

Parameter Optimisation for Texture Completion

Hoang M. Nguyen, Burkhard C. Wünsche, Patrice Delmas, and Christof Lutteroth
Dept of Computer Science, University of Auckland, Auckland, New Zealand
justin.nguyen@auckland.ac.nz {burkhard, p.delmas, lutteroth}@cs.auckland.ac.nz

Abstract—In 3D reconstruction applications frequently large mesh sections are missing. While the geometry can be created with mesh completion techniques, reconstructing the texture is more difficult. An exemplar-based synthesis usually requires that the missing region contains only one texture pattern. Image inpainting techniques work for more complex textures, but were designed primarily to fill narrow missing regions and tend to produce undesirable results when the missing region is large. We solve this problem by exploiting the fact that most 3D objects contain different surface regions with similar textures. We use an appearance space to identify suitable texture patches fitting with the boundary of the missing region, and fill the region by using a patch-based synthesis combined with Poisson interpolation. We evaluate the effect of different parameters and demonstrate its improved performance over existing inpainting techniques.

I. INTRODUCTION

Over the past couple of years 3D reconstruction technologies have improved tremendously and are now available in the consumer-level domain. Examples include structured lighting (Kinect), laser scanners, and image-based modeling. In many applications the resulting mesh contains large missing regions, e.g. because the surface was invisible to the utilised sensor. Surface reconstruction techniques, such as Poisson Surface Reconstruction, and mesh completion techniques can create a watertight surface, but are unable to reconstruct texture information.

Missing textures can be obtained with texture synthesis methods, but these usually employ a single exemplar and assume a consistent texture over the missing region. Image inpainting techniques work for more complex textures, but were designed primarily to fill narrow missing regions and tend to produce undesirable results when the missing region is large.

3D objects in the real-world usually contain repeating geometries and textures, e.g. the sides of an animal, components of a plant (stem, branch, blossom), or man-made objects which are usually designed using symmetries and self-similarities (e.g. walls and windows of a house).

In this paper we perform texture completion by exploiting the fact that most 3D objects contain different surface regions with similar textures. We use an appearance space to identify suitable texture patches fitting with the boundary of the missing regions, and fill the region by using a patch-based synthesis combined with Poisson interpolation.

The remainder of this paper is organised as follows: After a brief discussion of existing image inpainting techniques in

section II, we describe our inpainting algorithm in section III. Section IV presents results and section V concludes this paper.

II. RELATED WORK

Image in-painting techniques can be divided into two classes. Pixel-based methods fill the missing region pixel-by-pixel starting from the boundary, whereas patch-based methods usually add entire patches and minimise discontinuities between texture regions.

The arguably best known and most successful algorithm amongst pixel-based inpainting methods was proposed by Bertalmio *et al.* [1]. The authors attempt to replicate manual inpainting by propagating the known color values into a missing region along so called isophotes, representing the smallest spatial change of color values and structures.

Drori *et al.* [2] use adaptive circular fragments to operate on different scales to capture both global and local structures and approximate the missing region.

Telea *et al.* [3] present an inpainting technique based on a fast marching method for level set applications. The method is simple and considerably more efficient than other pixel-based methods.

Pixel-based methods often fail to properly reconstruct larger structures with semantic meaning, e.g. leaves. Criminisi *et al.* [4] iteratively select a “best-fit” rectangular patch and copy it over to the target region. The order in which boundary pixels of the missing region are processed is based on the amount of information available for that pixel and whether it has any prominent features.

Cheng *et al.* [5] update the priority equation of [4] and made it adjustable to the structural and textural information specific to an image. Ignacio *et al.* [6] extend the concept of Criminisi’s method and apply it in the wavelet domain.

III. ALGORITHM

We employ a similar patch search and insertion concept as Criminisi *et al.*, but use appearance space attributes and principal component analysis to improve region matching quality and speed. In addition, our method smoothly fuses patches together to remove all visible seams.

A. Candidate Patch Identification

The algorithm’s performance is significantly effected by the ability to identify the patch in the image that retains the highest resemblance to the processed patch. This is achieved by iteratively traversing through each pixel of the image outside the missing region and computing the similarity of the patch

centered around that pixel and the original patch. Instead of using the standard *Sum of Squared Differences (SSD)* to measure the similarity of two given patches, we employ *appearance space attributes*, which provide much more information and thus improve the search result.

When searching for a matching patch, we consider for each pixel an 11×11 pixel neighbourhood. For each pixel of this neighborhood we consider *RGB* colours, the *gradient vector*, as well as the signed Euclidean distance to the closest dominant feature in the original texture.

Gradient Vector: We estimate the gradient of an intensity image \mathbf{I} , $(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y})$ by using the convolution kernels of a standard Sobel operator. The two kernels for the x and y directions are:

$$S_x = (1 \ 2 \ 1)^T \cdot (1 \ 0 \ -1) \quad (1)$$

$$S_y = (1 \ 0 \ -1)^T \cdot (1 \ 2 \ 1) \quad (2)$$

The attribute gradient image is then defined as:

$$\mathbf{G}(p) = ((S_x * \mathbf{I})(p), (S_y * \mathbf{I})(p)) \quad (3)$$

where p denotes an image pixel.

Signed Feature Distance: The signed feature distance of a pixel p is defined as the distance of p to the closest pixel $q \in \mathbf{M}$ for which $\mathbf{M}(p) \neq \mathbf{M}(q)$ where \mathbf{M} is a binary image created by applying the *canny edge detection* method [7] on the input image.

The distance between two pixels is defined using the following equations (adapted from [8]):

$$\mathbf{F}(p) = s \cdot (\max_{r \in \mathbf{M}} \lambda(r) - \lambda(p)) \quad (4)$$

$$\lambda(x) = \frac{1}{|q_x - x|} \quad (5)$$

$$s = \begin{cases} 1 & \text{if } \mathbf{M}(p) = 1 \\ -1 & \text{if } \mathbf{M}(p) = 0 \end{cases} \quad (6)$$

The vector $(q_x - x)$ points from x to the closest pixel q in \mathbf{M} .

The entire attribute information is encapsulated into a $11 \times 11 \times (3 + 2 + 1) = 726$ -dimensional vector. Determining the similarity of two given patches by comparing two 726-dimensional vectors is not efficient. In order to make the appearance space more practicable, the 726 dimensional vectors are projected into low-dimensional vectors using *principal component analysis (PCA)* ([8], [9]). In our method, the dimensionality is reduced to 12, which in our experiments with different image types produced the best results.

The clear advantage of the attribute space over the conventional SSD is that the attribute space approach permits any meaningful information about the pixels and their surrounding to be embedded for matching purposes. By reducing the dimensionality, the computation time can be kept manageable.

B. Patch Fusion

The final step is to replicate the content of the candidate patch and smoothly blend it with the target region. We employ a *Poisson-guided interpolation* approach proposed by [10] for this task.

The goal is to adjust the colour information of patch Ψ_B , while preserving the *relative information* (image gradient) as much as possible, so that the transition between the newly modified patch Ψ_C and the rest of the image is gracefully blended.

IV. EVALUATION

In this section, we investigate the effect of different algorithm parameters and evaluate its performance over popular existing algorithms.

A. Effect of Parameters




Fig. 1. The original image is on the left, while the damaged image is on the right

Effect of appearance space attributes: Figure 2 shows an example in which the previous damaged “lizard” image (Figure 1) is inpainted using various different appearance space attributes. The following combinations of appearance space attributes are used for this testing purpose.

- Case 1: RGB color, HSB color, signed feature distance, horizontal and vertical gradient vectors.
- Case 2: RGB color, signed feature distance, horizontal and vertical gradient vectors.
- Case 3: RGB color and signed feature distance.
- Case 4: RGB color, horizontal and vertical gradient vectors.

As can be observed from figure 2, there is little or no visual difference between the first and second case. This indicates that adding the HSB color attribute does not result in finding better patches.

There is an obvious difference in the result quality when the gradient vectors are removed. Some parts of the texture of the head become fuzzy and appear incorrect. A small part of the line on the body of the lizard is missing.

The result quality further deteriorates when removing signed feature distance attribute instead of the gradient vectors. The textures become much more blurry and more missing features are noted. This is probably due to the fact that incorrect patches have been identified and blended in this case.

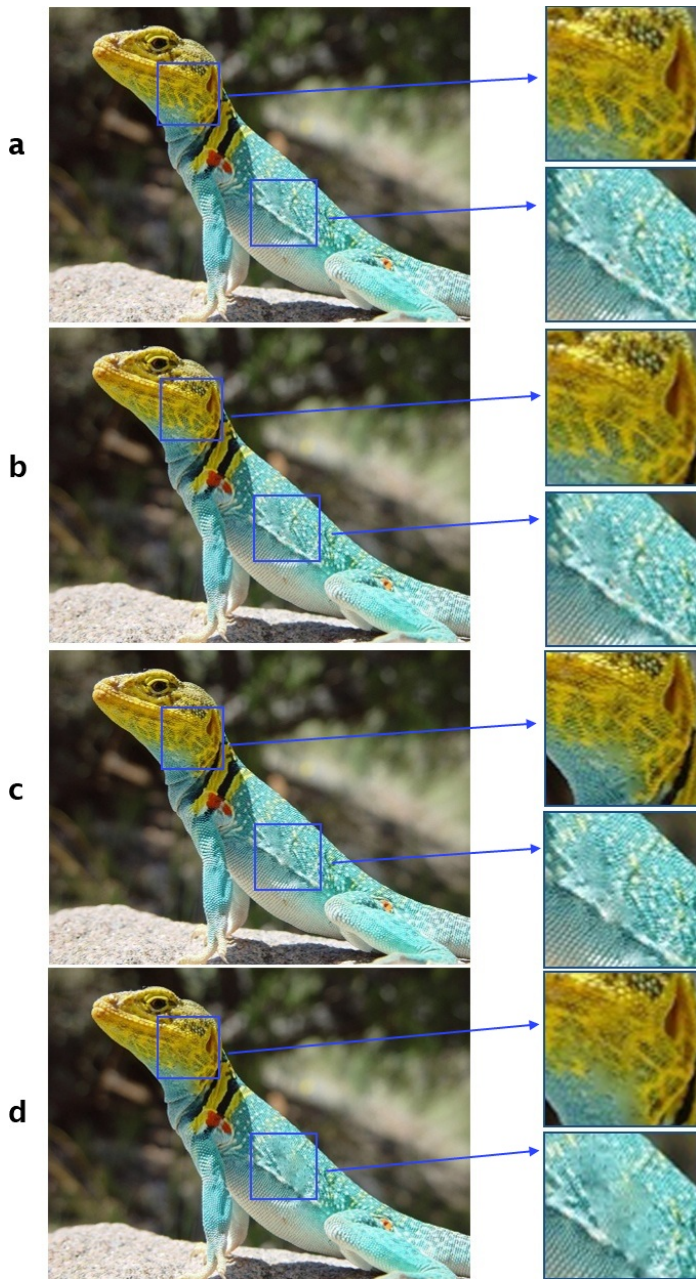


Fig. 2. Appearance Space Attribute Parameters. a) Color, HSB color, signed feature distance, horizontal and vertical gradient vectors. b) Color, signed feature distance, horizontal and vertical gradient vectors. c) Color and signed feature distance. d) Color, horizontal and vertical gradient vectors.

The results indicate that the optimal combination of appearance space attributes is *RGB color, signed feature distance, horizontal and vertical gradient vectors*. Adding more attributes often contributes little or no improvement and yet increases the computational cost.

Effect of attribute vectors' dimensionality: In order to improve the efficiency, appearance attribute vectors are projected to a low-dimensional vectors before being compared to others. In this section, we evaluate the effects of the dimensionality of reduced appearance attribute vectors on the result quality.

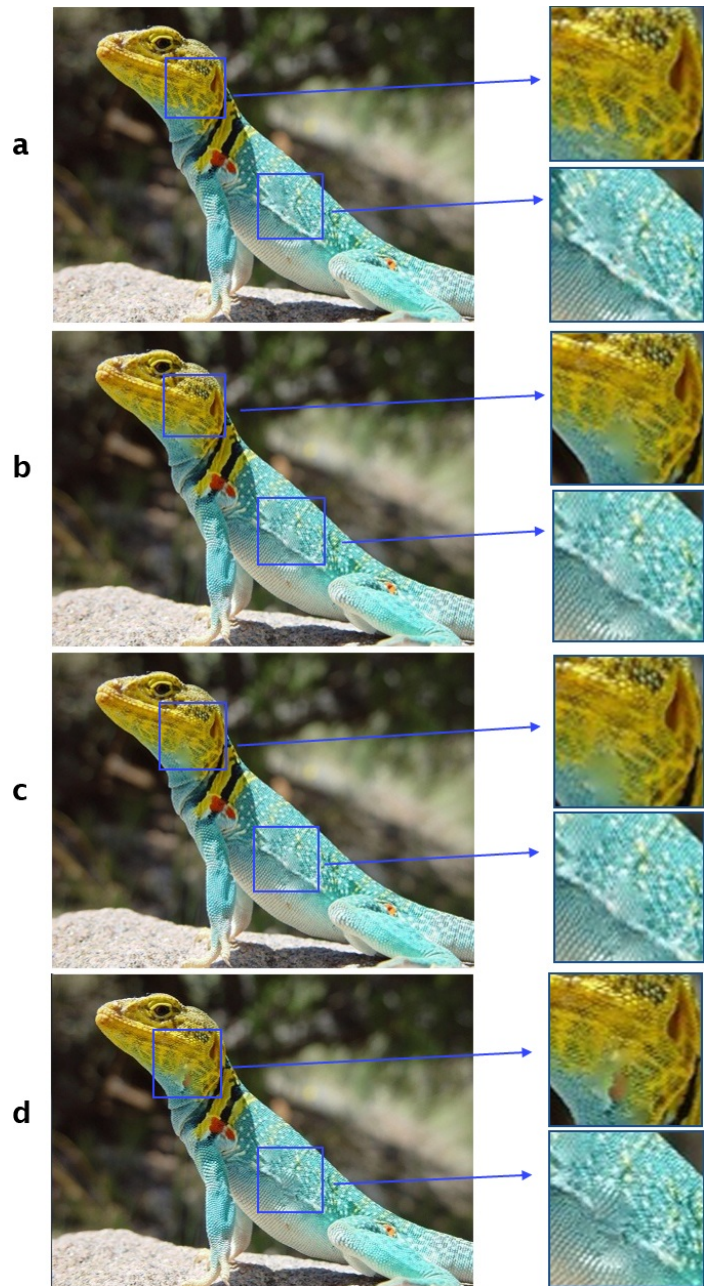


Fig. 3. Appearance Space Attribute Dimensions: a) 8-dimensional vectors. b) 7-dimensional vectors. c) 6-dimensional vectors. d) 5-dimensional vectors.

Figure 3 shows an example in which appearance attribute vectors were reduced to different dimensions (8, 7, 6, 5) before comparison.

In the first case, reducing the original 726-dimensional attribute vectors to a 8-dimensional vector yields good results. Missing textures are reconstructed well and seamlessly fused with the existing textures. We found that using more than 8 dimensions does not result in a visible improvement of the texture completion result.

Reducing the dimension further to 6 or 7 increases blurriness. This is probably due to the fact that the reduced vectors did not contain enough information for correct patches to be

found.

In the case where only 5 dimensions are retained, deformations have started to appear due to the lack of information required for correct patch detections.

Effect of Patch Size: Figure 4 demonstrates the difference in the result quality when varying the patch sizes. We found that in most cases the ideal patch size is the range of 7 and 11. Inpainted images in these cases are often well-reconstructed. As the size of the patches increases, the painted regions tend to become more blurry. This is probably because the larger the patch size is, the more unrelated features from surrounding regions are involuntarily taken into account leading to less accurate patch to be selected. Patch sizes smaller than 7 increase the computational cost while contribute little improvement in term of the quality.




Fig. 4. Effect of patch sizes on the result quality.

B. Evaluation against Other Inpainting Methods

In this section, we evaluate the performance of our method against some of the best known image inpainting methods described in the literature (Figure 5).

Bertalmio’s method [1] was not very successful in this test. It was able to reconstruct small parts of the missing regions

(near the boundary), but failed to interpolate textures further for the middle regions resulting in patches with colours not consistent with the regions neighborhood.

Telea’s method [3] outperformed Bertalmio’s method both in term of the reconstruction quality and efficiency. However, the inpainted regions appear very blurry and unrealistic. This is expected as pixel-based methods were designed to tackle only narrow missing regions.

Standard exemplar-based inpainting method (Criminisi [4]) performed reasonably well except for some anti-facts. These anti-facts are caused by incorrect patches that were selected as a result of employing the conventional *Sum of Squared Differences (SSD)*.

Our algorithm performed well in this test case. Although the inpainted region still exhibits slight blurriness, the overall structure of different scene components has been correctly recovered.

V. CONCLUSION AND FUTURE WORK

We have presented a novel image inpainting algorithm for synthesize large missing texture regions from digital images. The results of this inpainting process is a new image in which the deterioration has been “inpainted” and reverted in such a way that few visible traces of it remain. The basic idea of our approach is to replicate missing textures by looking for “best-fit” texture patches in the source regions and smoothly insert these patches into the missing region to produce the final result.

Our solution offers two major improvements compared to existing techniques. Patches for filling in missing regions are found using an appearance space vector, which not only encodes colour differences between regions, but also colour gradients, feature distances and other measures for image similarity. The second major improvement is the technique used combine patches filling in a missing region. We use a Poisson-guided interpolation to blend patche to avoid visible seams.

We have evaluated our method’s performance against some of the best known inpainting methods described in the literature and found that our results are superior.

REFERENCES

- [1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, “Image inpainting,” *In Proceeding SIGGRAPH ’00 Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pp. 417–424, 2000.
- [2] I. Drori, D. Cohen-Or, and H. Yeshurun, “Fragment-based image completion,” *ACM Transactions on Graphics*, pp. 303–312, 2003.
- [3] A. Telea, “An image inpainting technique based on the fast marching method,” *Journal of Graphics Tool*, vol. 9, pp. 23–34, 2004.
- [4] A. Criminisi, P. Perez, and K. Toyama, “Object removal by exemplar-based inpainting,” *ACM Transactions on Graphics*, pp. 721–728, 2003.
- [5] W.-H. Cheng, C.-W. Hsieh, S.-K. Lin, C.-W. Wang, and J.-L. Wu, “Robust algorithm for exemplar-based image inpainting,” *In Proceedings of CGIV*, pp. 64–69, 2005.
- [6] U. Ignecio and C. R. Jung, “Block-based image inpainting in the wavelet domain,” *Visual Computing*, pp. 733–741, 2007.
- [7] C. J, “A computational approach to edge detection,” in *IEEE Trans. Pattern Analysis and Machine Intelligence 1986*, vol. 8, no. 6, 1986, pp. 679–698.

- [8] F. Manke and B. Wunsche, "Analysis of appearance space attributes for texture synthesis and morphing," *Image and Vision Computing New Zealand 2009*, pp. 85–90, 2009.
- [9] S. Lefebvre and H. Hoppe, "Appearance-space texture synthesis," in *ACM SIGGRAPH 2006 Papers*, 2006, pp. 541–548.
- [10] P. Perez, M. Gangnet, and A. Blake, "Poisson image editing," in *ACM Transactions on Graphics*, vol. 22, no. 3, 2003, pp. 313–318.

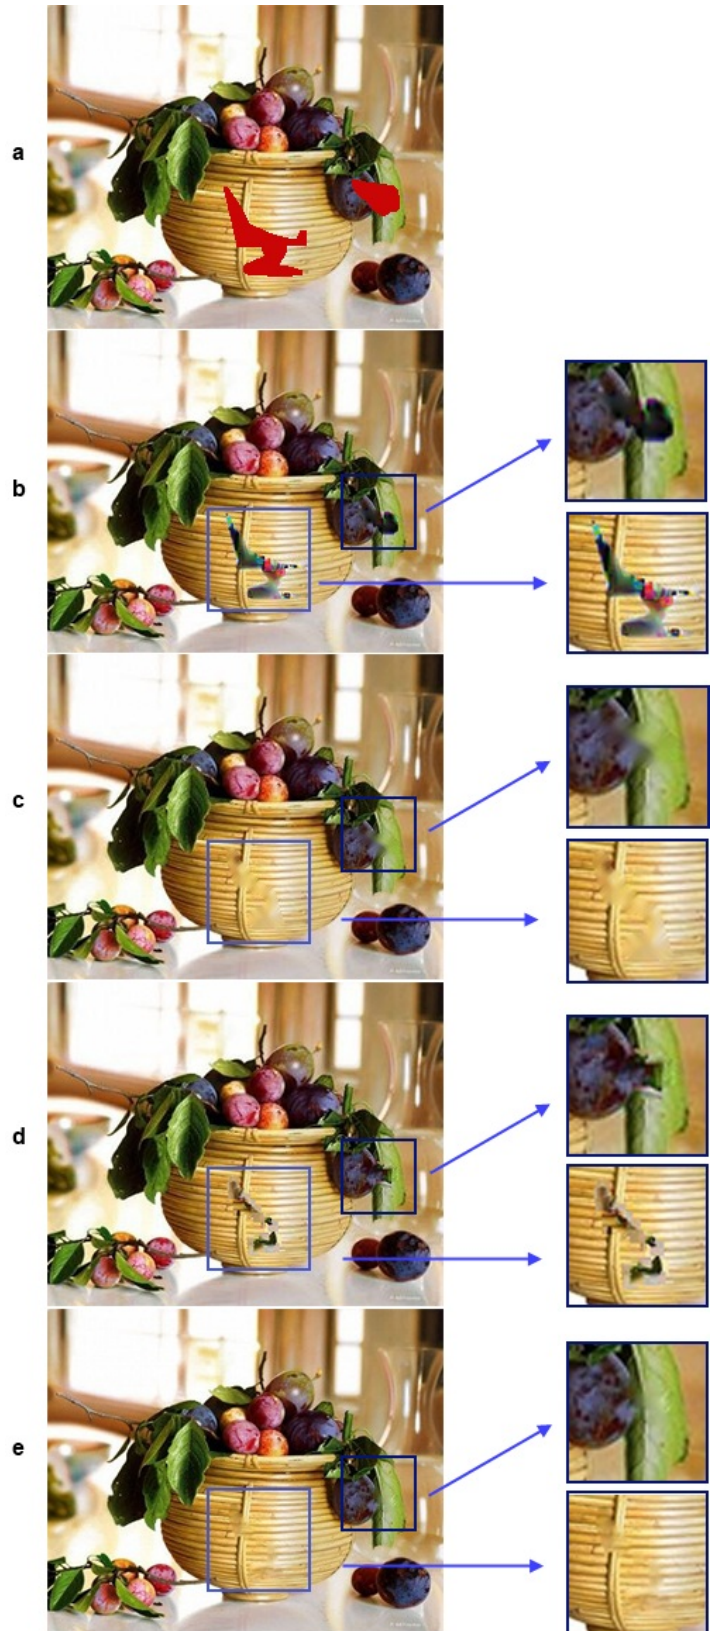


Fig. 5. a) The input image. Image inpainting results obtained using the algorithms from: (b) (Bertalmio et al. 2000), (c) (Telea et al. 2004), (d) (Criminisi et al. 2003) and (e) our method.