Stereoscopic image inpainting: distinct depth maps and images inpainting

Alexandre Hervieu, Nicolas Papadakis, Aurélie Bugeau, Pau Gargallo I
Piracés, Vicent Caselles

Stereoscopic image inpainting: distinct depth maps and images inpainting

Alexandre Hervieu, Nicolas Papadakis, Aurélie Bugeau, Pau Gargallo
Barcelona Media, Avenida Diagonal, 177, 08017 Barcelona, Spain
{*alexandre.hervieu, nicolas.papadakis, aurelie.bugeau, pau.gargallo*}@barcelonamedia.org

Vicent Caselles
DTIC, Universitat Pompeu Fabra, Roc Boronat, 138, 08018 Barcelona, Spain
vicent.caselles@upf.edu

Abstract

In this paper we propose an algorithm for inpainting of stereo images. The issue is to reconstruct the holes in a pair of stereo image as if they were the projection of a 3D scene. Hence, the reconstruction of the missing information has to produce a consistent visual perception of depth. Thus, first step of the algorithm consists in the computation and inpainting of disparity maps in the given holes. The second step of the algorithm is to fill-in missing regions using the complete disparity maps in a way that avoids the creation of 3D artifacts. We present some experiments on several pairs of stereo images.

1 Introduction

Image inpainting consists in recovering the missing or corrupted parts of an image so that the reconstructed image looks natural. We develop in this paper a method for depth consistent inpainting in pairs of stereoscopic images. We assume that the holes, or regions where we want to reconstruct the missing information, correspond to the image of objects in the scene, so that they are depth consistent. Then one has to design an inpainting method that produces visually depth consistent images. Inpainting methods can be organized into two main categories: geometry and texture oriented methods.

Geometry-based methods aim at propagating the geometric structure of the image from the boundary towards the interior of the holes. Usually these methods consist in finding the minimum of an energy functional or directly solving a Partial Differential Equation [1, 7]. The methods permit to recover the non-textured and smooth information when the missing region is thin.

Texture oriented methods applied to inpainting were developed as an application of texture synthesis [3]. The unknown region Ω, called the hole, is filled-in by copying content from the known part Ω^c of the image, the complement of the hole. In the seminal paper [3] a simple yet effective non-parametric texture synthesis method based on local image patches was proposed. In [3], the texture is modeled as a MRF by assuming that the probability distribution of brightness values for one pixel given the brightness values of its spatial neighbor-

hood is independent from the rest of the image. The neighborhood Ψ_p is a square window around the pixel p and its size $s \times s$ is fixed by hand. The task is then to select the patch that gives the smallest sum of squared differences distance (SSD) and to copy the color of its central pixel. This method is at the origin of many algorithms, one of the most effective one being [2]. In this paper, two improvements to the method in [3] were made. The first one is a speed-up process. Contrary to [3], that synthesizes pixels one by one, all the pixels in the neighborhood of a point are here filled-in. The second improvement is the order in which the pixels of Ω are filled based on priority assignments. This algorithm provides interesting results for inpainting of both textured and non-textured images.

For stereoscopic image inpainting, let us mention the method proposed by Wang et al. [8]. In this paper, disparity map estimation is required before a user is asked to choose the object to remove. On the contrary, our method directly deals with the estimation of disparities in the holes to be filled-in, thus enabling inpainting of *a priori* missing regions in the stereo images while ensuring 3D coherence of inpainting between stereo images.

Overview of the algorithm

Our method for filling-in missing parts in stereo images is based on a two-step algorithm, the first step dealing with computation and inpainting of depth maps, the second one achieving stereo image inpainting.

Stereo images are first used to get depth maps. To compute disparities, the pair of stereo images is rectified. Computation and inpainting of disparities are then performed simultaneously and provide complete disparity maps for each stereo image (see Section 2).

The second step consists in filling-in the images using the complete disparity maps obtained in the first step of the algorithm (Section 3). The holes to be inpainted in each image are first partially completed using the corresponding visible part in the other image. The two stereo images are finally inpainted simultaneously based upon a symmetric extension to stereo images of the inpainting method proposed by Criminisi et al. [2].

Rectifying the images

In order to compute disparities on the stereo images, rectification of the images is performed.

First, the computation of keypoints and their matching is done on the two stereo images using the algorithm proposed by Lowe in [6]. The fundamental matrix F is then computed using a classical eight points algorithm [4]. Image rectification is performed based upon a single value decomposition (SVD) of F [4].

2 Disparity estimation and inpainting with visibility constraint

In order to obtain complete disparity maps for each stereo image, disparities are computed outside the holes and simultaneously interpolated inside. Let us consider two fronto-parallel cameras C_i, $i = \{1,2\}$ observing a scene. Their rectified (color) images are denoted by $I_i : \Theta_i \rightarrow \mathbb{R}^3$, where $\Theta_i \subseteq \mathbb{R}^2$ is the domain of I_i.

As the images are rectified (the rows of the images are the corresponding epipolar lines), corresponding pixels in the 3D scene belong to the same lines of images I_1 and I_2. The difference of column coordinates of the corresponding pixels is the disparity. The disparity map going from I_i to I_j, $i, j \in \{1,2\}$, will be denoted by $D_i : \Theta_i \mapsto \mathcal{L}_i$, where $\mathcal{L}_i \subset \mathbb{R}$ is the set of possible disparity values of image I_i. Assuming that C_1 is located on the left and C_2 on the right, the set $\mathcal{L}_2$ will contain positive values, $\mathcal{L}_1$ will contain negative ones.

2.1 Disparity estimation

We want to estimate the disparity maps $D = [D_1, D_2]$ between images I_1 and I_2 for all their pixels $p = (x,y)$. For that, we use the method in [5] based on the minimization of an energy. As the problem is solved in a discrete framework via a graph cut approach, the disparity D takes its values in a discrete set of possible disparities contained in the range $\mathcal{L}_2 = -\mathcal{L}_1 = [D_{min}, D_{max}]$. The energy functional we use is composed of a photoconsistency data term and a total variation regularization:

$$\mathcal{E}(D) := \sum_{i=1}^{2} \sum_{p \in \Theta_i} \min(|I_i(p) - I_j(x + D_i(p), y)|, \gamma)$$
$$+ \alpha \sum_{i=1}^{2} \sum_{p \in \Theta_i} |\nabla D_i(p)|,$$
(1)

where the index j in the data term being the index of the image $\neq i$, and γ is the occlusion cost and α is the regularization parameter. This algorithm also enforces the visibility constraint between the disparity maps D_1 and D_2. This constraint can be formulated as

$$\forall (x', y) \in \Theta_j, j=1,2, \text{ for which } \exists (x,y) \in \Theta_i \text{ with } i \neq j$$
$$\text{such that } x' = x + D_i(x,y), \text{ then } |D_j(x',y)| \geq |D_i(x,y)|.$$
(2)

It means that a pixel (x,y) in image i can be occluded by a pixel (x',y) in image j if and only if the disparity of (x',y) is greater than the disparity of (x,y). This visibility constraint is fundamental in the case of stereo inpainting in order to synthesize stereo images respecting the occlusions. The class of disparity maps verifying the property (2) is denoted by $\mathcal{A}_{vis}$ and the algorithm in [5] finally solves the following problem

$$\min_{D \in \mathcal{A}_{vis}} \mathcal{E}(D).$$

2.2 Disparity inpainting

In the case of stereo inpainting, the method previously described can be adapted to provide an inpainting of disparities. We denote by Ω_1 and Ω_2 the holes of images I_1 and I_2, the data term of energy (1) can be modified as

$$\sum_{i=1}^{2} \sum_{p \in \Theta_i \backslash \Omega_i} \min(|I_i(p) - I_j(x + D_i(p), y)|, \gamma).$$
(3)

As a consequence, this model does not take into account the photometry inside the holes and the process will realize a total variation based diffusion inside these areas. As the problem is still solved on the set $\mathcal{A}_{vis}$, this adaptation of [5] inpaints the disparity map inside the holes, while respecting the visibility constraint.

3 Stereo image inpainting

Our next step is to fill-in the holes Ω_1 and Ω_2 in I_1 and I_2. To this aim, each image is first partially completed using the corresponding visible parts of the other. Then, a stereo inpainting algorithm is proposed based upon an extension of a popular inpainting algorithm [2].

Filling-in the occluded regions Having computed complete disparity maps, some unknown parts of one image may be obtained from the corresponding parts of the other image. Given a pixel $p = (x,y)$ in the hole Ω_i, if its corresponding pixel $q = (x + D_i(x,y), y)$ is known (i.e. $q \in \Theta_j \backslash \Omega_j$), then we compute the color of p as $I_i(p) = I_j(q)$. Since p could be occluded in I_j, we first check that q also corresponds to p by checking that $D_j(x + D_i(x,y)) = -D_i(x,y)$.

Inpainting the stereo images We propose a extension of the algorithm of Criminisi et al. [2] to the case of stereo pairs. In the original algorithm, a priority $P(p)$ is assigned to each pixel p of the boundary $\partial \Omega$ of Ω. The priority is the product of a confidence term, which measures the amount of reliable information surrounding the pixel, and a data term that encourages linear structures to be synthesized first.

In the stereo context, two images are considered. Therefore, the filling-in process is done here by considering the contours of Ω_1 and Ω_2. We keep the same definition of priority, but compute it on both $\partial \Omega_1$ and

$\partial\Omega_2$. The patch to be filled-in is the one surrounding the pixel in $\partial\Omega_1 \cup \partial\Omega_2$ with the highest priority.

When a patch Ψ_i^p centered at pixel $p = (x, y)$ of image I_i is filled-in, we also fill the patch Ψ_j^q centered at its correspondent pixel $q = (x + D_i(x, y), y)$ in I_j. The algorithm fills-in the uncovered regions of Ψ_i^p and Ψ_j^q with the corresponding parts of patch $\Psi_i^{\tilde{p}}$ defined by

$$\tilde{p} \in \underset{\tilde{p} \in \Theta_i \setminus \Omega_i}{\operatorname{argmin}} \ d(\Psi_i^p, \Psi_i^{\tilde{p}}) + d(\Psi_j^q, \Psi_j^{\tilde{q}}), \qquad (4)$$

where $d(\Psi_1, \Psi_2)$ is the SSD of known parts of patches Ψ_1 and Ψ_2, and pixel $\tilde{q}$ is the correspondent of pixel $\tilde{p}$.

To fill the pixels of Ψ_j^q, we warp the patch $\Psi_i^{\tilde{p}}$ according to the disparities of Ψ_i^p. Several pixels in $\Psi_i^{\tilde{p}}$ may be warped to the same pixel in Ψ_j^q. In this case, we choose the one closest to the camera (highest disparity).

To avoid coping texture from foreground objects to background, we assume that holes are to be filled with pixels that are deeper. Hence, we restrict the search in equation (4) to patches $\Psi_i^{\tilde{p}}$ composed of pixels with disparities smaller or equal than the disparity of pixel p.

We follow the above procedure for all patches that have a complete corresponding patch in the other image. Once the filling-in is completed for all such patches, the remaining patches (in the borders of the images) have a pixel that has no correspondence in the other image. These patches are then filled using the algorithm in [2].

4 Experimental Results and comparison

The results reported in this section illustrate the inpainting of both disparity maps and images. For the three examples shown in figures 1, 2 and 3, the first line presents the two initial images and the masks defining the regions to be inpainted. The second line of figures 1 and 2 contains the disparity maps computed without inpainting using eq. 1 (the two images on the left) and with inpainting using eq. 3 (images on the right). Last line of all figures shows the final image inpainting results. The two leftmost images are the results obtained with the method of Wang et al. [8] for comparison; the two on the right are obtained using our novel algorithm.

The results obtained show the efficiency of the disparity map inpainting scheme. Depths are correct and the inpainting correctly diffuses the neighboring disparities into the holes. The method provides, for each of the stereo images, a 3D-coherent disparity map.

Concerning stereo image inpainting, the results obtained are satisfactory. Geometrical structures have been maintained thanks to our exemplar-based scheme. For all the examples, the 3D coherence is well kept by our algorithm, ensuring a perfect correspondence between the inpainted colors of the two stereo images.

Comparison with the method in [8] highlights the effectiveness of our method and the interest of inpaint-

ing the disparity maps first. Indeed, in this method, 3D coherence is imposed by an iterative process that compares the results obtained on both sides. Results are considered as good if the inpainted disparity and color are close enough (relying on parameters). The method does not ensure convergence. Hence, the results may be very dependent on the parameters and on the number of iterations. The comparison method creates 3D artifacts: the inpainted part from one image does not fit, regarding disparity maps, with the inpainted part of the other image. This is observed in the three experiments shown here. For example, in figure 1, a pink part appears in one image and not in the other. On the contrary, our method ensures perfect correspondences of the inpainted parts of the two images according to the disparity maps.

5 Conclusion

This paper has presented a two-step algorithm for inpainting stereo images. First, a procedure for disparity maps inpainting has been proposed and tested positively. These disparities are then used to inpaint both stereo images using a symmetric patch-based method. Our scheme ensures the 3D coherence between both inpainted images. Comparison with the existing method has shown the benefits of our algorithm.

References

[1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *ACM SIGGRAPH*, 2000.

[2] A. Criminisi, P. Pérez, and K. Toyama. Region filling and object removal by exemplar-based inpainting. *IEEE Trans. on Im. Processing*, 13(9):1200–1212, 2004.

[3] A. Efros and T. Leung. Texture synthesis by nonparametric sampling. In *IEEE ICCV*, 1999.

[4] R. Hartley and A. Zisserman. Multiple view geometry in computer vision. In *Cambridge University Press*, 2004.

[5] V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. In *ECCV*, 2002.

[6] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int. Jour. on Comp. Vis.*, 60(2):91–110, 2004.

[7] S. Masnou and J. Morel. Level-lines based disocclusion. In *IEEE ICIP*, 1998.

[8] L. Wang, H. Lin, R. Yang, and M. Gong. Stereoscopic inpainting: joint color and depth completion from stereo images. In *IEEE CVPR*, 2008.

Figure 1. Inpainting of the toy

Figure 2. Inpainting of the lanterns

Figure 3. Inpainting of the bird