

This paper is published in the open archive of Mid Sweden University  
DIVA <http://miun.diva-portal.org> to ensure timely dissemination of scholarly and technical  
work. Copyright and all rights therein are retained by authors or by other copyright holders.  
All persons copying this information are expected to adhere to the terms and constraints  
invoked by each author's copyright. In most cases, these works may not be reposted  
without the explicit permission of the copyright holder.

Muddala, S. M.; Sjöström, M.; Olsson, R., "Depth-Based Inpainting For Disocclusion Filling ",  
in *3DTV-Conference*, 2-4 July 2014.

©2014 IEEE. Personal use of this material is permitted. However, permission to  
reprint/republish this material for advertising or promotional purposes of for creating new  
collective works for resale or redistribution to servers or lists, or to reuse any copyrighted  
component of this work in other works must be obtained from the IEEE.

# DEPTH-BASED INPAINTING FOR DISOCCLUSION FILLING

*Suryanarayana M. Muddala, Mårten Sjöström and Roger Olsson*

Dept. of Information and communication systems, Mid Sweden University, 85170, Sundsvall, Sweden

## ABSTRACT

Depth-based inpainting methods can solve disocclusion problems occurring in depth-image-based rendering. However, inpainting in this context suffers from artifacts along foreground objects due to foreground pixels in the patch matching. In this paper, we address the disocclusion problem by a refined depth-based inpainting method. The novelty is in classifying the foreground and background by using available local depth information. Thereby, the foreground information is excluded from both the source region and the target patch. In the proposed inpainting method, the local depth constraints imply inpainting only the background data and preserving the foreground object boundaries. The results from the proposed method are compared with those from the state-of-the-art inpainting methods. The experimental results demonstrate improved objective quality and a better visual quality along the object boundaries.

*Index Terms* — View synthesis, depth-image based rendering, image inpainting, disocclusions

## 1. INTRODUCTION

Three Dimensional (3D) video has become more popular with autostereoscopic display technology and free viewpoint television (FTV). Autostereoscopic displays provide depth impression without additional glasses, whereas FTV offers the user to freely select viewing perspective in the scene. To provide such an experience, these displays require multiple views. As it is inefficient to transmit multiple views due to bandwidth limitations, depth-image-based rendering (DIBR) method is one way to produce required number of views using video-plus-depth (V+D) and multiview-video-plus-depth (MVD) data formats [1].

The DIBR method requires the texture and depth of a specific view point to generate the desired view. The input to the DIBR we call the original view whereas the output is called the virtual view. The virtual view consists of warped texture and warped depth. The inherent problem with DIBR is the manifestation of holes in the virtual view, lacking both texture and depth information. These holes reduce the visual quality extensively. Holes can be either disocclusions or out-of-field areas. Disocclusions are consequences of the distance between the virtual camera and the original camera, and a distinct change in the depth between adjacent pixels, normally occurring at object borders. In this context, we define foreground to be the part of the scene closer to the camera and background to be the part of the scene farther from the camera which is partially occluded by foreground in the original image. Out-of-field areas are caused by virtual camera placement and the lack of information at the image boundaries. These artifacts need to be addressed in order to create a better visual quality and depth impression. Inpainting methods offer a solution to tackle the problems with holes. Earlier proposed inpainting methods for disocclusions were producing undesirable visual artifacts along the object boundaries. A number of later solutions

suggested incorporating the available depth data to address these artifacts. Despite many efforts in the field, there is still a need for improvements giving better objective and visual quality.

Image-inpainting methods can be generally classified into two methods: partial differential equation-based diffusion methods and exemplar-based methods. Oh et al. applied a simple diffusion method to fill holes in the virtual views [2]. In general, blurring artifacts arise by applying diffusion for large missing areas. Exemplar-based inpainting methods use both structure and texture to fill the holes in an image [3]. The key steps in this method are estimating the filling order and finding a best match between the target patch at the boundary of the hole and the source patch in the source region. The filling order is determined by the structure (data term) and the amount of reliable information (confidence term). The best source patch is found by minimizing the sum of square differences (SSD). The source region in the inpainting process is the region in the image where the best source patch is found to fill the missing regions. Although exemplar-based methods have potential to produce the structure details, they are inefficient to fill disocclusions since they propagate foreground into the disocclusion due to the lack of knowledge about depth.

Daribo et al. [4] extended the work in [3] by using depth information in the filling priority computation and patch matching step. Gautier et al. [5] also extended the exemplar-based method by restricting the filling priority to one direction. Although these methods are supported by depth information, inpainted disocclusions exhibit noticeable artifacts at the objects border. Jantet et al. [6] filled the warped depth map and then used the filled depth map as an input to the method proposed in [4]. However, the inpainted view still contained artifacts around foreground objects. The depth-based inpainting methods [7], [8] and [9] classified the source region into foreground and background regions during the patch matching process in order to fill the disocclusions from the background side. However, these methods show artifacts due to imperfections in source region classification and including details belonging to the foreground in the patch matching process; the disocclusion should only be filled with background information in DIBR disocclusions. Moreover, the methods [4], [5] and [7] rely on depth information at the virtual view position that is only available in reference test data, not in the original view.

In this paper, we propose an inpainting method for filling disoccluded regions in the virtual views that excludes foreground information in the patch matching procedure. The proposed method relies on the depth-based inpainting method in [9], but employs an enhanced source region classification and a patch matching process that utilize the available warped and original depth information. In contrast to methods [8] and [9], we have not relied on a fixed threshold in the source region classification; instead, we have considered the distribution of depth in the target patch. Furthermore, the method excludes the foreground pixels in the patch matching process.

The paper is organized as follows: Section 2 explains the fundamentals of disocclusion problems and the artifacts associated




Figure 1. Illustration of warped data (a), (b) and inpainting artifacts (c), (d).

with depth-based inpainting. This section also briefly presents the depth-included curvature inpainting method, which we use as the basis for our method. The proposed method is presented in Section 3. The experimental set-up and the evaluation criteria are described in Section 4. Results and analysis follow in Section 5 and finally Section 6 concludes this paper.

## 2. BACKGROUND

In general, a disocclusion due to DIBR is produced between a foreground object and its adjacent background; and the disoccluded area is part of the background. Disocclusions and out-of-field areas are illustrated in Figure 1(a) and (b). *Jaggedness* and *background-leaking* are the artifacts associated with conventional inpainting methods [8] and [9]. *Jaggedness* is the inconsistent texture along the foreground object (see Figure 1(c)). It is caused by a restriction of the path matching to a source region belonging to background when a target patch contains both foreground and background pixels. *Background-leaking* is a propagation of background data into different depth ranges (see Figure 1(d)). It appears when the patch matching is constrained to one depth level while the target patch contains background data but with different depth values.

The basis for our newly proposed method [9] is iterative, and the basic steps for one iteration are: (1) Identify the background boundary of a disocclusion, (2) determine the filling order on the background boundary and select a patch with the maximum priority, (3) fill holes in the target patch with an average of N-best patches from the source region, and (4) update the confidence term and the source region. Both texture and depth are filled in parallel.

## 3. PROPOSED METHOD

In this paper we propose a method to reduce the artifacts and propagate required background information by excluding the foreground pixels in the patch matching process. The method must therefore identify the foreground pixels in the target patch, which is carried out by finding the appropriate thresholds related to depth. The method can be described in two steps: (1) Foreground identification at patch level and (2) depth threshold for background selection. The steps are related to the patch matching and filling process (see Figure 2). Step (1) consists of foreground identification. In general, the depth values on the borders of a disocclusion




Figure 2. Block diagram of the [9] inpainting method; the highlighted block indicates the proposed method enhancements.




Figure 3. Exemplar Illustration of pixel locations in the warped and original depth.

are similar to depth values at depth discontinuity in the original view (see Figure 3). Therefore, we use the depth distribution of patches on both sides of a hole in the warped depth image, and the patch at the depth discontinuity in the original depth image. In Step (2), a depth threshold is selected by using the data from Step (1). The threshold is used to classify the source region and the target patch so that only background information is used in the filling of the disocclusion.

### 3.1. Foreground identification at patch level

The depth values from warped and original depth are compared in order to identify the foreground values. Assume that  $Z_p$  is a target depth patch centered at pixel  $p$  in the warped depth map, and  $f$  is a foreground pixel that relates to  $p$  by their original corresponding pixels  $f_o$  and  $p_o$  being neighbours. Depth patches centered at foreground pixels in original and warped depth are named  $Z_{f_o}$  and  $Z_f$  respectively (see Figure 4(a) and (b)). The known depth values of  $Z_p$  and  $Z_f$  are combined in to a set  $A = Z_f \cup Z_p$ , which is later used for determining depth distribution.

There are mainly two cases where holes can appear: disocclusions occurring (I) between foreground and background, and (II) between two different parts of foreground. An example of case (I) is shown in Figure 1(a) between the woman's leg and the floor, and case (II) is exemplified in Figure 1(a) between the woman's head and hand. In case (I), both the target and foreground patch share common information when the target patch  $Z_p$  positioned on a foreground object. Patches that contain foreground pixels are identified by comparing  $Z_p$  and  $Z_f$  according to (1). In case (II), the foreground patch  $Z_f$  has no relation with the target patch  $Z_p$ , because the depth values of  $Z_f$  differs from those in target patch, yet does not correspond to background. Foreground pixels are identified by comparing the depth values in  $A$  and the depth values in  $Z_{f_o}$  according to (1). We measure the similarity of the depth distributions, using the difference of local maxima positions in the histograms of  $A$  and  $Z_{f_o}$  (see Figure 4(c) and (d)). As this measure works only when the numbers of local maxima are equal in the two histograms, we include the difference in the standard deviation, as a measure for the spread in a given data. As a result, patches that contain foreground pixels are identified using the depth patch average and the histogram data.

$$Z_p \in \begin{cases} F_1 & \text{if } \overline{Z_f} \leq \max(Z_p); \\ F_2 & \text{if } (T_1 \neq T_2) \wedge (\Delta\sigma > \xi), \end{cases} \quad (1)$$

where  $F_1$  and  $F_2$  are foreground labels for two different depth ranges in a target patch in case (I) and case (II) respectively.  $\overline{Z_f}$  is




Figure 4. The patch positions in the warped and original depth (a), (b) and histograms of the patches (c), (d).

the average value of the depth patch  $Z_f$ .  $T_1$  and  $T_2$  denote the difference of local maxima positions in the histograms of  $A$  and  $Z_{f_o}$  respectively.  $\Delta\sigma$  is the difference between the standard deviation of  $A$  and  $Z_{f_o}$ .  $\xi$  is the threshold to identify the foreground objects pixels, not related to the disocclusion, calculated by computing the average of  $\Delta\sigma$  along one side of the disocclusions' boundary with respect to the virtual view.

### 3.2. Depth-threshold for background selection

After the foreground pixels in the target patch are identified, a depth threshold is selected using local depth values in order to segment the source region into foreground and background. The depth threshold has two different values depending on the depth values in the target patch. If the target patch contains foreground pixels, an average of depth values in the target patch distinguishes the foreground and background. In this case, the average of the local target patch is used as the threshold value for preventing foreground. In the case that the target patch contains background information and source region has gradient in depth, the depth values in the target patch are not sufficient to find the required background data. We solve this issue by selecting the threshold between the two local maxima in the histogram of the original depth patch, noting that the original depth patch at depth discontinuities contains at least two local maxima. In general, depth might have a number of layers, so the threshold is selected by taking the average of the depth values at the last two local maxima positions in the original depth patch histogram, as the last local maxima corresponds to the foreground. The depth threshold  $\mu$  is defined as:

$$\mu = \begin{cases} \overline{Z_p} & \text{if } Z_p \in F_1 \cup F_2; \\ \frac{d_1 + d_2}{2} & \text{otherwise,} \end{cases} \quad (2)$$

where  $\overline{Z_p}$  is average of known pixels in the target patch and  $d_1$  and  $d_2$  are the last two local maxima depth values in the histogram of  $Z_{f_o}$ . After the best-source patches are found, only background pixels are copied to the missing region in the target patch. It is important to classify the holes as disocclusion areas and out-of-boundary areas since the source region classification requires more computations. Therefore, disocclusions are inpainted with depth-based source region selection whereas out-of-field areas are filled without source region classification.

## 4. TEST ARRANGEMENT AND EVALUATION CRITERIA

The result of the proposed method is assessed using objective measurements and visual inspection. Mean structural similarity index (MSSIM) is used to assess the objective performance, given its good correspondence with visual quality. A set of 10 frames from three sequences “Ballet”, “Breakdancers” [10] and “Lovebird1” [11] is selected for quality evaluation, as this covers a range of depth distribution scenarios. The “Ballet” sequence has large depth discontinuities, wide camera baseline and several layers in the depth data, that results in larger disocclusion areas and objects at different depth levels. The “Breakdancers” has a large number of objects with gradual depth transitions. The “Lovebird1” has complex texture and structured background with sharp depth discontinuities.

The V+D input data is to produce extrapolated virtual views by DIBR. The virtual views are compared to the ground truth depth and texture at the virtual view point, which is available due selecting the virtual view point to coincide with a camera view. In the first two sequences, a virtual view positioned at camera view 4 is rendered from camera view 5 and in the third sequence the virtual view is selected to coincide with camera view 4 when rendered from camera view 6. The inpainting parameter settings are given in [9]. For comparison purposes we include results from the methods given in [3], [4], [5] and [9].

## 5. RESULTS AND ANALYSIS

We limit our study to the extrapolation scenario only, i.e., rendering the virtual view using a single V+D and we also limit the scope to images (not sequences). Results for the average MSSIM of the proposed and reference inpainting methods are listed in Table 1. The results show that the proposed method performs slightly better than the reference methods for the “Ballet” and “Breakdancers” sequences and gives similar results for the “Lovebird1” sequence. The objective improvements are in decimal precision because the disocclusions occupy only 3 – 15% of the total image. Further, we have also presented the MSSIM for the image subset (see Figure 5) in Table 2. The results show clear improvements compared to the reference methods.

In addition to the objective results, inpainted images are visually compared (examples are shown in Figure 5). The visual comparison consistently demonstrates the superior performance of the proposed method compared to the reference methods. The proposed method especially outperforms the other methods when considering the quality at the foreground object boundaries and so produces visually pleasing results (see Figure 5(e)), whereas reference methods show artifacts around the object boundaries (see Figure 5(b) to (d)).

## 6. CONCLUSIONS

In this paper, we have proposed a depth-based inpainting method for filling holes in synthesized views by focusing on source region classification. A novel technique has been presented for separating the foreground and background locally using original and warped depth in order to fill the holes with background information according to the depth range. The proposed method excludes the foreground pixels from both the source region and the target patch during the patch matching process and thus boundary artifacts are removed. Experimental results have demonstrated improved objective quality and better visual quality. Since problems in the




Figure 5. Results for visual comparison of “Ballet” (frame1) in the first two rows and “Breakdancers” (frame4) in the third row; Column (a) warped views with disocclusions; Column (b) Method [4]; Column (c) Method [5]; Column (d) Method [9]; Column (e) Proposed method.

Table 1. Average MSSIM for the whole image.

Sequence	[4]	[5]	[9]	Proposed
<b>Ballet</b>	0.860	0.865	0.874	<b>0.875</b>
<b>Breakdancers</b>	0.822	0.827	0.828	<b>0.829</b>
<b>Love bird1</b>	0.857	0.860	<b>0.861</b>	<b>0.861</b>

Table 2. MSSIM for the image subset shown in Figure 5(b) to (e)

Image	[4]	[5]	[9]	Proposed
<b>Ballet</b>	0.709	0.807	0.814	<b>0.840</b>
<b>Ballet</b>	0.826	0.848	0.846	<b>0.861</b>
<b>Breakdancers</b>	0.769	0.774	0.829	<b>0.841</b>

classification can lead to artifacts in the inpainted images, in the future work we will focus on robust methods for source region classification at presence of several depth layers.

## 7. ACKNOWLEDGEMENT

This work has been supported by grant 00156702 of the EU European Regional Development Fund, Mellersta Norrland, Sweden, and by grant 00155148 of Länsstyrelsen Västernorrland, Sweden.

## 8. REFERENCES

- [1] C. Fehn, “Depth-image-based rendering (DIBR), compression, and transmission for a new approach on 3D-TV,” *Proc. SPIE Stereoscopic Displays and Virtual Reality Systems XI*, pp. 93–104, Jan. 2004.
- [2] K. J. Oh, S. Yea, and Y. S. Ho, “Hole filling method using depth based inpainting for view synthesis in free viewpoint television and 3-d video,” in *PCS*, 2009, pp. 1–4.
- [3] A. Criminisi, P. Pérez, and K. Toyama, “Region filling and object removal by exemplar-based image inpainting,” *IEEE Transactions on Image Processing*, vol. 13, pp. 1200–1212, 2004.
- [4] I. Daribo and H. Saito, “A novel inpainting-based layered depth video for 3dtv,” *TBC*, vol. 57, no. 2, pp. 533–541, 2011.
- [5] J. Gautier, O. L. Meur, and C. Guillemot, “Depth-based image completion for view synthesis,” in *3DTV conference*, 2011, pp. 1–4.
- [6] V. Jantet, C. Guillemot, and L. Morin, “Joint projection filling method for occlusion handling in depth-image-based rendering,” in *3D Research*, 2011.
- [7] L. Ma, L. Do, and P. H. N. de With, “Depth-guided inpainting algorithm for free-viewpoint video,” in *ICIP*, 2012, pp. 1721–1724.
- [8] I. Ahn and C. Kim, “Depth-based disocclusion filling for virtual view synthesis,” in *ICME*, 2012, pp. 109–114.
- [9] S. M. Muddala, R. Olsson, and M. Sjöström, “Depth-included curvature inpainting for disocclusion filling in view synthesis,” *International Journal On Advances in Telecommunications*, vol. 6, no. 3 & 4, pp. 132–142, 2013.
- [10] C. L. Zitnick, S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski, “High-quality video view interpolation using a layered representation,” *ACM Trans. Graph.*, vol. 23, no. 3, pp. 600–608, Aug. 2004.
- [11] G. M. Um, G. Bang, N. Hur, J. Kim, and Y. S. Ho, “3d video test material of outdoor scene,” ISO/IEC JTC1/SC29/WG11/M15371, April 2008.