terms**— texture, art, painting, cradle.


## 1. INTRODUCTION

### 1.1. The digital cradle removal problem

Between the 12th to the 17th century, paintings in Europe were mostly created on wooden panels consisting of solid wooden boards. Until about 1950, cradling was a common restoration technique used by conservators to remediate or prevent structural or insect damage. In the process, the panels were first thinned, and then strengthened by (permanently) attaching to their backs hardwood lattices called *cradles*.

On the other hand, modern-day museum conservators and art historians rely heavily on X-ray imaging to study artist technique, painting fabrication, and the physical condition of the painting (e.g. cracks in the painted surface or wooden support). X-ray images of cradled works contain highly visible horizontal and vertical “bars”, caused by the higher X-ray absorption in the thicker wood layer due to the cradle (see Fig. 1). This obstructs the “reading” of the X-ray image by art experts, who would welcome a (semi-)automated procedure to remove these artifacts from the X-ray images.

In recent years, image processing algorithms have been developed to analyze high-resolution digital images of art paintings, e.g. to classify styles [1], to detect (and inpaint) cracks [2], or to remove canvas artifacts [3]. In [2] cradling artifacts in X-ray images constituted a hindrance, but since only a small portion of the painting needed to be processed, the corresponding artifacts were easy to isolate, and were then



**Fig. 1: Cradle of old paintings on wood panel:**

*Top left:* a sample patch from the X-ray image of a cradled painting panel, illustrating the artifacts caused by the cradling, with enlarged detail (showing cradle wood grain). *Top right:* sketch of the lattice of crossed perpendicular wooden slats: slats parallel to the panel wood grain are glued first; transverse slats are slotted in. *Left:* indication, on the sample patch, of the partition into different domains used in §2.

successfully removed by a *morphological component analysis* (MCA) [4] pre-processing step. To our knowledge, there is no algorithm yet to remove cradle artifacts from extended portions of the X-ray images of panel paintings; experienced conservators still do this manually, in Adobe Photoshop™ – a time-consuming process, applied to only a limited number of paintings among the large museum collections.<sup>1</sup>

### 1.2. Our Approach

We formulate the digital cradle removal as a source separation problem. According to the physical composition of the wooden panel and the cradle, a raw X-ray image  $\mathcal{I}$  can be viewed as the superposition of a *panel* image  $\mathcal{I}_p$ , containing all the information from the painting on the panel and the wooden panel itself, and a *cradle* image  $\mathcal{I}_c$  induced by the cradling alone. We decompose  $\mathcal{I}_c$  into two parts:  $\mathcal{I}_c^0$ , piecewise constant on the cradle support, corresponding to the cradle thickness<sup>2</sup>, and  $\mathcal{I}_c^{wg}$ , a texture component reflecting the

<sup>1</sup>Alternative X-ray techniques involve filling of the voids in the cradle lattice by pressing putty into them, in an attempt to even out the X-ray exposure, resulting in additional risk to the paintings.

<sup>2</sup>This assumes that the cradle has sharp boundaries; for less sharp boundaries, e.g. when the cradling supports have rounded corners,  $\mathcal{I}_c^0$  is not con-

We gratefully acknowledge support from SAMSI (RY), FWO (BC), NSF (DD and ID), . . .

cradle’s wood grain. The panel image can also be decomposed, into a “cartoon” part  $\mathcal{I}_p^{cart}$  (i.e. the different “color” planes and their boundaries), and a texture part  $\mathcal{I}_p^{text}$  (i.e. brushstroke details as well as the wood grain of the panel). Accordingly, the decomposition of the whole image into “texture” and “cartoon” parts is  $\mathcal{I} = (\mathcal{I}_p^{cart} + \mathcal{I}_c^0) + (\mathcal{I}_p^{text} + \mathcal{I}_c^{wg})$ . We use a three-stage approach to extract  $\mathcal{I}_c^0$  and  $\mathcal{I}_c^{wg}$  from  $\mathcal{I}$ .

We first estimate the locations of the cradles; to detect the cradle slat angles (close to but not exactly  $0^\circ$  or  $90^\circ$ ), we use the Radon transform (§2.1). We also estimate the intensity difference between each single cradled region and the surrounding panel-only region; crossing areas (where horizontal and vertical slats overlap) need to be treated separately, because they correspond to a different intensity change. We use these estimates to define the function  $\mathcal{I}_c^0$ , constant on the interior of each single or crossing cradled region (with sharp or simply modeled smoothed edges), and subtract it from  $\mathcal{I}$ . We define this intermediate result as  $\mathcal{I}^{intrm} := \mathcal{I} - \mathcal{I}_c^0$ .

The second stage uses MCA, inspired by the digital canvas removal [5], to separate  $\mathcal{I}^{intrm}$  into  $\mathcal{I}^{text}$  and  $\mathcal{I}_p^{cart}$ . To decompose  $F = A + B$  into its constituents, MCA uses that  $A$  and  $B$  each have sparse decompositions into different standard dictionaries,  $D_A$  and  $D_B$ . In our case, the textures in  $\mathcal{I}_c^{wg}$  and  $\mathcal{I}_p^{text}$  have sparse representations in a dictionary of high frequency curvelets or shearlets; for the complimentary  $\mathcal{I}_p^{cart}$ , we use complex dual-tree wavelets, the same dictionary as in [5] for the image “content”.

However, our problem is more complex than in [5], because  $\mathcal{I}_p$ , which we want to recover, has components in both constituents; the third stage in the algorithm deals with this complication. Let  $\mathcal{D}_p$  and  $\mathcal{D}_c$  be the *panel-only* and the *cradled* domains (see Fig.1;  $\mathcal{D}_c = \mathcal{D}_c^{int} \cup \mathcal{D}_c^{bd}$ ). Then  $\mathcal{I}^{text}|_{\mathcal{D}_p} = \mathcal{I}_p^{text}|_{\mathcal{D}_p}$ , since  $\mathcal{I}_c^{wg} = 0$  on  $\mathcal{D}_p$ . On  $\mathcal{D}_c$ , we have  $\mathcal{I}^{text}|_{\mathcal{D}_c} = \mathcal{I}_c^{wg}|_{\mathcal{D}_c} + \mathcal{I}_p^{text}|_{\mathcal{D}_c}$ . We shall use a variation of a sparse Bayesian factor model (see §2.2) to obtain  $\mathcal{I}_c^{wg}$ : we learn dictionaries for the panel and cradle textures in the feature (high-frequency curvelet or shearlet) space, using our observation of the different behavior of  $\mathcal{I}^{text}$  on  $\mathcal{D}_p$  and  $\mathcal{D}_c$ ;  $\mathcal{I}_c^{wg}$  is then approximated by MCMC sampling. In this particularly challenging unsupervised setting, the Bayesian approach avoids tuning parameters and the need for cross-validation.

## 2. CRADLE REMOVAL ALGORITHM

### 2.1. Radon transform and location estimation

Suppose, for instance, that the X-ray image  $\mathcal{I}$  on  $\mathcal{D} \subset \mathbb{R}^2$  contains one horizontal cradle slat.<sup>3</sup> To even out fine scale noise and texture effects, we first apply an elongated vertical Haar transform,  $\delta\mathcal{I}_{i,j} = \sum_{l=0}^{L-1} \mathcal{I}_{i-l,j} - \sum_{l=1}^L \mathcal{I}_{i+l,j}$ , ( $L$

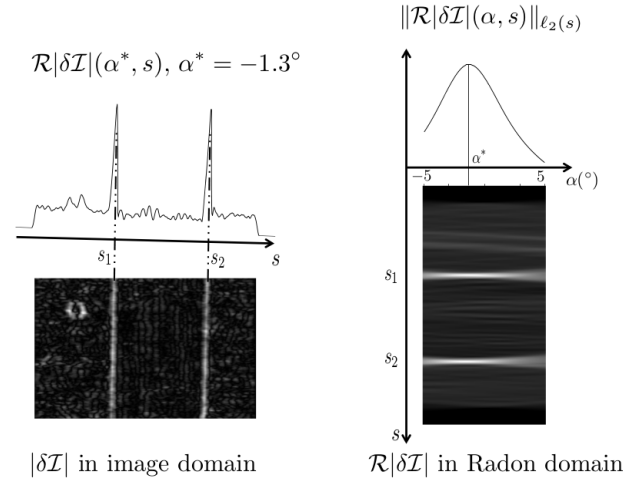
stant. However, it can then be modeled or approximated, as shown below.

<sup>3</sup>This assumption can be made without loss of generality: since cradle slats in the same direction are well separated, it is easy to consider patches that have at most one cradle slat in each direction. For a vertical slat, we just perform the same procedure after rotating the patch by  $90^\circ$ .

fixed), followed by a Radon transform  $\mathcal{R}$ :

$$\mathcal{R}|\delta\mathcal{I}|(\alpha, s) = \int_{-\infty}^{\infty} |\delta\mathcal{I}((s \cos \alpha, s \sin \alpha) + t\tau_\alpha)| dt \quad (1)$$

where  $\tau_\alpha = (-\sin \alpha, \cos \alpha)$ , with  $\alpha \in [-\theta, \theta]$ . (See Fig. 2.)



**Fig. 2: Radon transform to locate the cradle:**  $\mathcal{I}$  here is the left part of Fig. 1. *Left:* peaks in  $\mathcal{R}|\delta\mathcal{I}|(\alpha^*, s)$ , for the optimal angle  $\alpha^*$ , indicate the boundaries of the cradle. *Right:* a constant-angle section of the Radon transform of  $|\delta\mathcal{I}|$  attains maximal  $\ell_2$  energy for  $\alpha = \alpha^*$ .

For any  $\alpha$ ,  $\|\mathcal{R}|\delta\mathcal{I}|(\alpha, \cdot)\|_{\ell_1(\mathbb{R})} = \int_{\mathbb{R}^2} |\delta\mathcal{I}(x, y)| dx dy = \|\delta\mathcal{I}\|_{\ell_1(\mathbb{R}^2)}$ . On the other hand, for the optimal  $\alpha^*$ ,  $\mathcal{R}|\delta\mathcal{I}|(\alpha^*, s)$  is more concentrated in  $s$  than for other  $\alpha$ . Thus the optimal  $\alpha^*$  can be found by maximizing the  $\ell_2(\mathbb{R})$ -norm of  $\mathcal{R}|\delta\mathcal{I}|(\alpha, \cdot)$ ; we set the angle of the horizontal cradle slat as

$$\alpha^* = \arg \max_{\alpha \in [-\theta, \theta]} \|\mathcal{R}|\delta\mathcal{I}|(\alpha, \cdot)\|_{\ell_2(\mathbb{R})}.$$

To estimate sharp boundaries we back-project the hard-thresholded Radon data  $\mathcal{H}_\lambda(\mathcal{R}|\delta\mathcal{I}|(\alpha^*, s))$  to the image domain. In the case of smooth boundaries, the profiles are modeled by smooth shape (rather than step) functions  $\{\Phi_a\}_{a \in A}$ , with a uniform parameter choice (estimated from  $\mathcal{R}|\delta\mathcal{I}|(\alpha^*, \cdot)$ ) for each transition region. The intensity difference across the cradle boundaries is also estimated from  $\mathcal{R}|\delta\mathcal{I}|(\alpha^*, \cdot)$ .

### 2.2. Additive factor model

The second stage is a fairly standard application of MCA (see [4]) to separate out  $\mathcal{I}^{text}$  from  $\mathcal{I}^{intrm} = \mathcal{I} - \mathcal{I}_c^0$ . The input to the third stage, explained here, consists of  $\mathcal{I}^{text}$  and the partition of the image domain  $\mathcal{D} = \mathcal{D}_p \cup \mathcal{D}_c$  obtained in the first stage. Let  $\mathcal{S}$  be the linear curvelet (or shearlet) transform that encodes  $\mathcal{I}^{text}$  into a low dimensional feature vector of curvelet (or shearlet) coefficients at every pixel in

$\mathcal{D}$ . Since the wood grain texture of the cradle is highly directional, we can further reduce the dimensionality by considering only curvelets or shearlets within a (small) angle range  $\Theta = [-\theta, \theta]$ , such that  $\mathcal{I}_c^{text} = \mathcal{S}_\Theta^* \mathcal{S}_\Theta \mathcal{I}_c^{text}$ . ( $\mathcal{S}_\Theta^*$  is the adjoint operator of  $\mathcal{S}_\Theta$ .) Set  $Y_p = \mathcal{S}_\Theta \mathcal{I}_c^{text}|_{\mathcal{D}_p} \in \mathbb{R}^{K \times N}$ ,  $Y_c = \mathcal{S}_\Theta \mathcal{I}_c^{text}|_{\mathcal{D}_c} \in \mathbb{R}^{K \times M}$ , and  $Y = [Y_p, Y_c] = [y_i], i = 1, \dots, M + N$ , where  $N$  and  $M$  are the numbers of panel-only and cradled pixels (samples), and  $K$  is the dimension of the feature space associated to  $\mathcal{S}_\Theta$ .

On  $\mathcal{D}_c$ ,  $\mathcal{I}^{text}$  consists of two parts; accordingly we have  $Y_c = \mathcal{S}_\Theta(\mathcal{I}_c^{wg} + \mathcal{I}_p^{text}|_{\mathcal{D}_c}) = Y_c^* + Y_c^{**}$ ; it then follows that  $\mathcal{I}_c^{wg} = \mathcal{S}_\Theta^* Y_c^*$ . Our task is thus to separate  $Y_c$  into the sum  $Y_c^* + Y_c^{**}$ , where  $Y_c^{**}$  is ‘‘similar to’’  $Y_p$  (since both correspond to the texture content of the panel painting, albeit in different regions of the panel).

We build an additive factor model, with dictionaries (factor loading matrices) for  $[Y_c^{**}, Y_p]$  and  $Y_c^*$  introduced separately. For the panel-only samples  $Y_p = [y_i], i = 1, \dots, N$ , we use the formulation of a sparse Bayesian factor model in [6],

$$y_i = \Lambda \eta_i + \epsilon_i, \quad (2)$$

where  $\Lambda \in \mathbb{R}^{p \times k}$  is the factor loadings matrix,  $\eta_i$  is the  $i$ th latent factor and  $\epsilon_i$  is a residual. The columns  $\lambda_h$  of  $\Lambda$  can be viewed as dictionary elements. In our case, each dictionary element is a specific combination of shearlets or curvelets of different frequencies and directions. Performance can be improved by incorporating increasing sparsity in the dictionary with respect to the index  $h$ ; an appropriate size  $k$  of dictionary is adaptively learned as well as the sparse level. Instead of setting elements of  $\Lambda$  exactly to zero, we put a continuous shrinkage prior, which approximates the exact sparsity setting but has substantial computational improvement.

For cradled samples  $Y_c = [y_j], j = N + 1, \dots, N + M$ , we have the decomposition  $Y_c = Y_c^* + Y_c^{**}$ . Moreover,  $Y_c^{**}$  and  $Y_p$  are both generated from the panel-only sources  $\mathcal{I}_p^{text}|_{\mathcal{D}_c}$  and  $\mathcal{I}_p^{text}|_{\mathcal{D}_p}$ , although  $\mathcal{I}_p^{text}|_{\mathcal{D}_c}$  may be attenuated somewhat because of the strong signal of the cradle. Therefore, assuming that  $Y_p/\rho = [y_i/\rho]_{i=1, \dots, N}$  and  $Y_c^{**} = [y_j^{**}]_{j=N+1, \dots, N+M}$  obey the same distribution, where  $\rho$  is a constant attenuation factor in  $(0, 1]$ , we construct the following additive factor model,

$$y_j = \rho \Lambda \eta_j + \Gamma \xi_j + \epsilon_j; \quad (3)$$

here  $\Lambda$  is the same factor loading matrix as in (2) such that  $y_j^{**} \approx \rho \Lambda \eta_j$ .  $\Gamma$  and  $\xi_j$  are the factor loading matrix and the factor for  $y_j^*$ . The residual  $\epsilon_j$  in (3) has the same distribution as  $\epsilon_i$  in (2). We also put the same priors on  $\Gamma$ ,  $\eta_j$  in (3) as those on  $\Lambda$ ,  $\xi_i$  in (2), but allow different sparsity in  $\Gamma$  and  $\Lambda$ .

To learn (2) + (3) simultaneously, a block Gibbs sampler cycles through the conditionally conjugate posteriors of the parameters; This Gibbs sampler is analogous to the one in [6], where the explicit expressions of conjugate posteriors and further discussion about the construction can be found. However, to achieve the best result, we cut the feedback of cradle

data  $Y_c$  in updating  $\lambda_i$ , so that we implement the conditional posterior of  $\lambda_i$  as  $\pi(\lambda_i | Y_p, \Omega^c)$ , where  $\Omega^c$  stands for all the other parameters in the model; by using panel-only data, the chance of picking up cradled elements as  $\lambda_i$ 's is negligible.

Let  $\Gamma^{(n)}$ ,  $\xi_j^{(n)}$ , and  $\epsilon_j^{(n)}$ ,  $j = N + 1, \dots, N + M$ ,  $n = 1, \dots, N_0$  be the posterior draws of  $\Gamma$ ,  $\xi_j$ , and  $\epsilon_j$  from the above Gibbs sampler. We get the approximation  $\widehat{y}_j^* = \frac{1}{N_0} \sum_{n=1}^{N_0} (\Gamma^{(n)} \xi_j^{(n)} + \epsilon_j^{(n)})$ , and  $\widehat{\mathcal{I}_c^{text}} = \mathcal{S}_\Theta^* \widehat{Y}_c^*$ .

### 2.3. The Algorithm

---

#### Algorithm 1 Digital Cradle Removal Algorithm


---

- 1: **Input:**  $\mathcal{I}$
  - 2: Estimate horizontal cradle location  $\mathcal{D}_c^{hor}$
  - 3: Estimate boundary support  $\mathcal{D}_c^{hor, bd}$ , profile  $\mathcal{B}_h$  and intensity difference  $\delta_h$
  - 4: Compute  $\mathcal{D}_c^{hor, int} = \mathcal{D}_c^{hor} \setminus \mathcal{D}_c^{hor, bd}$  and  $\mathcal{I}_c^0|_{\mathcal{D}_c^{hor}} = \mathbf{1}_{\mathcal{D}_c^{hor, int}} \delta_h + \mathcal{B}_h$
  - 5: Rotate  $\mathcal{I}$  clock-wise  $90^\circ$
  - 6: Repeat step 2 - 3, output  $\mathcal{D}_c^{ver}$ ,  $\delta_v$
  - 7: Rotate  $\mathcal{D}_c^{ver}$  counter clock-wise  $90^\circ$
  - 8: Estimate cross-section region  $\mathcal{D}_c^\times = \mathcal{D}_c^{hor} \cap \mathcal{D}_c^{ver}$
  - 9: Estimate  $\delta_\times$ , compute  $\mathcal{I}_c^0|_{\mathcal{D}_c^\times} = \mathbf{1}_{\mathcal{D}_c^\times} \delta_\times$
  - 10: Compute  $\mathcal{I}_c^0|_{\mathcal{D}_c^{ver}} = (\mathbf{1}_{\mathcal{D}_c^{ver}} - \mathbf{1}_{\mathcal{D}_c^\times}) \delta_v + \mathcal{B}_v$
  - 11:  $\mathcal{D}_c = \mathcal{D}_c^{hor} \cup \mathcal{D}_c^{ver}$ ,  $\mathcal{D}_p = \mathcal{D} \setminus \mathcal{D}_c$
  - 12: Compute  $\mathcal{I}_c^0 = \mathcal{I}_c^0|_{\mathcal{D}_c^{hor}} + \mathcal{I}_c^0|_{\mathcal{D}_c^{ver}} + \mathcal{I}_c^0|_{\mathcal{D}_c^\times}$
  - 13: Do MCA on  $\mathcal{I} - \mathcal{I}_c^0$ , output  $\mathcal{I}^{text}$ ,  $\mathcal{I}_p^{cart}$
  - 14: Compute angle restricted shearlet/curvelet transform,  $Y_c = \mathcal{S}_\Theta \mathcal{I}^{text}|_{\mathcal{D}_c}$ ,  $Y_p = \mathcal{S}_\Theta \mathcal{I}^{text}|_{\mathcal{D}_p}$
  - 15: Run MCMC over the additive factor model (2) + (3), output  $\widehat{y}_j^* = \Gamma^{est} \xi_j^{est} + \epsilon_j^{est}$
  - 16: Compute  $\widehat{\mathcal{I}_c^{text}} = \mathcal{S}_\Theta^* \widehat{Y}_c^*$ ,  $\mathcal{I}_c = \mathcal{I}_c^0 + \mathcal{I}_c^{text}$
  - 17: Compute  $\mathcal{I}_p = \mathcal{I} - \mathcal{I}_c$
  - 18: **Output:**  $\mathcal{I}_p, \mathcal{I}_c, \mathcal{I}_c^0, \mathcal{I}_c^{text}$
- 

The steps listed in Algorithm 1 estimate the cradle image  $\mathcal{I}_c$  and remove it from  $\mathcal{I}$  as explained above. Steps 2 - 11 are the first stage of constant intensity difference  $\mathcal{I}_c^0$  approximation (possibly with smooth transition); steps 2 - 3 use the Radon transform technique described in §2.1. Usually,  $\mathcal{I}_c^0$  is not simply piecewise constant: it is constant on a large portion  $\mathcal{D}_c^{int}$  of its support  $\mathcal{D}_c$ , but on the boundary regions  $\mathcal{D}_c^{bd} = \mathcal{D}_c \setminus \mathcal{D}_c^{int}$ , consisting of narrow strips, it makes a smooth transition to 0, with a fairly uniform profile along every cradle edge segment. Step 13 carries out the second stage (MCA), and steps 14 - 16 cover the third stage of cradle wood grain texture separation and approximation.

### 3. IMPLEMENTATION AND EXPERIMENTS

We implement the first stage of our algorithm in MATLAB 2013b, using the Matlab function `radon` for the Radon trans-



**Fig. 3: Comparing cradle removal results** for the patch (and detail) shown in Fig. 1: *Left pair*: This paper’s algorithm; *Right pair*: Manual processing by Noelle Ocon, using Adobe Photoshop™.

form, allowing an angle deviation of  $5^\circ$  from the horizontal or vertical direction. For the shape functions  $\Phi_a$  in the smooth boundary approximation, we used scaled Gauss error functions  $a_1 \text{Erf}(a_2 \cdot)$ . For the MCA part, we use the MATLAB toolbox MCA Lab 110 from [7]. For the transforms involved in MCA and the feature extraction of the texture image (curvelet, dual-tree complex wavelet and shearlet transforms), we use CurveLab-2.1.3 from [8], the Dual-Tree Complex Wavelet Transform (Pack-version 2.1) introduced in [9] and the Fast Finite Shearlet Transform (FFST) toolbox from [10]. To obtain pixel-wise feature vectors via the curvelet transform, coefficients of low frequency curvelets are upsampled by interpolation<sup>4</sup> to the same resolution as the high frequency ones. To implement the Gibbs sampler of the additive factor model, we modified the code of [6], adopting the same values for the hyper-parameters of the priors. The attenuation factor  $\rho$  is set to 0.8 according to visual judgment, based on results with different  $\rho$  values. The Gibbs sampler is run for 1800 iterations with a burn-in of 1000 cycles; these numbers are chosen such that a stationary state is reached after burn-in and the result is visually the same if more iterations are executed.

We applied our method to the X-ray image of one panel in the *Ghissi Altarpiece* by Franciscussio Ghissi, a 14<sup>th</sup> century Italian painter, from the North Carolina Museum of Art (NCMA) collection. The image size is 7288-by-4376 pixels; the first stage algorithm is run on sub-images containing each one specific cradled region. After subtracting the constant (or smoothly modeled) part of the cradled image, the image is divided into overlapped sub-patches of size 512-by-512, and the second and the third stage algorithm is run on each sub-patch. The input cradled and panel-only feature vectors of the additive factor model are randomly sampled in the feature vectors belonging to the target sub-patch or to its immediate neighbors, so that the models learned are locally consistent. Those feature vectors not used as samples of the model are decomposed into “panel-only” and “cradled” vectors by posterior inference. The panel-only components of these sub-patches

<sup>4</sup>This is not necessary if FFST is used instead, where the low frequency coefficients are not downsampled in the forward transform.

are reconstructed in image space and merged to obtain the final result.

Fig. 3 shows the result of our algorithm on the patch illustrated in Figures 1 and 2; it compares favorably with the result obtained “manually” by one of us (N.O.), the conservator at NCMA in charge of removing X-ray image artifacts of this type, up to the standards required for art history publication. High-resolution versions for a much larger segment are at URL <https://www.math.duke.edu/~rachel/research/CradleRemoval/CradleRemoval.html>

#### 4. CONCLUSION

In this paper we propose a systematic digital cradle removal method for X-rays of paintings on wooden panel. The constant intensity, due to the thickness of the cradling, and the wood grain texture of the cradle are estimated and subtracted. Our method is the first to incorporate a probability model for the extraction of the wood grain coming only from the cradle without removing the wood grain of the panel itself. Our results compare favorably with the best existing practice.

This digital method has the potential to make X-ray images better-read and reveal hidden information not easily seen. We continue to investigate possible improvements and extensions to our current algorithm. Comparison with existing (manual) practice is only a substitute for ground truth. For some paintings, cradling-induced tensions have caused structural damage; museums are starting to physically remove these cradles. Applying our algorithm to the X-ray of such a cradled panel and comparing with an X-ray after “de-cradling” will make comparison possible with a “real” ground truth, when such datasets become available. It will also help us verify the assumption that the texture of the wooden panel is attenuated in areas that are superposed with the cradle.

#### 5. REFERENCES

- [1] C Richard Johnson, Ella Hendriks, Igor J Berezhnoy, Eugene Brevdo, Shannon M Hughes, Ingrid

- Daubechies, Jia Li, Eric Postma, and James Z Wang, "Image processing for artist identification," *Signal Processing Magazine, IEEE*, vol. 25, no. 4, pp. 37–48, 2008.
- [2] Bruno Cornelis, Yun Yang, Joshua T Vogelstein, Ann Dooms, Ingrid Daubechies, and David Dunson, "Bayesian crack detection in ultra high resolution multimodal images of paintings," *arXiv preprint arXiv:1304.5894*, 2013.
- [3] Bruno Cornelis, Ann Dooms, Jan Cornelis, and Peter Schelkens, "Digital canvas removal in paintings," *Signal Processing*, vol. 92, no. 4, pp. 1166–1171, 2012.
- [4] Michael Elad, J-L Starck, Philippe Querre, and David L Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis (mca)," *Applied and Computational Harmonic Analysis*, vol. 19, no. 3, pp. 340–358, 2005.
- [5] Bruno Cornelis and Alex Goodfriend, "Removal of canvas patterns in digital acquisitions of paintings: A source separation problem," .
- [6] Anirban Bhattacharya and David B Dunson, "Sparse bayesian infinite factor models," *Biometrika*, vol. 98, no. 2, pp. 291–306, 2011.
- [7] "MCALAB," <https://fadili.users.greyc.fr/demos/WaveRestore/downloads/mcalab/Home.html>.
- [8] "curvelet.org," <http://www.curvelet.org/>.
- [9] Nick G Kingsbury, "The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters," in *Proc. 8th IEEE DSP Workshop*. Citeseer, 1998, vol. 8, p. 86.
- [10] Sören Häuser, "Fast finite shearlet transform," *arXiv preprint arXiv:1202.1773*, 2012.