A TWO STAGE PATCH-BASED MARKOV RANDOM FIELD APPROACH TO STRUCTURE-AWARE IMAGE INPAINTING

Attila Fésűs, Tijana Ružić and Aleksandra Pižurica

Ghent University, TELIN-IPI-iMinds, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium

ABSTRACT

We introduce a two stage patch-based structure aware inpainting method within a Bayesian framework. An original conditional Markov random field (MRF) model is introduced to encode spatial clustering of image patches with missing parts. Within this model an anisotropic MRF encodes spreading of given classes of image content in locally dominant directions. We infer posterior probabilities for each patch to belong to different classes, and use these to select plausible candidate replacement patches. The second stage solves an optimization problem with multiple-candidates for each damaged patch, as in recent global inpainting methods. The results demonstrate a clear advantage of our structure-aware candidate selection.

I. INTRODUCTION

There are two main approaches to image inpainting: diffusion-based [1], and patch-based [2]–[6]. Diffusion-based methods propagate (diffuse) local structures in a given direction, in order to continue straight lines, curves or to fill in narrow holes. Patch-based methods are superior in inpainting larger missing regions, as they are capable of replicating both structure and texture of the surrounding areas. A recent overview of patch-based methods can be found in [6]. In a nutshell, these methods operate as follows. For each patch of the missing region (target patch), a suitable replacement patch is found in the available part of the image (source region) and copied to the corresponding location. Propagating image structures (i.e., continuing lines, object boundaries etc.) is typically treated by defining the filling order [2], [6]. Another important aspect is optimizing the search for plausible candidates (potential replacement patches) in order to accelerate the inpainting process and to prevent introduction of wrong textures into the target region. Solutions include using image segmentation [3], [7] and contextual descriptors [8], [9] to guide the search for candidate patches.

The approaches for optimizing filling order and candidate selection can be combined with both greedy [2], [5] and global [4], [9] methods. While greedy methods search for only one, best match for each target patch, global methods allow multiple candidates for each target patch and subsequently solve an optimization problem. This optimization problem can be understood as solving a puzzle with multiple pieces at each position, such that the selected patches agree with the underlying undamaged regions, and among themselves.

While the state-of-the-art in digital image inpainting is very advanced, such that even large missing regions can be filled in a visually plausible way [6], [9], certain limitations were recently reported in the application of virtual painting restoration such as crack inpainting [10]. In particular, in cases where painted structures, like fine inscriptions are relatively small compared to the crack width, it is very difficult for the inpainting algorithms to infer the correct structure locally. The solution proposed in [10] improves the candidate selection process by detecting locally likely directions of structure propagation and prioritizing accordingly certain types of source patches. In particular, directional neighborhoods $\mathcal{N}_i(p)$ are formed around each target patch ϕ_p using a finite number of discrete directions (e.g., horizontal, vertical, diagonal) and prior preference $P_{p,x}$ for selecting a source patch ϕ_x as a candidate replacement for ϕ_p was defined as $P_{p,x} = S(\phi_p, \phi_x) + \max_i \sum_{q \in \mathcal{N}_i(p)} S(\phi_x, \phi_q)$, where $S(\phi_p, \phi_x)$ is a certain measure of similarity between ϕ_p and ϕ_x. This idea was only briefly outlined in [10] and showed encouraging initial results.

In this paper, we cast a similar idea into a Bayesian framework and develop a novel algorithm for structure-aware candidate selection. We employ an anisotropic Markov random field (MRF) model to encode existence of locally dominant directional structures. In particular, we allow a given number (k) of (user-defined or inferred by clustering) classes $\rho \in \{1, 2, \ldots, k\}$ of image content and we estimate for each target patch ϕ_p marginal posterior probabilities that it belongs to these classes. The selection of plausible candidates ϕ_x is guided by these posterior probabilities $P(\rho_p \mid \phi_p)$. We employ this algorithm as the first stage of the inpainting process, and in the second stage we solve the resulting optimization problem (constructing a puzzle out of the selected candidates) using the MRF-based method of [8], an extended version of which was recently reported in [9].

This way we obtain a two-stage inpainting method, where both stages are within a unified Bayesian framework, using MRF priors and message-passing inference, albeit with different types of clique potentials.

The main novelties are in *(i)* introducing ideas for selecting candidate patches based on the posterior marginal probabilities of content classes of target patches; *(ii)* introducing an original conditional MRF model defined over image patches with missing parts and *(iii)* defining a concrete fitting criterion between the source and the target patches that employs the inferred posterior probabilities. Apart from this, as a by-product of our approach we also obtain a segmentation result, where each image patch, including those with partially or fully missing pixels, is assigned to one of the content classes. This result may in itself be of interest, even though we do not use it here directly. Experimental results on some parts of the digitized *Ghent Altarpiece* show a clear benefit from our two-stage method in virtual crack inpainting.

The paper is organized as follows. In Section 2, an outline of the complete proposed two-stage method is presented and the notation is set. Section 3 addresses the first stage and introduces a method for inferring posterior marginal probabilities of content classes of target patches. Section 4 establishes the fitting criterion based on the marginal posterior probabilities, and sets the way for the second stage of actual inpainting. Section 5 presents the results, and Section 6 concludes the paper.

II. A TWO-STAGE PATCH-BASED INPAINTING METHOD

In order to infer correctly the geometry of image structures in an inpainting process, we propose a novel two-stage patch-based impainting method. Here we give the basic outline of the algorithm. The role of each stage is as follows.

1) **Bird's view**
   Determining how image objects spread in the missing regions, following their geometry in a wider context. This yields plausible candidates for each target patch.

2) **Fine tuning**
   Solving a puzzle out of the selected candidates for replacement patches, such that they fit with each other and the known part of the image.

The problem defined under 2) has been well studied and can be solved by various inpainting algorithms. Our focus is on stage 1.

II-A. Notation

Let S denote an image grid. For a set of positions $Q \subset S$, $I(Q)$ denotes the set of the corresponding pixel values of the image I. Further on, let $\Omega \subset S$ be the missing part (target region), and $\Psi \subset S$ the known part of the image (source region), where $\Omega \cup \Psi = S$.

For a fixed integer ω we define a set of positions: $U = \{(x, y) \mid x = k\omega + 1,\ y = l\omega + 1,\ k, l \in \mathbb{N}_0,\ (x, y) \in S\}$. From now on, we will refer to pixel positions with p, instead of (x, y) for a simpler notation.

We define a $(2\omega + 1) \times (2\omega + 1)$ square mask ψ as a set of positions centered at the origin $(0, 0)$, and denote by $\psi_s := \psi + s$ a translated mask centered at $s \in S$.

Further on, $G = \{p \mid p \in U \ \wedge\ \psi_p \subset S \ \wedge\ \psi_p \cap \Omega \neq \emptyset\}$ is the set of positions from U which are centers of overlapping patches containing unknown pixels (the patches to be inpainted), and the set of central positions of all possible source patches is $\Lambda := \{s \in S \mid \psi_s \subset \Psi\}$.

We denote by ϕ_p the content of the patch occupying the positions in ψ_p.

III. A STRUCTURE-AWARE APPROACH: A BAYESIAN MODEL

In this section, we address the problem of inferring how probable it is that a given target patch (with partially or fully missing pixels) belongs to a given class of image content. This is made possible by taking into account the geometry of the classes by using wide neighborhoods of the patches.

We propose the neighborhood system shown on the left of Fig. 1. However, the proposed approach can be applied with an arbitrary neighborhood as well.

Let $G' \subset U$ be a set of positions consisting of the elements of G and their neighbors (denoted by $\tilde{G}$) according to the chosen neighborhood system. We refer to the elements of $\tilde{G}$ as *quasi-nodes*. See the image in the center of Fig.1 for an illustration. $\mathcal{N}(p)$ denotes the neighborhood of $p \in G'$.

The class of a patch is inferred based on the dominant class of its pixels and the pixels of the neighboring patches. The image content classes are modeled in a color or feature space $\mathcal{F}$.

Let $\rho_p \in \{1, 2, \ldots, k\}$ be a random variable denoting the class of the patch centered in $p \in G'$. The problem we are addressing can be formulated in the Bayesian framework. Denote by $P(\boldsymbol{\rho} \mid \boldsymbol{\phi})$ the joint a posteriori probability over all patches in our MRF model. With the assumption that the variables ϕ_p for $p \in G'$ are conditionally independent given their classes ρ_p, the likelihood is expressed as $P(\boldsymbol{\phi} \mid \boldsymbol{\rho}) = \prod_{p \in G'} P(\phi_p \mid \rho_p)$. Further on, we have $P(\phi_p \mid \rho_p) \propto e^{-D(\phi_p \mid \rho_p)}$, where $D(\phi_p \mid \rho_p)$ is the data cost.

The prior probability $P(\boldsymbol{\rho})$ is expressed by modeling $\boldsymbol{\rho}$ as a Markov random field. Let $\mathcal{G} = (G', E_{G'})$ be a graph, where G' is the set of nodes, and the set of edges $E_{G'} \subset G' \times G'$ is defined according to the chosen neighborhood system. Assuming that local Markov properties hold, the set of random variables $\{\rho_p\}_{p \in G'}$ forms a MRF with respect to the graph $\mathcal{G}$. The prior can be expressed as $P(\boldsymbol{\rho}) \propto e^{-E(\boldsymbol{\rho})}$, where $E(\boldsymbol{\rho})$ is the prior energy, which is a sum of clique potentials.

By Bayes' formula we have $P(\boldsymbol{\rho} \mid \boldsymbol{\phi}) \propto P(\boldsymbol{\phi} \mid \boldsymbol{\rho})P(\boldsymbol{\rho})$, and using $P(\boldsymbol{\rho} \mid \boldsymbol{\phi}) \propto e^{-E(\boldsymbol{\rho} \mid \boldsymbol{\phi})}$, the posterior energy can

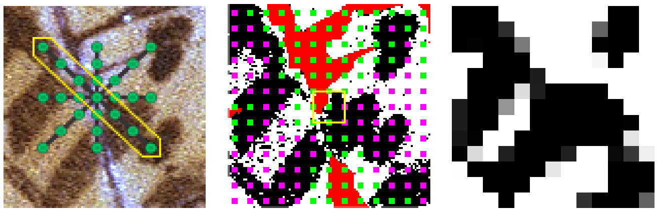

Fig. 1: **Left:** the proposed neighborhood: green dots represent centers of patches (nodes), different directions (horisontal, vertical, diagonal) are separate sub-neighborhoods. The sub-neighborhood marked is the one indicating the presence of a letter in the central patch. **Center:** Illustration of nodes: nodes subject to inpainting (green) and quasi-nodes (purple). The yellow frame indicates the size of each patch. Black and white indicate the class of the pixels, while cracks are marked by red; **Right:** Illustration of the result of stage 1: different shades represent different a posteriori probabilities of belonging to the classes.

be expressed as:

$$E(\boldsymbol{\rho} \mid \boldsymbol{\phi}) = \sum_{p \in G'} \gamma \cdot D(\phi_p \mid \rho_p) + \sum_{c \in \mathcal{C}} \delta \cdot V_c(\boldsymbol{\rho}), \qquad (1)$$

where $\gamma > 0$ and $\delta > 0$ express the parameters within the two terms, and $V_c(\boldsymbol{\rho})$ is the potential of the clique c.

Let us remark here that $\sum_{i=1}^{k} P(\rho_p = i \mid \phi_p) = 1$ holds for all $p \in G'$, as each node belongs to one of the classes.

We can rewrite (1) by defining a neighborhood potential $V_N(\rho_p, \boldsymbol{\rho}_{\mathcal{N}(p)})$ determined by the potentials of cliques belonging to the neighborhood of a given central node p. In this setting the second sum in the above expression can be replaced by summing $V_N(\rho_p, \boldsymbol{\rho}_{\mathcal{N}(p)})$ over $p \in G'$, as in [11] and [12].

We propose a natural choice of the parameters that balance the two terms of the posterior energy, based on the fraction of the known pixels α_p within a given image patch ϕ_p as follows:

$$E(\boldsymbol{\rho} \mid \boldsymbol{\phi}) = \sum_{p \in G'} \alpha_p \cdot D(\phi_p \mid \rho_p) + (1 - \alpha_p) \cdot V_N(\rho_p, \boldsymbol{\rho}_{\mathcal{N}(p)}),$$

For patches with fully unknown content ($\alpha_p = 0$), we rely solely on the neighboring content, while for fully known patches ($\alpha_p = 1$), we rely on the likelihood term.

For easier further discussion, in the above sum we denote the terms by $E_p := \alpha_p \cdot D(\phi_p \mid \rho_p) + (1 - \alpha_p) \cdot V_N(\rho_p, \boldsymbol{\rho}_{\mathcal{N}(p)})$.

III-A. Data cost

We model the local evidence by a multivariate Gaussian distribution. Let patch ϕ_p be represented by a vector $\hat{\boldsymbol{\phi}}_p$ of the mean value of its known pixels in the feature space $\mathcal{F}$. We have:

$$P(\phi_p \mid \rho_p = i) = \frac{1}{\sqrt{(2\pi)^l |\boldsymbol{\Sigma}_i^2|}} e^{-\frac{1}{2}(\hat{\boldsymbol{\phi}}_p - \boldsymbol{m}_i)^T \boldsymbol{\Sigma}_i^{-1} (\hat{\boldsymbol{\phi}}_p - \boldsymbol{m}_i)},$$

where l is the number of dimensions of the feature space, k denotes the number of classes as before, and $\boldsymbol{m}_i, \boldsymbol{\Sigma}_i$ for $i = 1, 2, \ldots, k$, are the means (vectors of size l) and covariance matrices of the data sets:

$$T_i = \{\hat{\boldsymbol{\phi}}_s \mid \psi_s \subset \Psi \ \wedge \ s \in U \ \wedge \ \hat{\rho}_s = i\},$$

where $\hat{\rho}_s$ denotes the class to which most of the pixels of patch ϕ_s belong.

Using the relationship between the energy and the posterior probability, the data cost function can be expressed as:

$$D(\phi_p \mid \rho_p = i) = c_{1,i} \cdot (\hat{\boldsymbol{\phi}}_p - \boldsymbol{m}_i)^T \boldsymbol{\Sigma}_i^{-1} (\hat{\boldsymbol{\phi}}_p - \boldsymbol{m}_i) + c_{2,i},$$

where $c_{1,i}, c_{2,i}$ are constants.

III-B. Neighbourhood potential: an anisotropic MRF approach

In order to infer the geometric structure of objects in the damaged regions, we turn to an anisotopic approach: the neighborhood potential (for each central node $p \in G$) will be determined by a sub-neighborhood that indicates presence of dominant local structures (edges,lines, etc.). The basic idea is that each sub-neighborhood represents a direction (or a pattern) in which the geometrical structure of interest can possibly spread. For example, the sub-neighborhood marked in Fig. 1 (left), is dominant, as it suggests the propagation of the letter in the central node.

The potential of each sub-neighborhood is the sum of clique potentials of its cliques. For simplicity, we will consider only pair-wise cliques in our model. This idea originates from [13], where it was introduced in a form of a binary anisotropic model (two classes) for pixels, and used for noise removal in images.

Let $\mathcal{N}_1(p), \ldots, \mathcal{N}_j(p) \subset \mathcal{N}(p)$ be sub-neighborhoods of the neighborhood $\mathcal{N}(p)$ of node p, such that all $\mathcal{N}_i(p)$ are of the same size and $\cup_{i=1}^{j} \mathcal{N}_i(p) = \mathcal{N}(p)$.

The pair-wise clique potential function is defined as:

$$V_2(\rho_p, \rho_q) := \begin{cases} -\beta & , \quad \rho_p = \rho_q \\ 0 & , \quad \rho_p \neq \rho_q \end{cases}, \qquad (2)$$

where β is a positive constant. We define sub-neighborhood potential functions V_{N_i} as

$$V_{N_i}(\rho_p, \rho_{\mathcal{N}_i(p)}) := \sum_{q \in \mathcal{N}_i(p)} V_2(\rho_p, \rho_q).$$

We distinguish two types of classes: *directed* and *non-directed*. An example of a *non-directed* class is background, or any class with structures that spread uniformly throughout different directions, while the *directed* classes are those with a structured geometry, relating to the defined sub-neighborhoods. Let $\mathcal{D} \subset \{1, 2, \ldots, k\}$ be the set of all *directed* classes. Finally, we define neighborhood potential as:

$$V_N(\rho_p, \rho_{\mathcal{N}(p)}) := \begin{cases} \min_i V_{N_i}(\rho_p, \rho_{\mathcal{N}_i(p)}) & ,\rho_p \ \in \mathcal{D} \\ \max_i V_{N_i}(\rho_p, \rho_{\mathcal{N}_i(p)}) & ,\neg\rho_p \ \in \mathcal{D} \end{cases}$$

so that if any sub-neighborhood indicates the presence of an object (structure) belonging to a *directed* class in the patch centered at p, that class is encouraged at node p. Otherwise, if there is no indication of any *directed* class, *non-directed* classes take over.

III-C. The optimization problem

Note that in the energy function $E(\rho \mid \phi)$ the term E_p for a quasi-node $p \in \tilde{G}$, reduces to $E_p = D(\phi_p \mid \rho_p)$, as $\alpha_p = 1$. Therefore, in the optimization process, we don't visit (update) quasi-nodes, in each iteration, as their contribution E_p to the energy remains unchanged. This leads to less computational complexity and faster executing time.

Further on, the normalization issues are solved naturally by using $P(\rho_p \mid \phi_p) \propto e^{-E_p}$ and $\sum_{i=1}^{k} P(\rho_p = i \mid \phi_p) = 1$. Hence, in practice we solve

$$\hat{\boldsymbol{\rho}} = \text{argmax}_{\boldsymbol{\rho}} P(\boldsymbol{\rho} \mid \boldsymbol{\phi}). \quad (3)$$

The resulting marginal posterior distributions are illustrated on the right of Fig. 1, in a case of 2 classes (one directed, and one non-directed).

IV. SELECTING CANDIDATE REPLACEMENT PATCHES

The marginal posterior probability $P(\rho_p|\ \phi_p)$ regarding node p is a vector of length k (number of classes), where the i-th coordinate is $P(\rho_p = i|\phi_p)$, for $i \in \{1, 2, \ldots, k\}$. Recall that $\sum_i P(\rho_p = i|\phi_p) = 1$. For the sake of finding suitable candidates for inpainting at node $p \in G$, a vector comparable to $P(\rho_p|\ \phi_p)$ needs to be defined for patches in the source region (potential replacement patches). For $x \in \Lambda$, this is the weighted (normalized) likelihood vector $L(\rho_x\, \phi_x)$ of length k, where the i-th coordinate is $L(\rho_x = i \mid \phi_x) \propto e^{-D(\phi_x = i \mid \rho_x)}$, and $\sum_i L(\rho_x = i \mid \phi_x) = 1$ holds.

$x \in \Lambda$ is considered a suitable candidate for inpainting at the node $p \in G$ if:

$$\|P(\rho_p \mid \phi_p) - L(\rho_x\, \phi_x)\|_2 < \epsilon, \quad (4)$$

for some small $\epsilon > 0$. Note that for a quasi-node $q \in \tilde{G}$ we have $P(\rho_p|\ \phi_p) = L(\rho_x\, \phi_x)$, as $\tilde{G} \subset \Lambda$ (patches represented by quasi-nodes are themselves potential replacement patches). This makes the defined vector $L(\rho_x\, \phi_x)$ and fitting criterion (4) fully consistent.

The search for candidates satisfying (4) can be carried out in different ways from exhaustive search to searching efficiently structured data using various approaches. For simplicity, we choose to divide the interval $[0, 1]$ into (say m) disjoint subintervals of equal lengths. Denote the set of these subintervals by $\mathcal{I}$. $\mathcal{I}^k$ is the Cartesian power of the set $\mathcal{I}$. We consider patch x with center in Λ a suitable candidate for inpainting at node $p \in G$ if:

$$(\exists J \in \mathcal{I}^k)(P(\rho_x \mid \phi_x) \in J \wedge P(\rho_p \mid \phi_p) \in J). \quad (5)$$

All the possible replacement patches (patches with centers in Λ) are classified into these subintervals. This can be done also per large enough non-overlapping blocks of the image, so the classification is more efficient. Such implementation fits very well if a context aware inpainting, based on selection of the appropriate blocks, is used in the second stage. If there is no such block selection mechanism, we simply take the candidates from the whole image.

When the appropriate candidates are selected as described, they are further pruned according to the inpainting method used in stage 2. As we are using the inpainting algorithm from [8] in the second stage, the final selection of candidate replacement patches is based on SSD (sum of squared differences) between the known part of the target patch and the same part of the candidate patch.

V. IMPLEMENTATION AND RESULTS

We solve the optimization problem (3) by Neighborhood Consensus Message Passing (NCMP) [14], which is an efficient, low-complexity alternative to standard loopy belief propagation (LBP). Convergence is achieved in less than 20 iterations in each experiment.

The second stage, where the actual inpainting takes place, formulated as an optimization problem within a MFR model as well, is also implemented by NCMP as in [8]. In each experiment we used 100×100 blocks.

NCMP is not only simpler and much faster than belief propagation, but it was also proven to give good results in applications that employ models related to ours. For example, in [14] it was shown that NCMP can yield in much shorter time similar quality of the results as LBP in applications like segmentation of a noisy image and patch based super-resolution with MRF priors. However, our framework is not tuned to a particular inference method and allows using LBP as well. We do not expect that this would affect significantly the quality of the results, only the number of required iterations to achieve convergence and the complexity per iteration is likely to change.

The number of final candidate replacement patches per node is chosen to be k^4, where k is the number of classes. In the bottom image of Fig. 2 we have 3 classes, while 2 in the other experiments (top of Fig. 2, and Fig. 3). Patch size $(2\omega + 1) \times (2\omega + 1)$, varies according to the size of images, and the size of unknown regions and structures to be followed. In our experiments $\omega = 8$ for the top image in Fig. 2, $\omega = 6$ for the bottom image in the same figure, and $\omega = 10$ for Fig. 3.

The algorithm is tested on different parts of the *Ghent Altarpiece*, to inpaint cracks and regions of removed overpaint. We compare the results with [8] and with *Content*

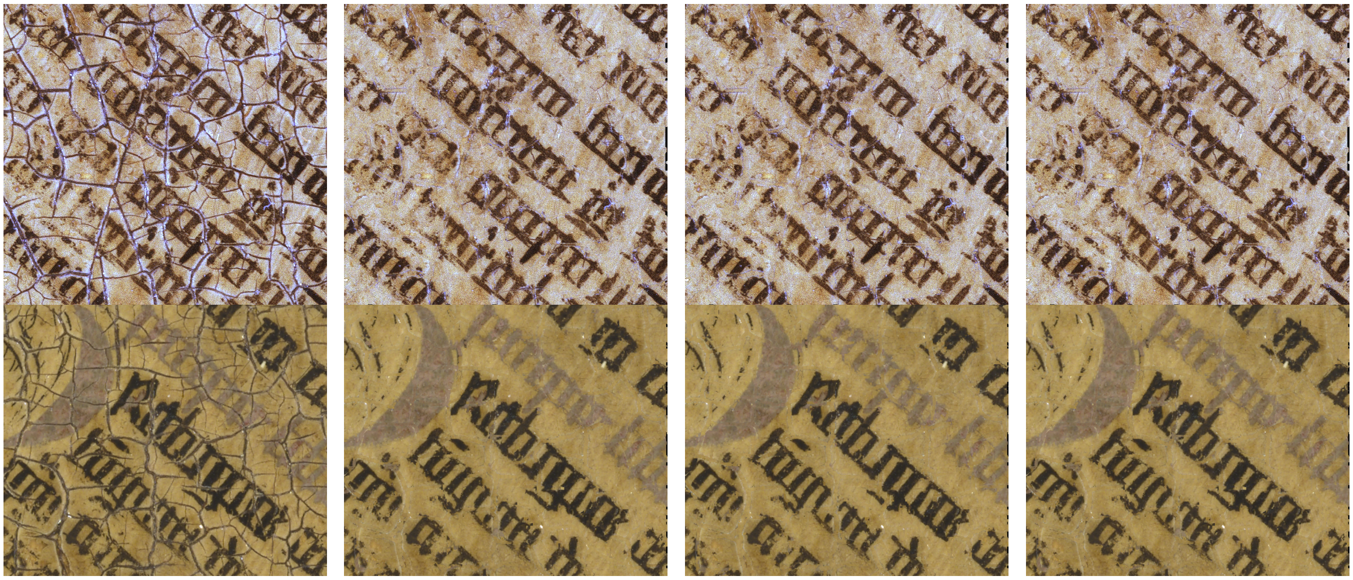

Fig. 2: In the two subfigures different parts of the book from the panel *Annunciation to Marry* from two different collections of scans are shown. **Left to right:** the original scan (up: 758×707; down: 495×415), crack inpainting result using *Adobe Photoshop's Content Aware Fill*, crack inpainting result using method from [8], and crack inpainting results of the proposed method.

Fig. 3: A small part of the *John the Evangelist*, panel where overpaint is removed. **Left to right:** the original scan (827×550), the result using *Adobe Photoshop's Content Aware Fill*, the result of [8], the result of the proposed method.

Aware Filling available in the commercial software *Adobe Photoshop CS5* based on [15], [16]. The results of *Content Aware Filling* show more artefacts in comparison to [8], but in the results of [8] some parts of letters are deleted as well. The improvements achieved by our method are clearly visible in Fig. 2 - Fig. 4. The structure of the letters is better inferred, as hence the letters are better inpainted (see enlarged parts in Fig. 4).

As a feature space we used a concatenation of RGB, Lab and HSV color spaces (9 channels in total). A more general and perhaps more effective solution could be found, however this is not a central question in the proposed method, but rather an implementation detail.

For initialization, we used pixel-wise segmentation using standard methods such as *k-means*. The parameter β was optimized experimentally (stable results are achieved with values between 1 and 1.2).

A major portion of the computation time in the complete inpainting process is taken by the second stage. A comprehensive analysis of the time and space complexity in this stage (also in comparison to related inpainting methods) was already reported in [9], showing, e.g., significant savings in comparison to the MRF-based method of [4]. The analysis in [9] also shows that the computation time of the second stage is mostly influenced by the number of the required SSD computations. The overall number of calculations depends on the number of potential candidate replacement patches before pruning, and on the number of final candidates per node. As our method limits the number of potential candidates by allowing only those consistent with the inference from the first stage of the algorithm, the extent of the SSD comparisons in the pruning process can be reduced greatly. This reduces the overall complexity of the algorithm, as the first stage is much less computationally expensive than the second one. We leave a more precise quantification of the resulting computation speed-up for a follow-up paper, because certain parts of our algorithms are not yet optimized in terms of speed. In particular, (5) is a computationally

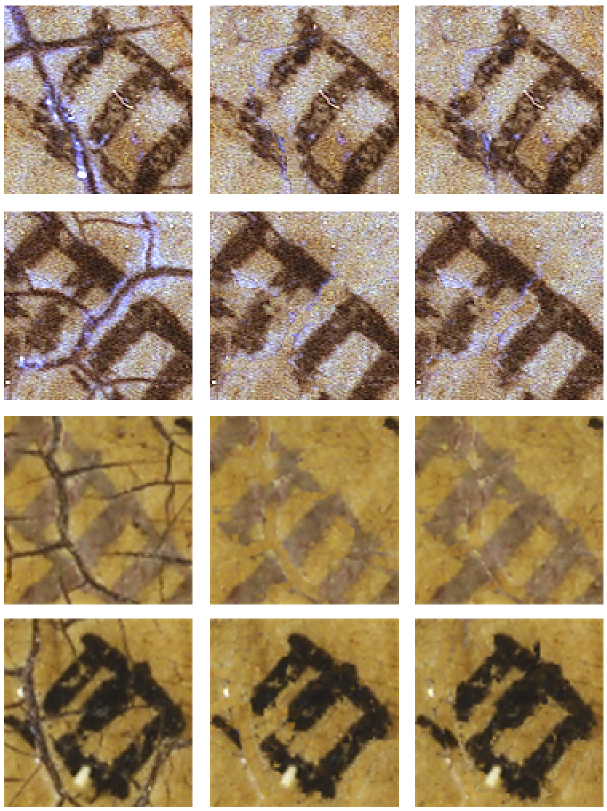

Fig. 4: Some enlarged parts of Fig. 2: **rows:** different enlarged parts (1-2 upper, 3-4 lower image); **columns (left to right):** original image, inpainting results using the method from [8], and the proposed method.

expensive approach to find candidate replacement patches satisfying (4), and depends greatly on the number of classes. Optimizing the implementation of this criterion will be a part of a future work.

VI. CONCLUSION

In this paper we introduced a structure-aware candidate selection mechanism based on the posterior marginal probabilities of content classes of target patches, which can be employed to improve the results of existing patch based inpainting methods. The potential of such an approach is evident from the presented experiments. The inference can be carried out based on different feature spaces, wherever the desired classes are distinguishable. Furthermore, the proposed selection criterion of replacement patch candidates can be further constrained by combining it with other existing approaches, what gives the proposed method lots of possible applications and ways for future developments.

VII. REFERENCES

[1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," *SIGGRAPH*, 2000.

[2] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1200–1212, Sept. 2004.

[3] J. Jia and C.-K. Tang, "Image repairing: robotus image synthesis by adaptive and tensor voting," *CVPR*, pp. 643–650, 2003.

[4] N. Komodakis and G. Tziritas, "Image completion using efficient belief propagation via priority sheduling and dinamic pruning," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2649–2661, Nov. 2007.

[5] O. L. Meur, J. Gautier, and C. Guillemont, "Examplar-based inpainting based on local geometry," *in ICIP*, pp. 3462–3465, 2011.

[6] C. Guillemot and O. L. Meur, "Image inpainting: Overview and recent advances," *IEEE Signal Processing Magazine*, vol. 31 (1), pp. 127–144, Jan. 2014.

[7] J. Lee, D.-K. Lee, and R.-H. Park, "Robotus examplar-based inpainting algorithm using region segmentation," *IEEE Transactions on Consumer Electronics*, vol. 58, no. 2, pp. 553–561, May 2012.

[8] T. Ružić, A. Pižurica, and W. Philips, "Markov Random Field based image inpainting with context-aware label selection," *ICIP*, pp. 1733–1736, 2012.

[9] T. Ružić and A. Pižurica, "Context-aware patch-based image inpainting using Markov random field modelling," *IEEE Transactions on Image Processing*, vol. 24, no. 1, pp. 444–456, Jan. 2015.

[10] A. Pižurica, L. Platiša, T. Ružić, B. Cornelis, A. Dooms, M. Martens, H. Dubois, B. Devolder, M. D. Mey, and I. Daubechies, "Digital Image Processing of the Ghent Altarpiece," *IEEE Signal Processing Magazine*, 2015 (in press).

[11] S. Z. Li, *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.

[12] G. Winkler, *Image Analysis, Random Fields and Dynamic Monte Carlo Methods*. Springer-Verlag, 1995.

[13] A. Pižurica, W. Philips, I. Lemahieu, and M. Acheroy, "A joint inter- and intrascale statistical model for bayesian wavelet based image denoising," *IEEE Transactions on Image Processing*, vol. 11, no. 5, pp. 545–557, May 2002.

[14] T. Ružić, A. Pižurica, and W. Philips, "Neighborhood-consensus message passing as a framework for generalized iterated conditional expectations," *Pattern Recognition Letters*, vol. 33, no. 3, pp. 309–3018, Feb. 2012.

[15] Y. Wexler, E. Shechtman, and M. Irani, "Space-time completion of video," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 3, pp. 463–476, Mar. 2007.

[16] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "Patch-match: a randrandomized correspondance algorithm for structural image editing," *SIGGRAPH*, 2009.