Back to Butterworth - a Fourier Basis for 3D Surface Relief Hole Filling within RGB-D Imagery

Amir Atapour-Abarghouei, Gregoire Payen de La Garanderie, and Toby P. Breckon

School of Engineering and Computing Sciences, Durham University, UK

Abstract—**We address the problem of hole filling in RGB-D (color and depth) images, obtained from either active or stereo based sensing, for the purposes of object removal and missing depth estimation. This is performed independently on the low frequency depth information (surface shape) and the high frequency depth detail (relief) by way of a Fourier space transform and classical Butterworth high/low pass filtering. The high frequency detail is then filled using a texture synthesis method, whilst the low frequency shape information is inpainted using structural inpainting. Here, a classical non-parametric sampling approach is extended, using the concept of query expansion, to perform high frequency depth synthesis with the final output then recombined in Fourier space. In order to improve the overall depth relief (D) and edge detail accuracy, color information (RGB) is also used to constrain the sampling process within high frequency component completion. Experimental results demonstrate the efficacy of the proposed method outperforming prior work for generalized depth filling in the presence of high frequency surface relief detail.**

I. INTRODUCTION

The main contribution of this paper is to present a method to complete an RGB-D image after a selected object or region is removed by combining texture synthesis (which mostly deals with simplistic texture [1–3]) and image inpainting (which works best with underlying image structure [4, 5]). We can similarly apply such an approach to generalized depth image completion for the task of filling the incomplete depth information resulting from sensing deficiencies [6, 7] or dynamic object removal [8].

Many texture synthesis methods follow [1], a non-parametric Markov chain synthesis algorithm that uses exhaustive nearest neighbor searching. However, one of its major drawbacks is that it can expand undesirable local features and degenerate texture (i.e. grow "garbage", [1]) or generate verbatim samples of the original image (see example in Figure 2b). This limits its general applicability to large area hole filling and the infilling of underlying scene structure. In this paper, we propose using a technique inspired by the query expansion methods in image retrieval [9, 10] to avoid this phenomenon within the context of synthesizing depth relief detail.

Furthermore, a number of approaches have been proposed to address the problem of naturally-occurring holes in depth images [11, 12]. However, work on object removal and hole filling within depth images in the presence of depth discontinuities and surface relief is still in its infancy [8, 13–15].

With an increasing demand for hole-free 3D scene capture [8], the approach presented here can be used not only to consistently remove a specified object from a depth image (see example, Figure 5), but also to fill the already existing holes, resulting from limitations in 3D scene capture [6, 7] (see example, Figure 4).

Our primary contribution is to fill the holes in scene depth by first decomposing the image using the Discrete Fourier Transform and applying the Butterworth filter [16] in the frequency domain to separate the high frequency information (relief) from the low frequency information (continuous surface shape). After filling the holes in the high frequency image via our improved texture synthesis method and inpainting the low frequency image using the method presented in [17], the two results are subsequently recombined in Fourier space to generate the final depth-filled output.

The results demonstrate that the proposed method is effective even upon depth surfaces oblique to the camera and does not introduce any further artifacts into the depth output.

II. PRIOR WORK

Despite significant prior work in color (RGB) image completion [1, 5, 17–21], very little can be found in the literature on object removal and hole-filling in depth (D) images [8, 14].

Liang et al. [14] propose a color and depth inpainting method via a segmentation-based approach in stereo images using the constraint that the missing information in one stereo image may still be visible in the other to fill in both color and depth via depth-assisted texture synthesis. However, in cases where stereo or multi-camera views are not available, yet missing depth data is abundant, other approaches have to be used to fill the naturally occurring holes in depth images.

Within active sensing, there are three approaches commonly used to deal with holes [22]. Firstly, the use of temporal and motion information to fill holes in depth maps from adjacent frames [12, 23]. Secondly, the use of neighboring pixel values and/or filtering techniques to interpolate the missing data [11, 24], and finally, a hybrid of these two approaches in an attempt to fill holes using spatio-temporal information in depth images [22, 25]. However, these approaches commonly fail at object boundary contours, produce poor surface relief and suffer from degenerate artifacts.

By contrast, our approach fills depth holes, whether they are naturally-occurring or dynamic object originated, by dealing with the underlying smooth surface continuation and fine relief/contour separately, resulting in an accurate and plausible depth image. This contribution is enabled both by the work in texture synthesis [1] and structure-based inpainting [17].
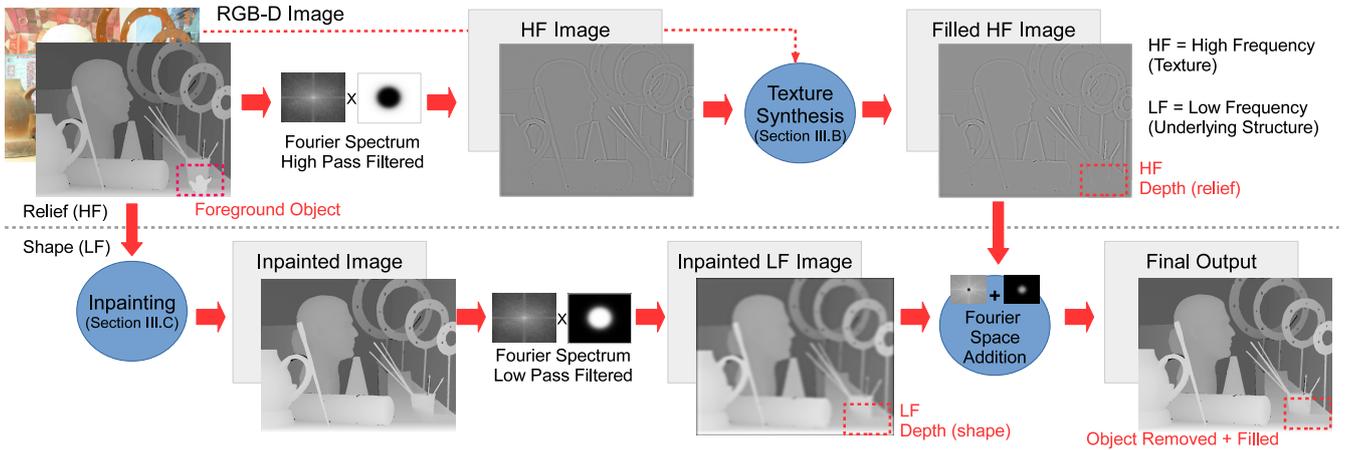
Fig. 1: A Fourier basis for the completion of scene depth via Butterworth low/high pass filtering.

In most structure-based inpainting techniques, the isophotes (lines where the intensity value is the same) are continued smoothly into the target area that is to be inpainted. Chan and Shen [19] proposed the Total Variational (TV) Model where isophotes are propagated via a Euler-Lagrange equation with diffusion. The TV Model was improved in [26] to inpaint larger regions driving diffusion along the direction of the isophotes. In [20], Oliveira et al. uses a filter convolved over the target region to diffuse the information into the target region. Telea [21] begins the inpainting from the region boundary to make sure the pixels near the known pixels are inpainted first via fast marching. Arias et al. [17] provides a variational framework for non-local inpainting which combines both diffusion methods and exemplar-based filling.

Texture synthesis methods, however, follow a different outline, where texture is generated by copying values from the known regions of the image [1, 27–29]. Although most texture synthesis methods were originally proposed to expand a smaller known region synthetically, a few can also be used in an infilling context [1].

The seminal work of Efros and Leung [1] inspired many other works [27, 30] and is used here as the basis of how a simplistic pixel-wise sampling method, with inherent stability issues, can be enhanced with reference to query expansion from the object retrieval literature [9, 10].

Other texture synthesis methods address stability by alternatively sampling and copying whole patches (texel regions) generating faster results (e.g. [28]). The notable work in [29] makes further use of adaptive patches to fill a lattice and then calculates the mismatch errors on the patch borders with overlapping parts resynthesized. Nevertheless, most of these stability enhancements into patch-wise synthesis are ill-suited for the highly constrained synthesis problem of hole filling within a randomized hole boundary in the image.

III. PROPOSED METHOD

Our approach uses three key elements in order to complete a depth image (Section III-A): separation of depth relief detail from the underlying depth shape (Section III-A), extension of texture synthesis via the integration of standard query expansion for the completion of that relief detail (Section III-B), and a structural inpainting technique used to fill the underlying continuous surfaces and structures (Section III-C).

A. Completing Depth Images

Our approach consists of four simple stages (Figure 1):

1) Separating the depth image into two images; one with only high spatial frequency information (including edges and surface relief), and one with only low spatial frequency information such as continuous surfaces shape and underlying depth gradients.
2) Filling the holes in the high frequency (HF) image via an extended synthesis method (Section III-B).
3) Using structural inpainting [17] to propagate the underlying shape structures in the low frequency (LF) image (Section III-C).
4) Final recombination of the low frequency shape information obtained from structural inpainting and the high frequency detail from the extended texture synthesis in Fourier space.

High and low frequency components are isolated via the Discrete Fourier Transform [31] with the use of a piecewise continuous circularly symmetric filter to avoid ringing artifacts. Here, we use the classical Butterworth high pass filter [16] as follows:

$$B(x,y) = \frac{1}{1 + (\frac{k}{\sqrt{x^2+y^2}})^{2n}} \quad (1)$$

and the corresponding low pass Butterworth filter defined as follows:

$$B(x,y) = \frac{1}{1 + (\frac{\sqrt{x^2+y^2}}{k})^{2n}} \quad (2)$$

where x and y are pixel positions in the Fourier image, k is the cut-off frequency, and n is the order.

Figure 1 (upper) illustrates how the original image is thus transformed into the frequency domain where the high-pass Butterworth filter is applied (Eqn. 1). The resulting high
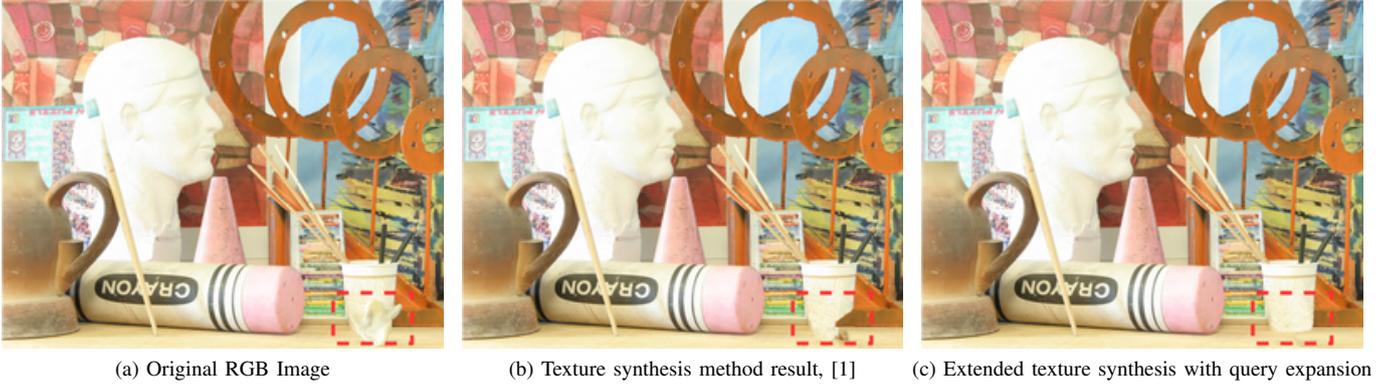
(a) Original RGB Image   (b) Texture synthesis method result, [1]   (c) Extended texture synthesis with query expansion

Fig. 2: Results of the texture synthesis method in [1] with and without query expansion.

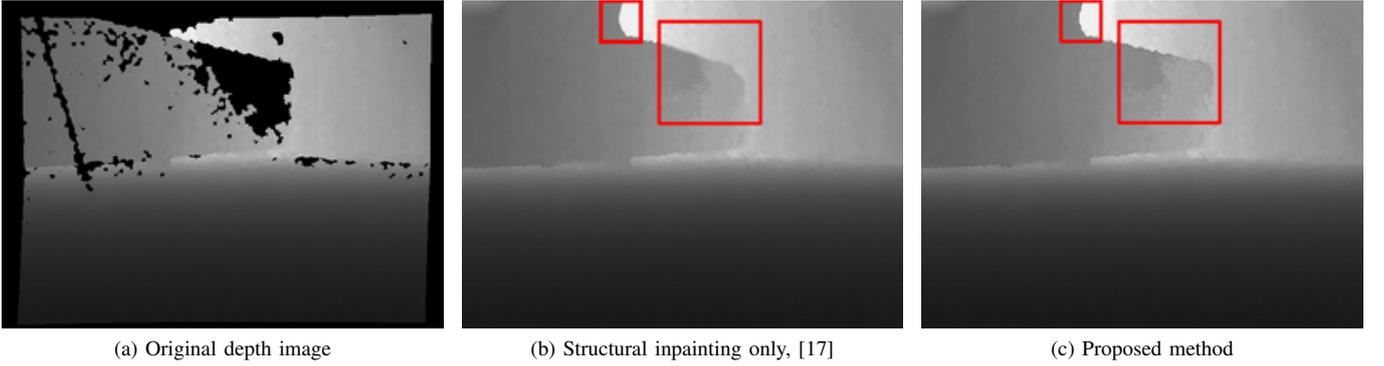(a) Original depth image   (b) Structural inpainting only, [17]   (c) Proposed method

Fig. 3: Effectiveness of the proposed method on depth maps.

frequency depth content (relief) is then completed via our extended texture synthesis approach (Section III-B), whilst the low frequency surface shape is completed via structural inpainting akin to [17] (Section III-C).

Although it is possible to structurally inpaint the low frequency components after the decomposition of the image, experiments demonstrate superior results if the low-pass filter is applied post structural inpainting (Figure 1, lower).

In either case, both the synthesis and inpainting approaches could be equally replaced with an alternate constrained texture synthesis and another structural inpainting method that does not use fronto-parallel translational patches respectively.

After the two images have been filled, they are added back together in the Fourier domain to generate the final image with the missing depth regions plausibly filled. The stages of this process are illustrated in Figure 1 for the removal of a foreground object (inset, top left) from which the resulting depth hole is subsequently filled across both the high and low frequency spatial components of the scene.

B. Texture Synthesis via Query Expansion

In this classic method [1], texture is modeled as a Markov Random Field. To generate texture, the neighborhood of a pixel is considered to be a square window around the pixel. For every unknown pixel to be filled a $w \times w$ pixel neighborhood N_q is created which is then used to query the sample (known) region of the image to find all sample neighborhoods, N_s,

that are similar. The similarity of the two neighborhoods is determined by calculating the Gaussian weighted sum of squared difference (SSD) between the two (Eqn. 3 and 4):

$$SSD(N_q, N_s) = \sum_{1 \leq i,j \leq w} d[p_{i,j}(N_q) - p_{i,j}(N_s)] \times G(i,j) \quad (3)$$

$$G(i,j) = \frac{1}{2\pi\sigma^2} e^{-\frac{(i-i_0)^2 + (j-j_0)^2}{2\sigma^2}} \quad (4)$$

where $d[p_{i,j}(N_q) - p_{i,j}(N_s)]$ is the squared Euclidean difference between the corresponding pixels in neighborhoods N_q and N_s, w is the neighborhood size, G is a two-dimensional Gaussian kernel, σ is the standard deviation, and $i_0 = j_0 = \frac{w-1}{2}$ marks the center of the neighborhood.

It should be noted that the SSDs can be calculated for a one-channel image (as in a single depth image, Figures 3 and 4), a three-channel image (as in an RGB image, Figure 2), or a four-channel image (as in an RGB-D image, Figures 5 and 6).

All windows N_s that meet the following criterion are placed in a list, from which one is randomly selected to fill the pixel:

$$SSD(N_q, N_s) < (1 + \epsilon)d_{min} \quad (5)$$

where $\epsilon = 0.1$, and d_{min} is the minimum of the SSD between N_s and N_q. It is here where this approach can fall into a narrow subspace of the overall query space and produce substandard results. To reduce this likelihood, we propose

using a technique from the object retrieval literature [9] to improve comparative query results. The solution is to simply expand the original query, issuing new queries, and hence avoiding a narrowing of the search space towards degenerate sampling of the query space.

In our proposed method, after generating the set of neighborhoods that satisfy the condition in Eqn. (5), instead of randomly choosing one instantaneously, new queries are performed based on each of the neighborhoods already in the list. These results of the search are added to the list once again. This can be repeated several times; however, experiments have shown better results are achieved when the search space is expanded once. By expanding the search for each pixel, even if the neighboring pixels were synthesized incorrectly, the overall probability of degenerate sampling due to a few bad pixels is greatly reduced. Exemplar comparative results, operating solely on RGB color at this stage, are shown in Figure 2.

Within our framework, this extended texture synthesis is performed on RGB-D imagery with the SSD (Eqn. 3) extended to a four channel $\{R,B,G,D\}$ formulation to operate on both color (RGB) and high frequency depth (D) pixel values (Figure 1, upper). In general, the concept can be applied to RGB and/or depth information in isolation (Figures 2 and 3).

Furthermore, this expansion of search space can be applied to any other approach that issues queries in a similar fashion, be it for various types of patch-based or pixel-based texture synthesis techniques or inpainting methods [27, 28, 32].

C. Structural Inpainting

The variational framework for inpainting proposed in [17] is mainly inspired by the energy function used by Gilboa and Osher in [33]. In [17], however, a second term is added to the energy equation, which improves the estimation of the image function.

Let u be the image function $u : \Omega \to \mathbb{R}$, where Ω denotes the image domain. $O \subset \Omega$ is the target region (hole), and $O^c := \Omega \backslash O$. $\tilde{O}$ is taken as the extended target region, which consists of the centers of the patches that intersect the target region. As a result, the patches in $\tilde{O}^c$ are completely outside the target region O.

The energy function proposed in [17] measures the consistency between the patches in $\tilde{O}$ and $\tilde{O}^c$ for a similarity weight function $w : \tilde{O} \times \tilde{O}^c \to \mathbb{R}$ (Eqn. 6).

$$E(u, w) = \frac{1}{h}\tilde{F}_w(u) - \int_{\tilde{O}} H_w(x)dx \qquad (6)$$

subject to: $\qquad \int_{O^c} w(x, \hat{x})d\hat{x} = 1$

where h is the selectivity parameter of the Gaussian weights,

$$\tilde{F}_w(u) := \int_{\tilde{O}} \int_{\tilde{O}^c} w(x, \hat{x})\epsilon(p_u(x) - p_{\hat{u}}(\hat{x}))d\hat{x}dx \qquad (7)$$

and

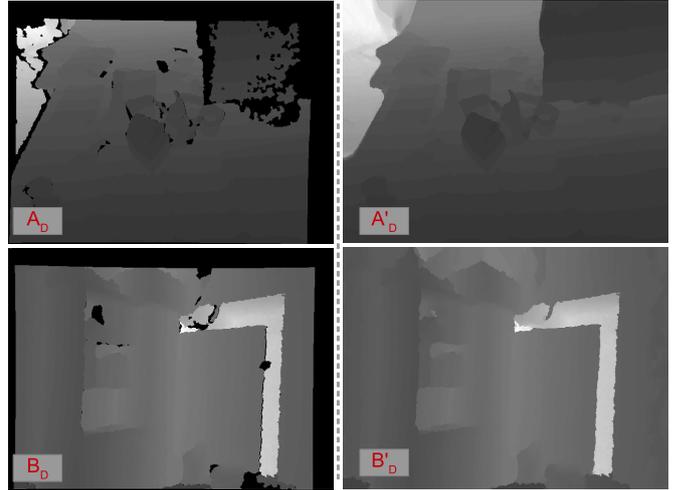$$H_w(x) := -\int_{\tilde{O}^c} w(x, \hat{x}) \log w(x, \hat{x})d\hat{x} \qquad (8)$$

Fig. 4: The proposed method used to fill holes in depth maps.

where ϵ is an error function for image patches, and $H_w(x)$ is the entropy of the probability function $w(x, .)$. For further details on the error function ϵ, please see [17]. Minimizing Eqn. 7 will result in patches $p_u(x)$ and $p_{\hat{u}}(\hat{x})$ being similar when the weight $w(x, \hat{x})$ is high. By taking the limit $h \to 0$, the energy is approximated as:

$$E(u, w) \simeq \int_{\tilde{O}} \int_{\tilde{O}^c} w(x, \hat{x})\epsilon(p_u(x) - p_{\hat{u}}(\hat{x}))d\hat{x}dx \qquad (9)$$

Therefore, inpainting is now formulated as an optimization problem:

$$(u^*, w^*) = \underset{u,w}{\arg\min}\, E(u, w) \qquad (10)$$

subject to: $\qquad \int_{\tilde{O}^c} w(x, \hat{x})d\hat{x} = 1 \quad \forall x \in \tilde{O}$

where E is the energy defined in Eqn. 6. The minimization of E is described in [17]. The inpainting is then applied on a Gaussian image pyramid [34]. The results at each scale are upsampled, starting from the coarse scale and ending with the finest scale, and fed as initial values to the next scale.

Let there be S scales with A_0 being the size of the image at the finest scale and A_{S-1} at the coarsest. The process is initialized at the coarsest scale:

$$(u^{S-1}, w^{S-1}) = \underset{u,w}{\arg\min}\, E_{a_0}(u, w) \qquad (11)$$

where a_0 is the size of the patch and E_{a_0} the corresponding energy. The subsequent upsampling from one scale to the next is done as per [35].

Here, this method [17] provides a robust approach for the structural inpainting of low frequency shape information within the scene (Figure 1, lower). However, in general, our overall depth filling proposal remains agnostic to the exact inpainting routine in use for this task, and many other inpainting techniques can yield similarly satisfactory results.

IV. EXPERIMENTAL RESULTS

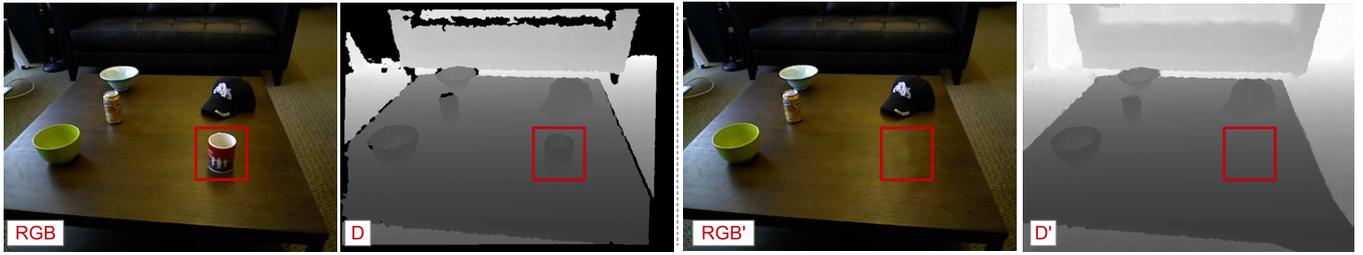Results showing the efficacy of our proposed approach across both interior and exterior scenes are presented in

Fig. 5: Object removed from an RGB-D image captured via a structured light depth sensor (Microsoft Kinect).

Figures 2 - 6. These include depth relief completion (Figures 3 and 4), object removal (Figures 5 and 2), in addition to combined depth filling and object removal in complex exterior environments within the context of vehicle sensing (Figure 6).

As indicated in Figure 3, the proposed method for depth infilling can be used to fill the existing holes in depth maps. Figure 3b depicts how the structural inpainting method [17] used alone can create additional artifacts in the image around the edges and is incapable of producing accurate texture, whereas the proposed method (Figure 3c) can correctly identify texture and edges more accurately and produce more satisfactory results with sharper and crisper depth relief.

The proposed method has also been applied to single depth images with missing data. As seen in Figure 4, our approach is capable of filling the holes in the images obtained through active 3D sensors without being affected by any noise or depth blurring as is usually the case in most hole-filling strategies or inpainting techniques applied to depth maps [11, 17, 24]. Moreover, even when there are holes in or around the boundaries of narrow objects, our method functions effectively, unlike many other depth hole-filling approaches [12, 22, 23].

Figure 5 demonstrates what results the proposed method can yield when applied to an actively sensed RGB-D image (Microsoft Kinect). An object has been removed from the both the color (RGB) and the depth (D) image, and the naturally-occurring holes in the depth image have also been filled successfully as seen in the images on the right (Figure 5, RGB' & D').

The algorithm was also applied to real-world images captured using stereo cameras for automotive applications [6, 8]. An example of a major challenge in vision-based automotive systems is being able to remove undesirable dynamic objects while trying to map a static scene through which the automotive system attempts to navigate [8]. Figure 6 demonstrates how selected objects, such as other cars or pedestrians, have been removed from RGB and disparity images and the existing holes in the disparity maps have been successfully filled.

V. CONCLUSION

We present a novel method to fill holes within RGB-D imagery, resulting from either object removal [8, 14] or limitations in 3D sensor capabilities [12, 22, 23], uniquely performed in Fourier space to facilitate independent completion of both high frequency depth detail (relief) and low frequency surface continuation (shape).

Our approach makes use of the Discrete Fourier Transform to decompose a depth map into separate high and low frequency component via Butterworth filtering. An improved texture synthesis method is then proposed to fill the high frequency depth detail, and a structural inpainting approach is used to fill the underlying structures in the image. The two are then recombined in Fourier space to create a depth image, where the holes are plausibly filled and the surface relief and edges preserved.

The majority of the current hole-filling techniques fall short when in comes to the depth filling of surface relief and edge detail. Not only do most strategies fail to estimate object boundaries and fine texture correctly [22, 23, 25], many can introduce additional artifacts and/or ringing effects into the image [11, 24]. Our approach preserves the texture and the boundaries successfully, even in case of narrow object, and can create a sharp, crisp depth output as illustrated over a range of exemplar scenarios.

REFERENCES

[1] A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," in *Proc. Int. Conf. Computer Vision*, vol. 2. IEEE, 1999, pp. 1033–1038.

[2] D. J. Heeger and J. R. Bergen, "Pyramid-based texture analysis/synthesis," in *Proc. Conf. Computer Graphics and Interactive Techniques*, 1995, pp. 229–238.

[3] E. P. Simoncelli and J. Portilla, "Texture characterization via joint statistics of wavelet coefficient magnitudes," in *Int. Proc. Conf. Image Processing*, vol. 1. IEEE, 1998, pp. 62–66.

[4] M. Bertalmio, A. L. Bertozzi, and G. Sapiro, "Navier-stokes, fluid dynamics, and image and video inpainting," in *Proc. Computer Vision and Pattern Recognition*, vol. 1. IEEE, 2001, pp. I–355.

[5] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proc. Conf. Computer Graphics and Interactive Techniques*, 2000, pp. 417–424.

[6] F. Mroz and T. Breckon, "An empirical comparison of real-time dense stereo approaches for use in the automotive environment," *Image and Video Processing*, vol. 2012, no. 13, pp. 1–19, 2012.

[7] G. Payen de La Garanderie and T. Breckon, "Improved depth recovery in consumer depth cameras via disparity space fusion within cross-spectral stereo," in *Proc. British Machine Vision Conference*, September 2014, pp. 417.1–417.12.

[8] O. Hamilton and T. Breckon, "Generalized dynamic object removal for dense stereo vision based scene mapping using synthesised optical flow," in *Proc. Int. Conf. Image Processing*. IEEE, September 2016.

[9] O. Chum, A. Mikulik, M. Perdoch, and J. Matas, "Total recall ii: Query expansion revisited," in *Proc. Computer Vision and Pattern Recognition*. IEEE, 2011, pp. 889–896.

[10] R. Arandjelović and A. Zisserman, "Three things everyone should know to improve object retrieval," in *Proc. Computer Vision and Pattern Recognition*. IEEE, 2012, pp. 2911–2918.
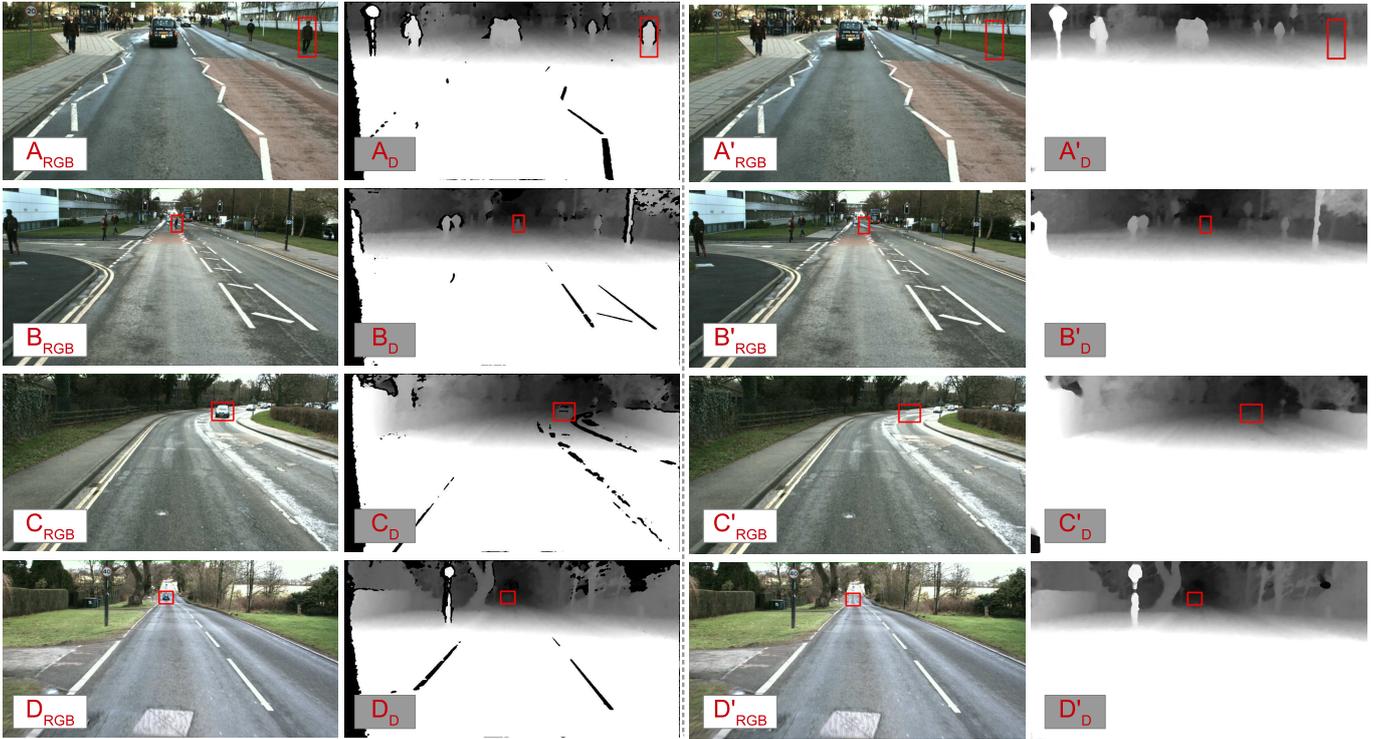
Fig. 6: Examples of the method applied to real-world images captured for automotive applications.

[11] N.-E. Yang, Y.-G. Kim, and R.-H. Park, "Depth hole filling using the depth distribution of neighboring regions of depth holes in the kinect sensor," in *Int. Conf. Signal Processing, Communication and Computing*. IEEE, 2012, pp. 658–661.

[12] S. Matyunin, D. Vatolin, Y. Berdnikov, and M. Smirnov, "Temporal filtering for depth maps generated by kinect depth camera," in *3DTV Conference*. IEEE, 2011, pp. 1–4.

[13] T. Breckon and R. Fisher, "3D surface relief completion via non-parametric techniques," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 30, no. 12, pp. 2249 – 2255, 2008.

[14] L. Wang, H. Jin, R. Yang, and M. Gong, "Stereoscopic inpainting: Joint color and depth completion from stereo images," in *Proc. Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.

[15] T. Breckon and R. Fisher, "A hierarchical extension to 3D non-parametric surface relief completion," *Pattern Recognition*, vol. 45, pp. 172–185, September 2012.

[16] C. Butterworth, "Filter approximation theory," *Engineer*, vol. 7, pp. 536–541, 1930.

[17] P. Arias, G. Facciolo, V. Caselles, and G. Sapiro, "A variational framework for exemplar-based image inpainting," *Int. J. Computer Vision*, vol. 93, no. 3, pp. 319–347, 2011.

[18] A. Criminisi, P. Pérez, 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.

[19] T. Chan and J. Shen, "Mathematical models for local deterministic inpaintings," Technical Report CAM TR 00-11, UCLA, Tech. Rep., 2000.

[20] M. M. O. B. B. Richard and M. Y.-S. Chang, "Fast digital image inpainting," in *Proc. Int. Conf. Visualization, Imaging and Image Processing*, 2001, pp. 106–107.

[21] A. Telea, "An image inpainting technique based on the fast marching method," *Graphics Tools*, vol. 9, no. 1, pp. 23–34, 2004.

[22] J. Wang, P. An, Y. Zuo, Z. You, and Z. Zhang, "High accuracy hole filling for kinect depth maps," in *SPIE/COS Photonics Asia*, 2014, pp. 92 732L–92 732L.

[23] Y. Berdnikov and D. Vatolin, "Real-time depth map occlusion filling and scene background restoration for projected-pattern based depth cameras," in *Graphic Conf. IETP*, 2011.

[24] D. Min, J. Lu, and M. N. Do, "Depth video enhancement based on weighted mode filtering," *IEEE Trans. Image Processing*, vol. 21, no. 3, pp. 1176–1190, 2012.

[25] M. Camplani and L. Salgado, "Efficient spatio-temporal hole filling strategy for kinect depth maps," in *IS&T/SPIE Electronic Imaging*, 2012, pp. 82 900E–82 900E.

[26] T. F. Chan and J. Shen, "Nontexture inpainting by curvature-driven diffusions," *Visual Communication and Image Representation*, vol. 12, no. 4, pp. 436–449, 2001.

[27] A. A. Efros and W. T. Freeman, "Image quilting for texture synthesis and transfer," in *Proc. Conf. Computer Graphics and Interactive Techniques*, 2001, pp. 341–346.

[28] L. Liang, C. Liu, Y.-Q. Xu, B. Guo, and H.-Y. Shum, "Real-time texture synthesis by patch-based sampling," *ACM Trans. Graphics*, vol. 20, no. 3, pp. 127–150, 2001.

[29] A. Nealen and M. Alexa, "Hybrid texture synthesis," in *Proc. 14th Eurographics Workshop on Rendering*, 2003.

[30] M. Ashikhmin, "Synthesizing natural textures," in *Symp. Interactive 3D graphics*, 2001, pp. 217–226.

[31] C. Solomon and T. Breckon, *Fundamentals of Digital Image Processing: A Practical Approach with Examples in Matlab*. Wiley-Blackwell, 2010.

[32] A. Criminisi, P. Pérez, 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.

[33] G. Gilboa and S. Osher, "Nonlocal linear image regularization and supervised segmentation," *Multiscale Modeling & Simulation*, vol. 6, no. 2, pp. 595–630, 2007.

[34] C.-W. Fang and J.-J. J. Lien, "Rapid image completion system using multiresolution patch-based directional and nondirectional approaches," *IEEE Trans. Image Processing*, vol. 18, no. 12, pp. 2769–2779, 2009.

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

[36] H. Hirschmuller and D. Scharstein, "Evaluation of cost functions for stereo matching," in *Proc. Computer Vision and Pattern Recognition*, 2007, pp. 1–8.

[37] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *Proc. Int. Conf. Intelligent Robot Systems*, 2012.

[38] K. Lai, L. Bo, and D. Fox, "Unsupervised feature learning for 3d scene labeling," in *IEEE Int. Conf. Robotics and Automation*, 2014, pp. 3050–3057.