

A NONLOCAL-MEANS APPROACH TO EXEMPLAR-BASED INPAINTING

Alexander Wong¹ and Jeff Orchard²

University of Waterloo

¹Systems Design Engineering, ²Cheriton School of Computer Science
Waterloo, Ontario, Canada

ABSTRACT

This paper introduces a novel approach to the problem of image inpainting through the use of nonlocal-means. In traditional inpainting techniques, only local information around the target regions are used to fill in the missing information, which is insufficient in many cases. More recent inpainting techniques based on the concept of exemplar-based synthesis utilize nonlocal information but in a very limited way. In the proposed algorithm, we use nonlocal image information from multiple samples within the image. The contribution of each sample to the reconstruction of a target pixel is determined using an weighted similarity function and aggregated to form the missing information. Experimental results show that the proposed method yields quantitative and qualitative improvements compared to the current exemplar-based approach. The proposed approach can also be integrated into existing exemplar-based inpainting techniques to provide improved visual quality.

Index Terms— image inpainting, nonlocal-means

1. INTRODUCTION

An increasingly popular area of research in the field of image processing is image inpainting, where the underlying goal is to reconstruct the missing regions within an image in such a way that it is visually plausible to an observer. In most cases, the missing region (called the target region) is filled in using information from the rest of the image (called the source region). Image inpainting is an important part of image processing applications such as scratch and object removal from a photograph. Much of the traditional work in inpainting focused on filling in missing regions through the diffusion of local information [1, 2, 3, 4]. One of the main issues with such techniques is that it is restricted to using local information as a prior. Therefore, in many situations where the local information does not characterize the missing information, the resulting reconstructed information in the missing region will not be visually consistent with the rest of the image.

Newer approaches to the problem of image inpainting have focused on the concept of exemplar-based synthe-

sis [5, 6, 7]. In these techniques, a best match sample from the source region is found and copied directly into the target region. The main benefit of techniques based on exemplar-based synthesis is that they can utilize nonlocal information. This is particularly useful in situations where the correct information for filling in a region is located far away from that region. The main drawback is that the nonlocal information is used in a very limited way. By using only the best match sample, the method runs the risk of choosing a sample that is corrupted, or not a perfect match. However, an image with redundant content could have several samples that could be combined to form a more robust estimate of the missing information. Therefore, we propose that improved inpainting quality can be achieved using nonlocal information from multiple samples within the image.

The main contribution of this paper is a novel approach to exemplar-based inpainting using the concept of nonlocal-means. In this framework, the relative contribution of each sample to the reconstruction of a target pixel is determined using a weighted similarity function and aggregated to form the missing information. Thus, the proposed method makes use of all relevant nonlocal image information. The method can also be integrated into existing exemplar-based inpainting techniques to provide improved visual quality.

2. EXEMPLAR-BASED INPAINTING

Before describing the proposed method, it is important to provide a brief overview of the concepts underlying exemplar-based inpainting techniques. The first attempts at inpainting propagated image information from the source region to the target region through the process of diffusion. Hence, only information that was local to the target region was used. This heavy reliance on local redundancy is problematic in situations where the missing information in the target region is not characterized by local information. Exemplar-based methods address this issue by using nonlocal information from the image to fill in the target region. The overall process of exemplar-based inpainting is illustrated in Fig. 1. Let us denote the target region as Ω and the source region as Θ . Suppose we wish to fill in a patch Ψ_p , centered on pixel p , that is at least partially within the target region Ω . The size of this

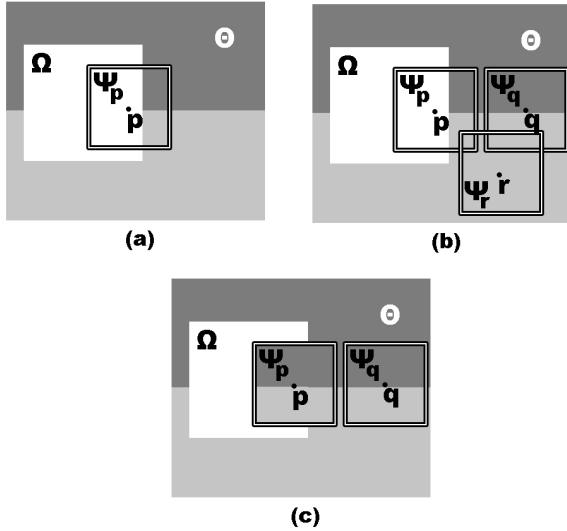


Fig. 1. Information propagation using exemplar-based inpainting: a) Original image, with source region Θ and target region Ω . We wish to fill in patch Ψ_p centered at p . b) Neighborhood surrounding Ψ_p is compared to other neighborhoods such as Ψ_q and Ψ_r . c) Ψ_q was found to be the best match and its contents is copied into Ψ_p .

patch may vary depending on the situations and can be greater than or equal to a single pixel. In exemplar-based inpainting, a neighborhood surrounding Ψ_p (whose size is greater than or equal to Ψ_p) is compared using a similarity measure to other neighborhoods centered around other points within the image (eg., Ψ_q and Ψ_r). For example, the sum of squared differences (SSD) cost function was used in [6]. The patch that minimizes the cost function is selected as the best match and its contents are copied into the part of Ψ_p that lies inside the target region.

3. INPAINTING USING ADAPTIVE NONLOCAL-MEANS

One of the main issues with the current approach to exemplar-based inpainting is the fact that they use image information from only a single neighborhood. They do not fully exploit content redundancy in an image and, thus, “put all their eggs in one basket”. A more robust approach of exemplar-based inpainting is to use image information from multiple samples within the image and weight their contribution according to their similarity to the neighborhood under evaluation. This concept of weighted aggregation of nonlocal information has proven effective for the purpose of image denoising [9, 8]. Let us illustrate the benefit of this nonlocal-means approach over exemplar inpainting. Suppose we have an image depicting a lattice of degraded circles, as shown in Fig. 2. Our goal is to fill in the missing information in the target region (depicted

as the gray square). The inpainting results for the current exemplar-based approach and the proposed nonlocal-means approach are shown in Fig. 3. Due to the way the circles are degraded, the neighborhood of the target (the outline in Fig. 3) best matches the degraded circle to its right. However, this best-match region is, itself, degraded and its error is copied to the target region. However, using the weighted aggregation of multiple neighborhoods yields a much improved result.




Fig. 2. Image of degraded circle lattice. The grey area represents the target region and the red window represents the search neighborhood



Fig. 3. Inpainting results. Left: Current approach of exemplar-based inpainting, Right: Proposed approach using nonlocal-means

The proposed inpainting method can be described in the following manner. Let us suppose that we are given an image consisting of a target region Ω and a source region Θ . Given a patch Ψ_p centered at point p in Ω , and another patch Ψ_q , we compare the patches by computing a (dis)similarity metric, $c(\Psi_p, \Psi_q)$. In our experiments, $c(\Psi_p, \Psi_q)$ is the sum of squared differences between corresponding pixel values in Ψ_p and Ψ_q , excluding any pixels that happen to lie in the target region.

Once the cost function has been evaluated for each point in Θ , the n patches with the lowest cost are selected for the weighted aggregation process. For testing purposes, n is set to 10. This effectively reduces the contribution of dissimilar patches.

The weights for each selected patch Ψ_i are then computed using the weighing function

$$w(\Psi_i) = e^{-\frac{c(\Psi_p, \Psi_i)}{h}}, \quad (1)$$

where h is a chosen decay coefficient. The weighting function is designed to ensure that the contribution of each selected patch is related to its similarity to the patch Ψ_p . Finally, the patch used to fill in the target portion of Ψ_p , denoted Ψ'_p , is

Table 1. PSNR for Test Images

Test	PSNR (dB)		PSNR Gain
	Traditional Approach	Proposed Approach	
TEST1 (19.6545dB)	31.4723	33.1413	+1.6690
TEST2 (20.0380dB)	45.1456	46.1803	+1.0347
TEST3 (27.7129dB)	40.0865	42.5439	+2.4574
TEST4 (18.8935dB)	44.5158	46.6985	+2.1827




Fig. 4. TEST1: Top-left: Original image, Top-right: Image with missing information, Bottom-left: Inpainted using current approach, Bottom-right: Inpainted using proposed approach

computed using the weighted aggregation equation,

$$\Psi_{p'} = \frac{\sum_{i=1}^n w(\Psi_i) \Psi_i}{\sum_{i=1}^n w(\Psi_i)}. \quad (2)$$

4. EXPERIMENTAL RESULTS

A set of four test images was used to evaluate the effectiveness of the proposed method. For comparison, we also ran the same test images through an exemplar-based inpainting technique such as [6]. We used local neighborhoods of size 10×10 for these experiments, and evaluated the performance of the two methods quantitatively using the peak signal-to-noise ratio (PSNR). A summary of the PSNR results is shown in Table 1. The proposed algorithm shows an improvement in PSNR over the current exemplar-based approach for all of the test images. Samples of inpainted images are shown in Figs. 4-8. The proposed method provides good visual quality and appears visually plausible without resulting in image blur. These results show that the proposed method is capable of providing good inpainting performance.




Fig. 5. TEST2: Top-left: Original image, Top-right: Image with missing information, Bottom-left: Inpainted using current approach, Bottom-right: Inpainted using proposed approach




Fig. 6. TEST2 Error: Left: Exemplar-based method, Right: Proposed method. The images have been contrast enhanced to make the differences more visible. The arrows point out places where the two error images differ.

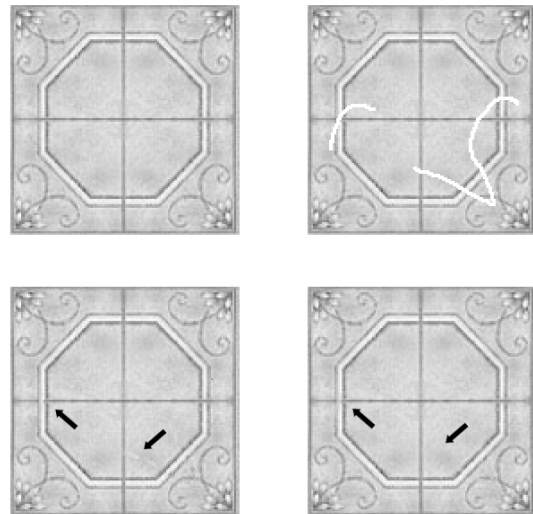


Fig. 7. TEST3: Top-left: Original image, Top-right: Image with missing information, Bottom-left: Inpainted using current approach, Bottom-right: Inpainted using proposed approach. The arrows point out places with very noticeable improvements.




Fig. 8. TEST4: Top-left: Original image, Top-right: Image with missing information, Bottom-left: Inpainted using current approach, Bottom-right: Inpainted using proposed approach. The arrows point out places with very noticeable improvements.

5. CONCLUSIONS

This paper introduced a novel approach to exemplar-based inpainting using nonlocal-means. By utilizing nonlocal image information from multiple samples and selecting the samples used based on the underlying image content, it was demonstrated that improved visual quality can be achieved over the approach used in current exemplar-based inpainting techniques. Finally, the proposed approach can be integrated into existing exemplar-based inpainting techniques to improve visual quality. Future work includes investigating alternative weighing functions and alternative methods for selecting samples being used in the weighted aggregation process.

Acknowledgment

The authors would like to thank the Natural Sciences and Engineering Research Council (NSERC) of Canada for funding this project, as well as the Vision and Image Processing Group at the University of Waterloo.

6. REFERENCES

[1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proc. SIGGRAPH*, pp. 417-424, 2000.

[2] T. Chan and J. Shen, "Non-texture inpainting by curvature-driven diffusions," *J. Visual Comm. Image Rep.*, vol. 4, no. 12, pp. 436-449, 2001.

[3] S. Rane, G. Sapiro, and M. Bertalmio, "Structure and texture filling-in of missing image blocks in wireless transmission and compression applications," in *IEEE Trans. Image Processing*, vol. 12, no. 3, pp. 296-303, 2003.

[4] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, J. Verdera, "Filling-In by Joint Interpolation of Vector Fields and Gray Levels," in *IEEE Trans. Image Processing*, vol. 10, no. 8, pp. 1200-1211, 2001.

[5] A. Criminisi, P. Perez, and K. Toyama, "Object Removal by Exemplar-Based Image Inpainting," *Proc. IEEE Intl. Conference on Computer Vision*, 2003.

[6] A. Criminisi, P. Perez, and K. Toyama, "Region Filling and Object Removal by Exemplar-Based Image Inpainting," *IEEE Trans. Image Processing*, vol. 13, no. 9, pp. 1200-1212, 2004.

[7] J. Wu and Q. Ruan, "Object Removal By Cross Isophotes Exemplar-based Inpainting," *Proc. 18th Intl. Conference on Pattern Recognition*, vol. 3, pp. 810-813, 2006.

[8] A. Buades, B. Coll, and J. Morel, "Nonlocal image and movie denoising," *Intl. J. of Computer Vision*, (to appear), 2007.

[9] M. Mahmoudi and G. Sapiro, "Fast image and video denoising via nonlocal means of similar neighborhoods," *IEEE Signal Processing Letters*, vol. 12, no. 12, pp. 839-842, December 2005.