



Published in Image Processing On Line on YYYY-MM-DD.
 Submitted on YYYY-MM-DD, accepted on YYYY-MM-DD.
 ISSN 2105-1232 © YYYY IPOL & the authors CC-BY-NC-SA
 This article is available online with supplementary materials,
 software, datasets and online demo at
<http://dx.doi.org/10.5201/ipol>

Non-local patch-based image inpainting

Alasdair Newson¹, Andrés Almansa², Yann Gousseau², Patrick Pérez³

¹ Université Paris Descartes (alasdair.newson@parisdescartes.fr)

² Telecom ParisTech, CNRS LTCI ({almansa,gousseau}@telecom-paristech.fr)

³ Technicolor (patrick.perez@technicolor.com)

PREPRINT November 11, 2015

Abstract

Image inpainting is the process of filling in missing regions in an image in a plausible way. In this contribution, we propose and describe an implementation of a patch-based image inpainting algorithm. The method is actually a two-dimensional version of our video inpainting algorithm proposed in [16]. The algorithm attempts to minimise a highly non-convex functional, first introduced by Wexler et al. in [18]. The functional specifies that a good solution to the inpainting problem should be an image where each patch is very similar to its nearest neighbour in the unoccluded area. Iterations are performed in a multi-scale framework which yields globally coherent results. In this manner two of the major goals of image inpainting, the correct reconstruction of textures and structures, are addressed. We address a series of important practical issues which arise when using such an approach. In particular, we reduce execution times by using the PatchMatch [3] algorithm for nearest neighbour searches, and we propose a modified patch distance which improves the comparison of textured patches. We address the crucial issue of initialisation and the choice of the number of pyramid levels, two points which are rarely discussed in such approaches. We provide several examples which illustrate the advantages of our algorithm, and compare our results with those of state-of-the-art methods.

1 Introduction

Image inpainting is the task of filling in an unknown region in an image. This is useful, for example, if an area is damaged, or if one wishes to remove an unwanted object from the image. Inpainting can be used for personal purposes or for professional image restoration. Two of the main goals of image inpainting are the convincing restitution of structures and textures which are coherent with the unoccluded part of the image.

The first approaches to image inpainting took the form of the minimisation of a functional [2, 4, 14], whose argument is the image to inpaint. The main goal of such approaches was to continue visual structures (such as edges) from one point on the occlusion boundary to another. While this produces good results for smooth images, they fail at reproducing textured content. Several algorithms which could maintain both structure and texture in an inpainting result were subsequently proposed [5, 6, 7].

These algorithms all hinged on the notion of image *patches* (small square neighbourhoods), which were first introduced by Efros and Leung [8] for texture synthesis.

The next step forward in inpainting was to reformulate these patch-based approaches (which were, for the most part, greedy methods) in a more rigorous minimisation framework. This was first done using discrete minimisation [11, 17], and then in a continuous context by Arias *et al.* [1]. One important element of such methods is that they use a multi-scale approach to converge to better inpainting solutions. To the authors’ knowledge, this class of algorithms (mutliscale minimisation of a patch-based energy) produces the best inpainting results to date.

In this paper, we take a similar approach. We minimise a patch-based functional taken from the video inpainting [18, 19] and texture synthesis [12] literature. We use the PatchMatch algorithm [3] to provide a crucial speedup for a nearest neighbour search. We identify an important problem concerning the correct inpainting of textures with patch-based methods, and modify the patch distance to address this problem, in a similar manner to that of Liu and Caselles [13] in image inpainting and Newson *et al.* [16] in video inpainting. We provide specific details concerning the influential initialisation step as well as the implementation of the the multi-scale scheme. These issues are rarely discussed in such approaches, and have important consequences on the quality of the inpainting results. A pseudo-code for our algorithm may be seen in Algorithm 2. To summarise, this work describes an image inpainting algorithm similar in its approach to our previous work on video inpainting [16]. In order to adhere to the execution time standard of IPOL (30s or less), we have restricted the scope of this work to the case of images. The algorithm produces high-quality results in the image case, and is shown to be as useful as in the cases of videos. The algorithm may be tested with the accompanying online demo as well as with the provided C++ code.

This publication accompanies our previous publication on video inpainting [16], which applies the same algorithm as presented here to videos. The corresponding video results and code (written in Matlab) can be seen at the following website:

http://www.telecom-paristech.fr/~gousseau/video_inpainting.

2 Proposed algorithm

As discussed above, the core of our algorithm is the minimisation of a patch-based non-local functional. Before describing the algorithm in detail, first of all we need to set out some notation.

2.1 Notation

The image support is denoted with Ω , and this support is separated into two regions: \mathcal{H} the occlusion (\mathcal{H} for “hole”) and \mathcal{D} the unoccluded region (\mathcal{D} for “data set”), such that $\mathcal{H} \cap \mathcal{D} = \emptyset$ and $\mathcal{H} \cup \mathcal{D} = \Omega$. A *position* in the support may be denoted with $p = (x, y)$. The image content itself is denoted with $u : \Omega \rightarrow \mathbb{R}^3$.

Another very important notion is that of a *patch*. We first define the patch neighborhood of a pixel p as a set of positions in Ω , denoted with \mathcal{N}_p , of cardinality N . The patch itself is denoted with $W_p = (u(x_1), \dots, u(x_N))$. A connected notion which we need to define is that of a *nearest neighbour* (NN for short) of a patch W_p . The NN of W_p is a patch W_q in another set of patches which minimises a certain patch distance $d(W_p, W_q)$. We shall choose NNs from the set of patches which contain no occluded pixels. For this, define $\tilde{\mathcal{D}}$ the set of positions such that $\forall p \in \tilde{\mathcal{D}}, \mathcal{N}_p \subset \mathcal{D}$. Therefore, the NN of W_p is W_q , with $q = \underset{k \in \tilde{\mathcal{D}}}{\operatorname{argmin}} d^2(W_p, W_k)$.

Finally, we define the *shift map* $\phi : \Omega \rightarrow \mathbb{N}^2$ as a vector field which shows where the NN of a patch is located. That is, the NN of W_p is $W_{p+\phi(p)}$.

2.2 Variational framework

The functional which we wish to minimise takes the following form:

$$E(u, \phi) = \sum_{p \in \mathcal{H}} d^2(W_p, W_{p+\phi(p)}). \quad (1)$$

Intuitively speaking, this highly non-convex functional specifies that for a “good” solution, each patch inside \mathcal{H} should be as similar as possible (in terms of the patch distance $d^2(\cdot, \cdot)$) to its nearest neighbour in the unoccluded region \mathcal{D} . We use the following distance:

$$d^2(W_p, W_{p+\phi(p)}) = \sum_{q \in \mathcal{N}_p} (u(q) - u(q + \phi(p)))^2. \quad (2)$$

Following the heuristic proposed in [19], a solution is obtained using an iterated alternating minimisation, as well as a multi-scale framework. At each scale, we minimise firstly with respect to the shift map ϕ , then with respect to the image content u . These two steps correspond, respectively, to the following operations:

- nearest neighbour search;
- image reconstruction.

The nearest neighbour search can be quite expensive, especially if the whole image is used as the search space. In Section 3.1, we specify the manner in which the NN search is carried out, and how the computational issue is addressed. The reconstruction process is described in Section 3.2. An improvement of this basic process is proposed in Section 3.3 in order to improve the reconstruction of textures. Section 3.4 deals with the initialisation step, and Section 3.5 deals with the multi-scale framework of the algorithm, both points being especially important given the non-convex nature of the problem.

3 Algorithm

Now that we have given a broad outline of our approach, we provide the specific algorithmic details necessary for the implementation of each component of our algorithm.

3.1 Approximate nearest neighbour search

As mentioned previously, the NN search can be quite computationally expensive. In such situations, it is common to use *approximate nearest neighbours* instead of exact nearest neighbours. In this implementation, we use the *PatchMatch* algorithm to search for nearest neighbours. This algorithm was introduced by Barnes *et al.* [3] and relies on the idea that good shift maps tend to be piecewise constant. In other words, a relative shift which leads to a good NN for a patch W_p has a good chance of leading to a good NN for the patches situated around W_p .

The algorithm consists of three steps:

- random initialisation;
- propagation;
- random search.

The first step is carried out once at the beginning of the algorithm. Each occluded pixel is randomly assigned a candidate NN in the region $\tilde{\mathcal{D}}$. The next two steps are carried out on each patch, and each patch is visited in lexicographical order. This is carried out a predefined number of times. The propagation step attempts to spread good shifts throughout the vector field ϕ . Finally, the random search step looks for better NNs randomly in an increasingly small window around the current NN. Since PatchMatch itself is not a novelty of our algorithm, we refer the reader to the original publication [3] for further details.

In practice, we only need the initialisation step during the filling in of the occlusion (at the beginning of our algorithm). After this, we may consider that we have a good initialisation of ϕ , which is simply the ϕ of the previous iteration. Let us emphasize that PatchMatch has already been used as a way to accelerate the search for ANN in a commercial implementation of the algorithm from [19], in the *content aware* tool from the software Photoshop. The pseudo-code for the PatchMatch algorithm may be seen in Algorithm 1.

Algorithm 1: ANN search with PatchMatch.

Data: Current inpainting configuration u (height: m , width: n), ϕ , $\tilde{\mathcal{H}}$
Parameters: r_{max} ($\max(m, n)$), ρ (0.5)
Result: ANN shift map ϕ

```

for  $k = 1$  to  $k_{max}$  do
  for  $p = p_1$  to  $p_{|\tilde{\mathcal{H}}|}$  (pixels in  $\tilde{\mathcal{H}}$  lexicographically ordered) do
    if  $k$  is even then                                     /* Propagation on even iteration */
       $a = p - (1, 0)$ ,  $b = p - (0, 1)$ ;
       $q = \arg \min_{r \in \{p, a, b\}} d(W_p^u, W_{p+\phi(r)}^u)$ ;
      if  $p + \phi(q) \in \tilde{\mathcal{D}}$  then  $\phi(p) \leftarrow \phi(q)$ ;
    else                                                 /* Propagation on odd iteration */
       $a = p + (1, 0)$ ,  $b = p + (0, 1)$ ;
       $q = \arg \min_{r \in \{p, a, b\}} d(W_p^u, W_{p+\phi(r)}^u)$ ;
      if  $p + \phi(q) \in \tilde{\mathcal{D}}$  then  $\phi(p) \leftarrow \phi(q)$ ;
    end
     $z_{max} \leftarrow \lceil -\frac{\log(\max(m, n))}{\log(\rho)} \rceil$ ;
    for  $z = 1$  to  $z_{max}$  do
       $q = p + \phi(p) + \lfloor r_{max} \rho^z \text{RandUniform}([-1, 1]^2) \rfloor$ ;
      if  $d(W_p^u, W_{p+\phi(q)}^u) < d(W_p^u, W_{p+\phi(p)}^u)$  and  $p + \phi(q) \in \tilde{\mathcal{D}}$  then  $\phi(p) \leftarrow \phi(q)$ ;
    end
  end
end
end

```

3.2 Reconstructing the image

The second step in the iterative algorithm is the minimisation of the functional in Equation (1) with respect to u . Intuitively speaking, this is the process of assigning a colour value to each position in the occluded area, or as we refer to it here, the *reconstruction* process. The minimisation of (1) with respect to u should lead to each pixel being reconstructed with an *unweighted mean* of several colour values in \mathcal{D} . However, we shall use a *weighted mean* scheme, initially proposed by Wexler *et al.* [19]. Experimentally, we have observed that this improves the convergence of the algorithm.

Given a shift map, each pixel is reconstructed as:

$$u(p) = \frac{\sum_{q \in \mathcal{N}_p} s_q^p u(p + \phi(q))}{\sum_{q \in \mathcal{N}_p} s_q^p}, \quad (3)$$

where s_q^p is the weight assigned to the pixel value for p indicated by the ANN of W_q ,

$$s_q^p = \exp\left(\frac{-d^2(W_q, W_{q+\phi(q)})}{2\sigma^2}\right). \quad (4)$$

Informally, it can be seen that each pixel p is reconstructed using the “versions” of p indicated by the ANNs of all the patches which contain p . This operation is carried out for all pixels $p \in \mathcal{D}$.

An inevitable problem of such a reconstruction scheme is blurred results, especially in textured areas. This was noticed by Wexler *et al.* in [19], who proposed a method to avoid this effect which used the mean shift algorithm. The idea is to gradually decide on the best value for each pixel by clustering the colour values and only using pixels which correspond to the dominant cluster for reconstruction. After experimentation, we found that such a sophisticated approach was unnecessarily complex, and a very similar result could be obtained simply by choosing the “best” pixel at the *end* of the algorithm (at the finest pyramid level) when it has converged. Therefore the final reconstruction step is the following:

$$u(p) = u(p + \phi(q^*)), \text{ with } q^* = \underset{q \in \mathcal{N}_p}{\operatorname{argmin}} d^2(W_q, W_{q+\phi(q)}). \quad (5)$$

3.3 Inpainting with textures

Even though patch-based methods were initially introduced for texture synthesis, patch-based reconstruction of textured areas is often problematic when dealing with composite scenes made of different regions.. More precisely, when several regions with a similar average colour but differing texture contents exist, three main problems can arise with multi-scale and iterative algorithms:

- ambiguities in patch comparisons at coarse pyramid levels;
- failure of the regular ℓ_2 patch distance to correctly compare textures;
- the weighted mean reconstruction tends to smooth textures during the algorithm, which in turn leads to smooth patches being chosen as NNs.

In order to address these problems, we introduce *texture features* into the patch distance. The goal of these features is to correctly match textured patches with similarly textured patches. After experimenting with several different features, we found that the following features, inspired by the work of Liu and Caselles [13], identified textures in a satisfactory manner for our purposes:

$$T_x(p) = \frac{1}{\operatorname{card}(\nu)} \sum_{q \in \nu} |I_x(p)| \quad (6)$$

$$T_y(p) = \frac{1}{\operatorname{card}(\nu)} \sum_{q \in \nu} |I_y(p)|,$$

where ν is a neighbourhood centred on p , and I_x and I_y are the derivatives of the image in the x and y directions. In the case of white noise patches with a Gaussian distribution, these attributes

tend to the variance of the noise (up to a multiplicative factor) as the size of ν increases. Therefore, we consider our texture to be well-represented by the variance of its pixel values. This is obviously quite a simple representation of a texture, but we found that it distinguishes well between textured and non textured areas with the same average colour.

Thus, the squared patch distance is now redefined as:

$$d^2(W_p, W_q) = \frac{1}{N} \sum_{r \in \mathcal{N}_p} \left(\|u(r) - u(r - p + q)\|_2^2 + \lambda \|T(r) - T(r - p + q)\|_2^2 \right). \quad (7)$$

It is important to note that we calculate these features at *the finest pyramid level* for the unoccluded pixels $p \in \tilde{\mathcal{D}}$, and propagate them to coarser levels by nearest neighbour subsampling. This is necessary since the textures are not present at coarser levels, the image having been smoothed and subsampled.

This solves the problem of correctly matching textured patches, but due to the patch averaging process in Equation 4, we still have a tendency to smooth the result. This gradually leads to inpainting with the wrong texture, in spite of the process detailed in Section 3.2. To deal with this, we create a separate image with the texture features, and this image is inpainted in parallel to the colour image.

$$T(p) = \frac{\sum_{q \in \mathcal{N}_p} s_p^q T(p + \phi(q))}{\sum_{q \in \mathcal{N}_p} s_p^q}, \quad \forall p \in \mathcal{H}. \quad (8)$$

A comparative example of the effect of using these features may be seen in Figure 2. It is clear that without these features, satisfactory results are difficult to obtain.

3.4 Initialisation

All iterative methods which attempt to optimise an objective function require an initial solution, and this initialisation turns out to be very important in our case (which is highly non-convex, as mentioned previously). It is also an issue which is seldom discussed in many works on the subject.

We propose an “onion-peel” approach which inpaints the occlusion one layer at a time, eroding the occlusion progressively. The layers are one-pixel thick, and located at the border of the occlusion. Each pixel in a layer is processed using the information of the inpainting solution of the previous layer, and thus there is no processing order for pixels within a single layer (they are processed “in parallel”). For such an approach, we obviously need to use a *partial* patch comparison, since the content of the patches will not be completely known.

Let $\mathcal{H}' \subset \mathcal{H}$ be the current occlusion, and $\partial\mathcal{H}' \subset \mathcal{H}'$ the current layer to inpaint. We now define the partial patch neighbourhood :

$$\mathcal{N}'_p = \{q \in \mathcal{N}_p, q \notin \mathcal{H}'\}. \quad (9)$$

The distance between two partially known patches W_p and $W_{p+\phi(p)}$ is defined as:

$$d^2(W_p, W_{p+\phi(p)}) = \frac{1}{|\mathcal{N}'_p|} \sum_{q \in \mathcal{N}'_p} \left(\|u(q) - u(q + \phi(p))\|_2^2 + \lambda \|T(q) - T(q + \phi(p))\|_2^2 \right). \quad (10)$$

Now that we are able to compare two partially known patches, we wish to reconstruct the information of a given pixel p in the current occlusion layer $\partial\mathcal{H}'$. For this, we need to choose which patches containing p we will use for the reconstruction. Indeed, the patches will contain more or less

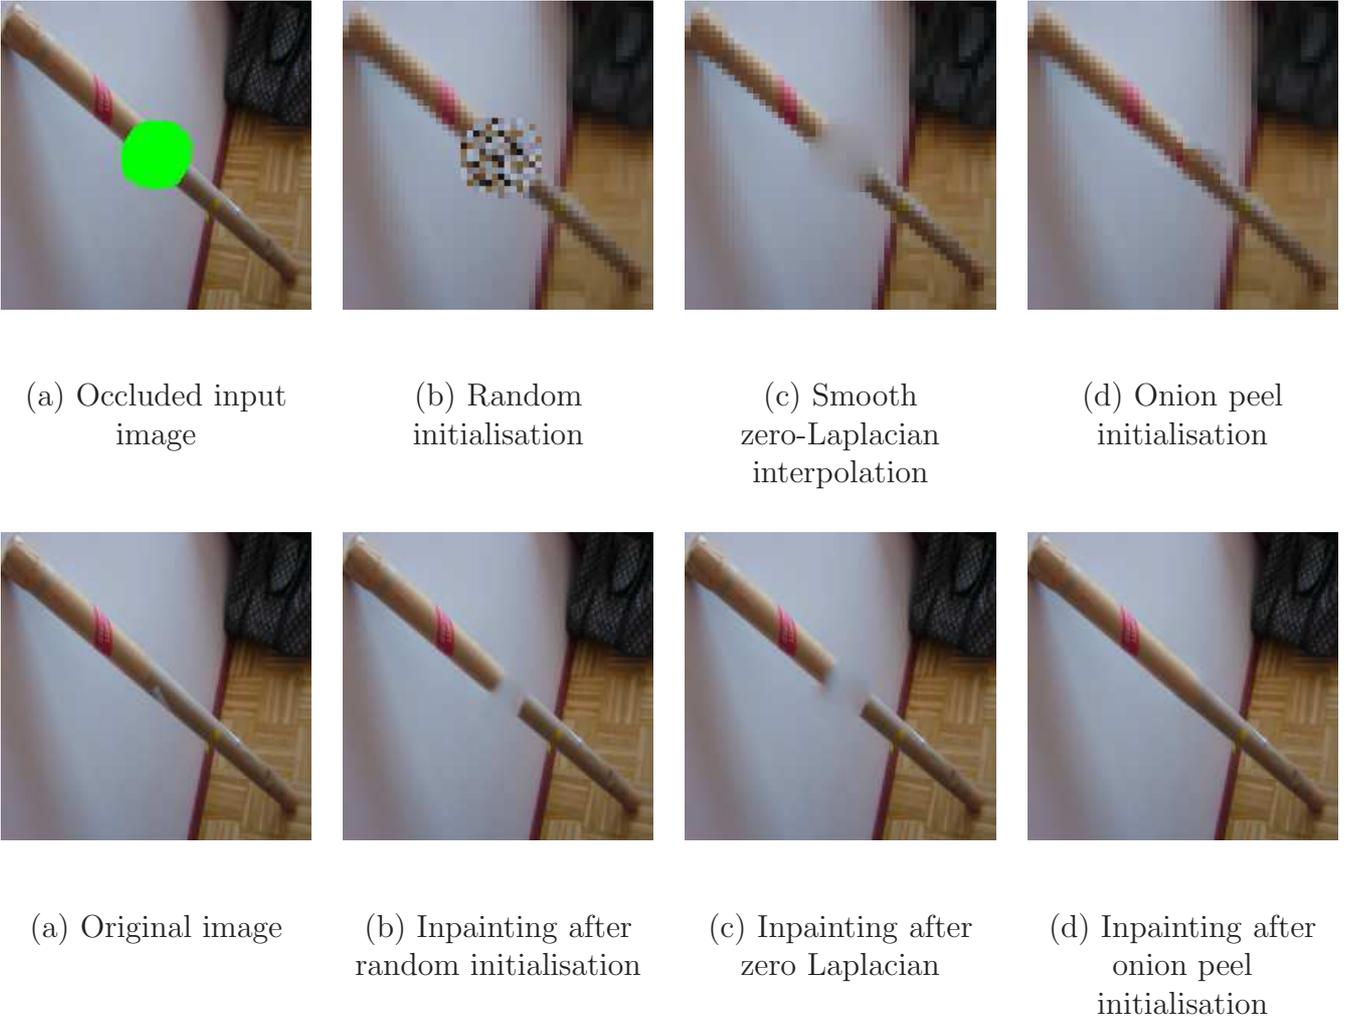


Figure 1: **Impact of initialisation schemes on inpainting results.** We compare the proposed approach with a random initialisation and a smooth interpolation of the conditions on the border of the occlusion (“zero-Laplacian interpolation”). The onion-peel approach correctly continues the cardboard object.

known pixels, and thus be more or less reliable. We propose to use only those patches which are centred at *known* pixels. Thus, we have:

$$u_p = \frac{\sum_{q \in \mathcal{N}'_p} s_p^q u(p + \phi(q))}{\sum_{q \in \mathcal{N}'_p} s_p^q}. \quad (11)$$

The texture features are reconstructed in the same fashion. In Figure 1, we compare the results of our algorithm using different initialisations. In this case, as in many of our other experiments, it is clear that the use of an onion peel approach provides the best initialisation when compared with other common approaches.

3.5 Multiscale scheme

The use of a multiscale scheme is of crucial importance in avoiding local minima of the energy functional (1), and in converging to more satisfactory solutions. Therefore, care must be taken in

Algorithm 2: Complete proposed inpainting algorithm.

Data: Input image u , occlusion \mathcal{H} **Result:** Inpainted image

```

{ $u^\ell$ }_{ $\ell=1$ } $L$   $\leftarrow$  ImagePyramid( $u$ );
{ $T^\ell$ }_{ $\ell=1$ } $L$   $\leftarrow$  TextureFeaturePyramid( $u$ ); // Eq.7
{ $\mathcal{H}^\ell$ }_{ $\ell=1$ } $L$   $\leftarrow$  OcclusionPyramid( $\mathcal{H}$ );
 $\phi^L \leftarrow$  Random;
( $u^L, T^L, \phi^L$ )  $\leftarrow$  Initialisation( $u^L, T^L, \phi^L, \mathcal{H}^L$ ); // Eq.
for  $\ell = L$  to 1 do
   $k = 0, e = 1;$ 
  while  $e > 0.1$  and  $k < 20$  do
     $v = u^\ell;$ 
     $\phi^\ell \leftarrow$  ANNsearch( $u^\ell, T^\ell, \phi^\ell, \mathcal{H}^\ell$ ); // Alg.1
     $u^\ell \leftarrow$  Reconstruction( $u^\ell, \phi^\ell, \mathcal{H}^\ell$ ); // Eq.3
     $T^\ell \leftarrow$  Reconstruction( $T^\ell, \phi^\ell, \mathcal{H}^\ell$ );
     $e = \frac{1}{3|\mathcal{H}^\ell|} \|u_{\mathcal{H}^\ell}^\ell - v_{\mathcal{H}^\ell}\|_2;$ 
     $k \leftarrow k + 1;$ 
  end
  if  $\ell = 1$  then
     $u \leftarrow$  FinalReconstruction( $u^1, \phi^1, \mathcal{H}$ ); // Eq.5
  else
     $\phi^{\ell-1} \leftarrow$  UpSample( $\phi^\ell, 2$ ); // Sec.3.5
     $u^{\ell-1} \leftarrow$  Reconstruction( $u^{\ell-1}, \phi^{\ell-1}, \mathcal{H}^{\ell-1}$ );
     $T^{\ell-1} \leftarrow$  Reconstruction( $T^{\ell-1}, \phi^{\ell-1}, \mathcal{H}^{\ell-1}$ );
  end
end

```

its implementation. The subsampling factor is $\frac{1}{2}$, and we use a Gaussian 3×3 filter with a standard deviation of 1.5 pixels for smoothing.

Two further details are of particular importance: the manner in which an inpainting solution is passed from a coarser pyramid level to a finer one, and the number of pyramid levels to use.

In our implementation, we upsample the *shift map* ϕ rather than the inpainting solution itself. We do this to avoid the inevitable smoothing which would happen if we took the second option. Indeed, the correct textures are not present in the coarser level (as we mentioned in Section 3.3), and therefore using this as an initialisation could encourage the algorithm to use smoother areas for inpainting. The shift map is upsampled using a simple nearest neighbours interpolation, with non-integer coordinates resulting from this interpolation being rounded down to the nearest integer. Then, the initial solution is produced by reconstructing the image at the new resolution with Equation (3). The same process is applied to the texture features.

The next choice which must be made is the number of pyramid levels. This subject has been given very little thought in the inpainting literature, and generally the number of levels is decided by fixing a minimum image resolution and subsampling until that resolution is attained. In reality, this heuristic is not a very good way to ensure that good minima are produced. For example, if the preset resolution is already attained in the original image, and the occlusion size is large, then structures will not be correctly continued into the occlusion.

A better criterion is the relationship between the patch size and the occlusion size. Roughly speaking, the patch size needs to “cover” a certain amount of the occlusion for structures to be well continued. In very simple cases it can be shown [15] that the occlusion size should be a little more than twice the patch size for maintaining structures.

In practice, the occlusion “size” is not necessarily trivial to calculate. To determine this, we iteratively erode the occlusion with a square structuring element of width three pixels until it disappears. Let N_o be the number of times that the occlusion needs to be eroded for it to disappear. We define the number of pyramid levels as:

$$L = \frac{\log \frac{2N_o}{N}}{\log(2)}. \quad (12)$$

This method was used for all of our experiments.

This concludes the algorithmic description of the proposed method. The pseudo-code for the complete algorithm may be seen in Algorithm 2.

3.6 Algorithm parameters

The main tunable parameter of our algorithm is the patch size. Unless specified otherwise, we use a default patch size of 7×7 , which is adequate for most images we used (from around 512×512 to 800×800 pixels). In Figure 3, we show the influence of the patch size on inpainting results.

The parameters of the PatchMatch algorithm are considered to be fixed. We set the random search reduction window to be $\rho = \frac{1}{2}$. We have used ten iterations of propagation/random search in PatchMatch, although this can be reduced if quicker results are required.

We use the procedure described in Section 3.5 to determine the number of pyramid levels automatically, although this can be set manually if required. We set the weight λ associated with the texture features to 50. Finally, we identify the convergence of the algorithm at a pyramid level by determining the average absolute pixel value difference for each colour channel between the current solution and the previous one. If this value drops below 0.1 or the number of iterations exceeds ten, we stop iterating at the current level.

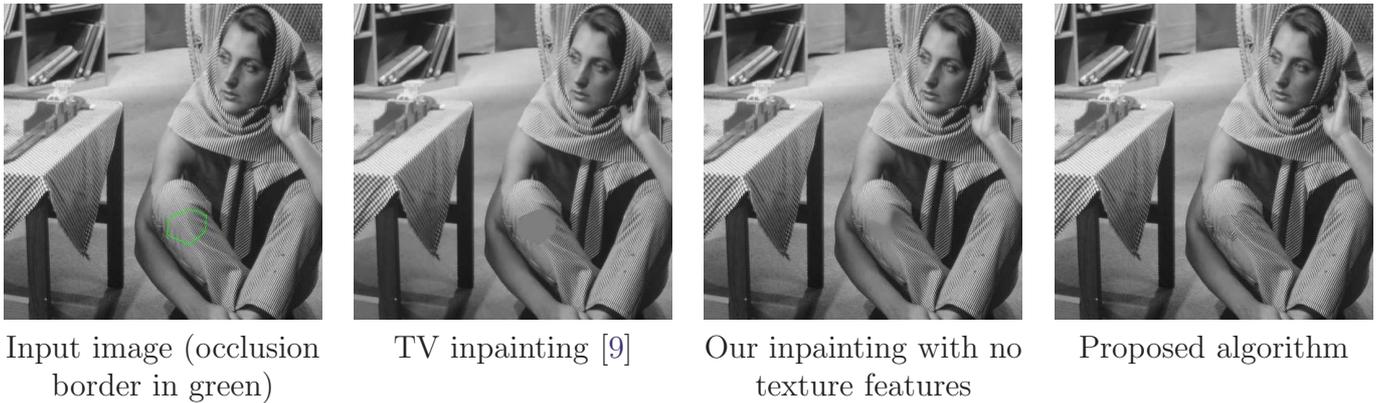


Figure 2: The importance of correctly inpainting textures. PDE-based inpainting methods, which are not designed for dealing with textures, fill the occlusion smoothly. Without the use of texture features, our algorithm has trouble identifying the correct patches to use for inpainting.

4 Results

We now show some visual inpainting results, and discuss the parameters of our algorithm. We show how the results vary according to these parameters. Firstly, we inspect the behaviour of our algorithm in different cases which reflect and justify the different concerns which we addressed in Section 3. In particular, we look at the importance of dealing with textures correctly and the influence of the patch size on inpainting results. Finally, we compare our results with other commonly used and successful inpainting methods.

4.1 Textures

In Figure 2 we show some results of inpainting a medium-sized textured zone with different approaches. The necessity of using an approach which can deal with textures is clear when we observe the result of a PDE-based approach, which is not designed for such situations. Furthermore, it is clear that even a non-local patch-based approach which does not account for textures (our algorithm without texture features) also fails at recreating textures. The proposed algorithm correctly identifies the textures needed for a good inpainting solution. While seams in the solution remain visible, our result is a clear improvement over a PDE-based approach or a patch-based multi-scale approach without texture features, in this case as in many others we have tested.

4.2 Influence of the patch size

The choice of the patch size is the main, tunable parameter of our algorithm. This patch size has a great influence on the ability of the algorithm to reconstruct structures in a satisfactory manner. If too small a size is chosen, then large, global structures in the image will be poorly continued. If the patch size is too large, the seams between different regions will contain visible inconsistencies. An illustration of these concerns may be seen in Figure 3. With a patch size of 3×3 , the continuation of the grass/concrete is done poorly, and the global coherence of the solution suffers. However, at higher patch size (9×9 and 11×11), the patch tends to bring large-scale objects (the man’s leg) into the occluded area, because the patches reach too far. Optimal results are achieved with a patch size of 5×5 - 7×7 . These are usually good patch sizes in most situations.

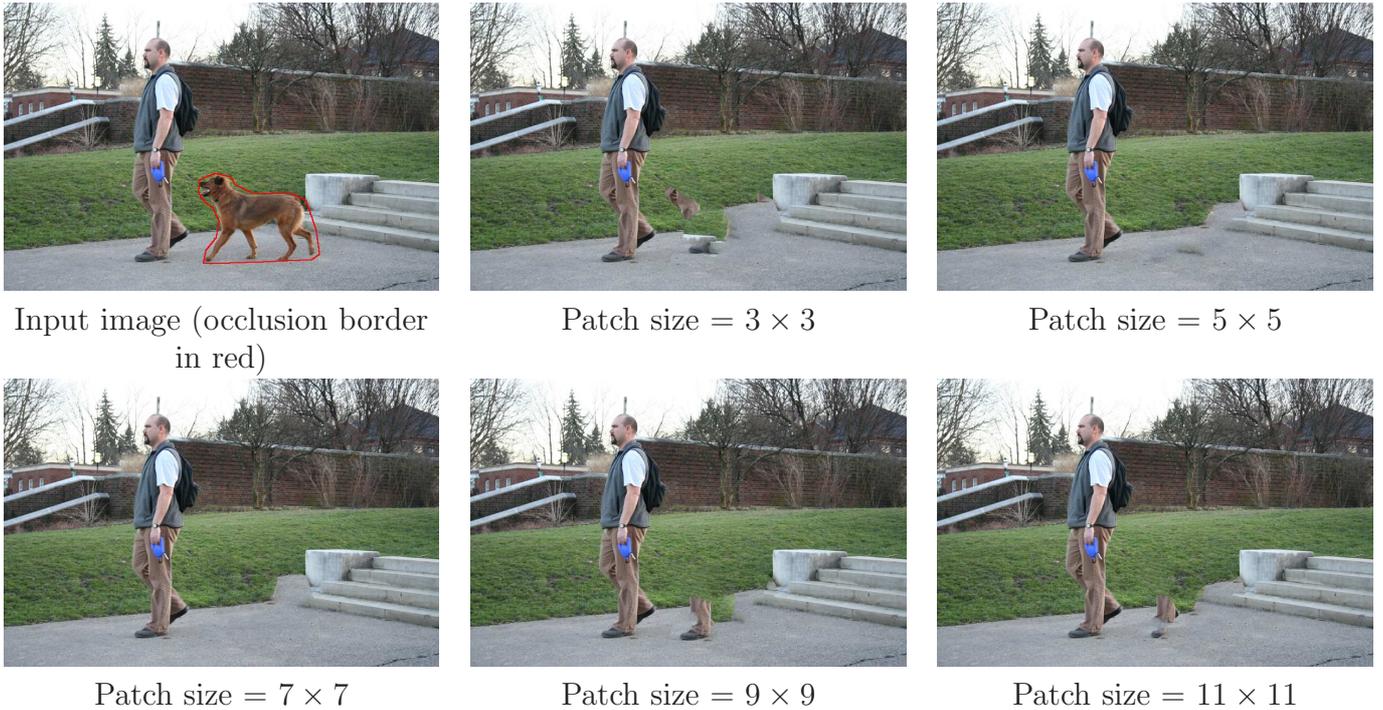


Figure 3: The influence of the patch size. The chosen patch size has a great influence on the ability of the algorithm to reconstruct structures correctly. We show the results of our algorithm with increasing patch sizes.

4.3 Comparison with other similar approaches

In Figure 2 we have compared our algorithm to a PDE-based method [9] in order to show that textures need to be dealt with carefully in image inpainting. However, this comparison is not quite fair, since PDE-based algorithms are not designed for this purpose. Therefore, we show in Figure 4 an example of inpainting in a very difficult situation using two very successful patch-based methods: that of Criminisi *et al.* [6] and the more recent approach of Liu and Caselles [13]. This example is extremely difficult, as there is not much “correct” information to use for inpainting. However, the results of the different algorithms provide significant insight into the advantages and disadvantages of the algorithms. We also provide comparative execution times for each algorithm, to illustrate the speed/quality tradeoff to be expected,

The work of Criminisi *et al.* is a greedy patch-based method. It is very successful at producing good results in the case of small to medium-sized occlusions. However, due to the lack of a global optimisation scheme, the results tend to be very repetitive for larger occlusions, which can be observed in Figure 4. The work of Liu and Caselles is very similar to ours, in that they employ a multi-scale patch-based approach. However, they use a discrete labelling optimisation approach, which is extremely slow. Therefore, they are forced to find a solution at the coarsest pyramid level and only modify it slightly as the resolution increases. Thus, their solution can easily be trapped in a “bad” local minimum, which translates in this case to a large piece of the image being copied into the occlusion. Our algorithm does not provide a perfect solution either, however it manages to avoid the traps of the previous methods. It provides a smoother transition between different occluded areas, which is due to the greater flexibility in the multi-scale optimisation scheme.

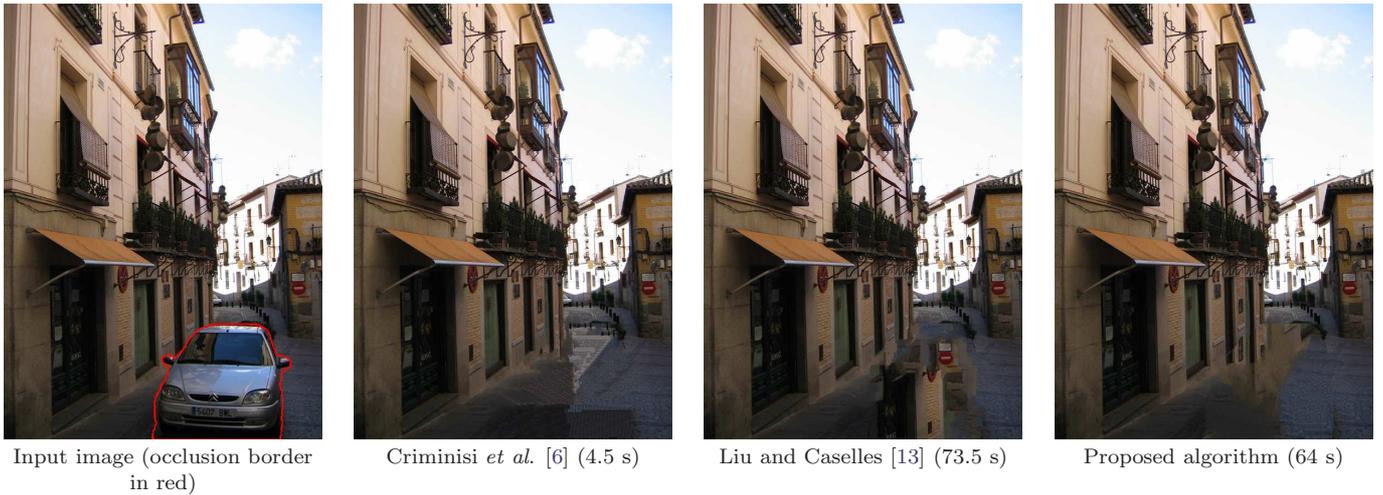


Figure 4: Behaviour of different algorithms in a challenging situation. This example was taken from the dataset of Hays and Efros [10].

5 Conclusion

We have proposed an iterative patch-based multi-scale approach to image inpainting. Our algorithm is based on the minimisation of a global patch-based functional. In order to obtain results of high quality with such an approach, several issues have been addressed. In particular, we describe the method in which we search for nearest neighbour patches, something which is crucial for achieving results in realistic times. We have addressed the issue of correctly comparing textured patches, and shown the importance of doing this in real examples. Finally, we provide specific details concerning initialisation issues and multi-scale choices, both of which have a great influence on the inpainting results. The resulting algorithm can deal with a variety of difficult inpainting situations, as demonstrated by the experimental results, and can be verified with the online demo.

References

- [1] P. ARIAS, G. FACCILOLO, V. CASELLES, AND G. SAPIRO, *A Variational Framework for Exemplar-Based Image Inpainting*, International Journal of Computer Vision, 93 (2011), pp. 319–347.
- [2] C. BALLESTER, M. BERTALMIO, V. CASELLES, G. SAPIRO, AND J. VERDERA, *Filling-in by joint interpolation of vector fields and gray levels*, IEEE Transactions on Image Processing, (2001), pp. 1200–1211.
- [3] C. BARNES, E. SHECHTMAN, A. FINKELSTEIN, AND D. B. GOLDMAN, *PatchMatch: a randomized correspondence algorithm for structural image editing*, ACM Transactions on Graphics, (2009).
- [4] M. BERTALMIO, G. SAPIRO, V. CASELLES, AND C. BALLESTER, *Image Inpainting*, in ACM Transactions on Graphics, 2000, pp. 417–424.
- [5] R. BORNARD, E. LECAN, L. LABORELLI, AND J. H. CHENOT, *Missing data correction in still images and image sequences*, in ACM International Conference on Multimedia, ACM, 2002, pp. 355–361.

- [6] A. CRIMINISI, P. PEREZ, AND K. TOYAMA, *Object removal by exemplar-based inpainting*, in Conference on Computer Vision and Pattern Recognition, vol. 2, June 2003, pp. 721–728.
- [7] IDDO DRORI, DANIEL C. OR, AND HEZY YESHURUN, *Fragment-based image completion*, ACM Transactions on Graphics, 22 (2003), pp. 303–312.
- [8] A. A. EFROS AND T. K. LEUNG, *Texture synthesis by non-parametric sampling*, in International Conference on Computer Vision, vol. 2, 1999, pp. 1033–1038.
- [9] PASCAL GETREUER, *Total Variation Inpainting using Split Bregman*, Image Processing On Line, 2 (2012), pp. 147–157.
- [10] JAMES HAYS AND ALEXEI A. EFROS, *Scene Completion Using Millions of Photographs*, in ACM Transactions on Graphics (TOG), ACM, 2007, p. 4.
- [11] N. KOMODAKIS AND G. TZIRITAS, *Image Completion Using Efficient Belief Propagation Via Priority Scheduling and Dynamic Pruning*, IEEE Transactions on Image Processing, 16 (2007), pp. 2649–2661.
- [12] VIVEK KWATRA, IRFAN ESSA, AARON BOBICK, AND NIPUN KWATRA, *Texture Optimization for Example-based Synthesis*, ACM Transactions on Graphics, 24 (2005), pp. 795–802.
- [13] YUNQIANG LIU AND V. CASELLES, *Exemplar-Based Image Inpainting Using Multiscale Graph Cuts*, IEEE Transactions on Image Processing, 22 (2013), pp. 1699–1711.
- [14] S. MASNOU AND J. M. MOREL, *Level lines based disocclusion*, in International Conference on Image Processing, vol. 3, Oct. 1998, pp. 259–263.
- [15] ALASDAIR NEWSON, *On Video Completion: Line Scratch Detection in Films and Video Inpainting of Complex Scenes*, PhD thesis, Télécom ParisTech, March 2014.
- [16] ALASDAIR NEWSON, ANDRÉS ALMANSA, MATTHIEU FRADET, YANN GOUSSEAU, AND PATRICK PÉREZ, *Video inpainting of complex scenes*, SIAM Journal of Imaging Sciences, 7 (2014).
- [17] Y. PRITCH, E. KAV-VENAKI, AND S. PELEG, *Shift-map image editing*, in International Conference on Computer Vision, Sept. 2009, pp. 151–158.
- [18] Y. WEXLER, E. SHECHTMAN, AND M. IRANI, *Space-time video completion*, in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, June 2004, pp. 120–127.
- [19] —, *Space-Time Completion of Video*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 29 (2007), pp. 463–476.