

Edge adaptive and energy preserving volume upscaling for high quality volume rendering

A. Giachetti¹, J.A. Iglesias Guitián² and E. Gobbetti²

¹VIPS lab, Dipartimento di Informatica, Università di Verona, Strada Le Grazie 15, 37134 Verona, Italy

²Visual Computing Group, CRS4, POLARIS, Edificio 1, 09010 PULA (CA - Italy)

Abstract

We describe an edge-directed optimization-based method for volumetric data supersampling. The method is based on voxel splitting and iterative refinement performed with a greedy optimization driven by the smoothing of second order gray level derivatives and the assumption that the average gray level in the original voxels region cannot change. Due to these assumptions, the method, which is the 3D extension of a recently proposed technique, is particularly suitable for upscaling medical imaging data creating physically reasonable voxel values and overcoming the so-called partial volume effect.

The good quality of the results obtained is demonstrated through experimental tests. Furthermore, we show how offline 3D upscaling of volumes can be coupled with recent techniques to perform high quality volume rendering of large datasets, obtaining a better inspection of medical volumetric data.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement/Sharpening and deblurring—

1. Introduction

A lot of research effort has been recently spent on digital image or video upscaling, and several different methods have been recently proposed for these purposes, like edge based [LO01], optimization based [Fat07, SXS08, GA08, TMSV09] or example based methods [FJP02, KTN*07, KK08]. These techniques, however, have been rarely applied to voxelized volumes captured by diagnostic modalities.

The upscaling of medical volumes is usually performed just in the Z direction to overcome the problem of relevant slice spacings causing highly non-isotropic voxels. Typical solutions applied consist of interleaving new "interpolated" slices between the original ones, computed usually with methods based on preserving continuity of segmented structures (e.g. Shape Based Interpolation [HZB92]) or through deformable registration of consecutive slices (e.g. Registration-based interpolation [PSR*04]).

In this way it is possible to reduce the so-called "partial volume" effect that makes the physical value measured at each voxel location not tissue-specific, but corresponding to

an averaged value due to the presence of different materials inside the volume.

In this paper we propose a volume upscaling method that we implemented not only for interleaving new slices, but also for isotropic enlargement. The method is the 3D extension of a recent algorithm for single image superresolution [Gia10].

It is based on voxel subdivision (so that integer zooming factors along the three dimensions can be obtained), the assumption that the energy or density measured inside each voxel is equal to the sum of those that would have been acquired in the region by the simulated high resolution sensor, and to the assumption of local smoothness of the second order derivatives of the image brightness.

The physical constraint of assuming the average spatial properties of the material unchanged in the original volume block seem the best choice to partially overcome the partial volume effect and obtain more accurate results in classification and segmentation tasks. Furthermore, this kind of processing can be useful in order to approximate high resolution detail in volume rendering, improving visual analy-

sis of the imaged anatomy. Recent multi-resolution methods are, in fact, able to handle huge voxelized volumes allowing their interactive visualization. This means that through an off-line volume upscaling and the use of such techniques it is possible to relevantly improve the quality of the 3D visual analysis of medical data. The paper is organized as follows: Section 2 presents the proposed algorithm and the implementations realized, Section 3 describes the multi-resolution volume rendering approach, Section 4 presents experimental results.

2. Proposed upscaling approach

The proposed technique enlarges images by subdividing original voxels into smaller ones. In this way it is possible to obtain integer scale factors along each directions: we implemented two schemes, one subdividing each voxel in two parts along the z direction (Fig. 1) obtaining a slice-based interpolation and another subdividing each voxel in eight smaller ones as in (Fig. 3) obtaining a $2\times$ zooming. The algorithm then computes the new voxel values through an optimization scheme using two constraints: the first is the constancy of the sum of the gray levels inside each splitted voxel, the second is the continuity of second order derivatives of the image brightness that has, in fact, shown good results in 2D image upscaling [GA08] and it is able to reduce the jagged artifacts of upscaled images. This assumption is also strictly related to the constant covariance constraint used by other well known edge directed interpolation methods [LO01].

The constancy of the gray level sum is a physically motivated constraint that assumes that the energy captured by the acquisition system in the original voxel location is the same that would have been captured by a sensor array with higher resolution. This assumption is correct for volumetric imaging only if the gray level is a linear function of the physical value measured by the device and if the so-called "slice thickness" of the imaging modality is equal to the "slice spacing", i.e. the effective voxel volumes do not overlap. This is not true, of course for all modalities and protocols, but it can be considered a reasonable approximation for many real cases.

The implementations of the method for the two cases considered (e.g. slice splitting and voxel splitting) are described in detail in the following subsections.

2.1. Slice splitting implementation

To split each slice we divide each original voxel $B(i, j, k)$ of the $W \times H \times N$ dataset in two smaller ones as in Figure 1 obtaining a new $W \times H \times 2N$ volume. Assuming that the gray level value in the original voxel is equal to the signal intensity per volume unit multiplied by the voxel size, we can assign to the two smaller voxels the same value one.



Figure 1: *Slice splitting: original voxels (A) are subdivided in two parts with half z -component (B); the optimization procedure smoothing derivatives keep the sum of the original voxels unchanged (C).*

The greedy optimization procedure consists of an iterative processing of splitted blocks. At each iteration, for each splitted voxel, that will be called in the following "block", a local energy function is computed in three different cases: unchanged gray level values, and the two "perturbed" configurations obtained adding a small δ value to one voxel and subtracting the same value to the other (Fig. 2). The average value in the block is in any case unchanged.

If i, j, k are the coordinates of the original grid corresponding to the subdivided block $B(i, j, k)$, for each perturbed configuration B_p , the computed local energy $E_p(i, j, k)$ is defined as the sum absolute values of the differences between second order derivatives inside and outside B (computed at the finer scale):

$$E(i, j, k)_p = \sum_{hk} \sum_{\vec{x} \in B(i, j, k)} \sum_{\vec{x}' \in N_{ijk}(\vec{x})} \left| H_{hk}(I_p(\vec{x})) - H_{hk}(I_p(\vec{x}')) \right| \quad (1)$$

where $\vec{x} = (x, y, z)$ are the coordinates of small voxels in the new grid, $N_{ijk}(\vec{x})$ is the set of neighboring voxels of \vec{x} that are not in the block B_{ijk} , $I_p(\vec{x})$ is the high-resolution volume created by the voxel splitting, modified only in the block (i, j, k) by replacing it with the "perturbed" version B_p . $H_{hk}(I_p(\vec{x}))$ are the Hessian matrix entries (second order derivatives). At the end of each complete blocks processing we finally update each block replacing the internal voxels values with the values of the modified block B_{min} corresponding to the lowest energy.

To avoid large memory usage, δ is an integer value and the representation of the original voxel value (unsigned char or short) is not changed. The first iteration are performed with $\delta = 4$ and then the values are reduced to $\delta = 1$.

2.2. Voxel subdivision in octants

In this case we split each original voxel in eight smaller ones (Fig. 3). Then energy terms are computed as before, by evaluating a sum of absolute values of the differences between first and second order derivatives of the gray levels inside and outside the splitted voxel. Possible voxel perturbations adding an integer value δ to one of the eight small voxels and subtracting δ from the value of another one are 56. Of




Figure 2: The three modified intensity distribution inside each splitted voxel tested at each iteration of the greedy procedure




Figure 3: Voxel subdivision in octants: original voxels (A) are subdivided in eight parts (B): the optimization procedure smoothing derivatives keep the sum of the original voxel values unchanged (C).

course, different sets of modified configurations with equal sum of intensities can be used, and actually we tested different options, obtaining negligible changes in the results.

At each iteration and for each splitted voxel, we compute for each of these configuration B_1, B_2, \dots, B_{56} and for the unchanged one B_0 the voxel energy, modifying then each block according to the perturbation minimizing it. To have a good convergence and smoother edges, we found in this case useful to first remove voxelization with another greedy iterative modification of the splitted voxel minimizing a local energy depending by the gray level continuity at the big voxel borders, i.e.:

$$E'(i, j, k)_p = \sum_{hk} \sum_{\vec{x} \in B(i, j, k)} \sum_{\vec{x}' \in N_{ijk}(\vec{x})} \left| I_p(\vec{x}) - I_p(\vec{x}') \right| \quad (2)$$

The iterative smoothing of second order derivatives performed using the same energy described by eq. 2 is performed after this step, providing the final result.

To speed up the computation, however, we test at each iteration only a random subset of the voxel perturbations: taking only one third of the possible configurations we found that the difference in the result is negligible.

3. Multi-resolution volume rendering

Edge adaptive image upscaling is applied usually to enhance printing quality or to upscale video for high resolution display. Medical data upscaling can be used as well to enhance 2D slices visualization on high resolution displays, but can,

of course used also to enhance volumetric visualization (e.g. ray casting) that is often used by radiologists and physician for diagnosis, being supported by recent radiology workstations. However, a big problem in handling upscaled volumes is the large amount of memory required to store the 3D grid, that can create problems even using high end graphics cards.

To avoid the necessity of storing the complete dataset in the graphics card memory and still have an interactive visualization of huge datasets, it is possible to exploit recent adaptive techniques like that presented in [GMI08]. This method uses an octree representation of the volume maintained out of core, and an adaptive loader executing on the CPU, that updates view and transfer function dependent subvolume maintained on GPU memory by asynchronously fetching data from the octree (Fig.4). In this way, even if the original data cannot be stored in the graphic card memory, the volume can still be interactively inspected, and it is possible to obtain the maximal visual quality for each part of the dataset just controlling viewpoint and camera parameters. The technique has shown very good results, but its medical application is still limited due to the fact that usual diagnostic dataset are not really huge and can be fitted into the memory of recent graphic cards. However, the increasing resolution of new imaging modalities, and the use of smart volume upscaling techniques like the one presented here (or, for example, a learning based method) could make the adaptive multi-resolution volume rendering quite useful.

4. Experimental results

4.1. Image quality measures

We tested both the z-upscaling and the isotropic upscaling on different test medical datasets: two MRI acquisition of the head, one with original resolution $256 \times 256 \times 60$ and voxel size $0.86 \times 0.86 \times 3mm.$, another brain MRI with original resolution $256 \times 256 \times 256$ and voxel size $1 \times 1 \times 1mm.$, a CT scan of the liver with original size $512 \times 512 \times 130$ and voxel size $0.89 \times 0.89 \times 1.25mm$ and a region of CT scan of the aortic bifurcation with original size $180 \times 180 \times 200$ and voxel size $0.74 \times 0.74 \times 1.25mm$. All the dataset have been converted to 8 bit grayscale.

For all the data sets we created two downsampled versions simulating the low resolution acquisition in the first case by merging couples of slices taking average values and in the second case by collapsing 8-neighboring voxels blocks into low resolution voxels taking the average value. We then upsampled the simulated low resolution acquisitions with different techniques, e.g trilinear interpolation, the 3D Spline interpolator implemented in Octave and the proposed method.

Figure 5 shows a visual comparison of corresponding interpolated slices of upscaled volumes: the new method creates evidently sharper images, even if it is not similarly effective in removing jaggies and smoothing contours.

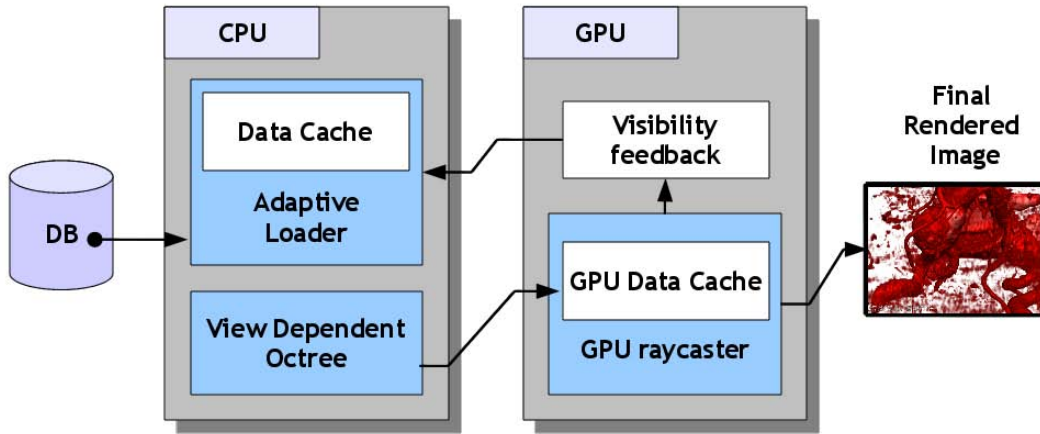


Figure 4: Scheme of the volume rendering technique

To give a quantitative evaluation of the similarity of down-scaled/upscaled volumes and original ones, we considered two error measures. The first is the Peak Signal to Noise Ratio (PSNR), defined as:

$$PSNR = 20 \log_{10} \frac{MAXVOX}{\sqrt{\frac{\sum_{i=1}^{dimx} \sum_{j=1}^{dimy} \sum_{k=1}^{dimz} (I_{up}(i,j,k) - I_{orig}(i,j,k))^2}{dimx * dimy * dimz}}} \quad (3)$$

where $I_{up}(i, j)$ is the upscaled subsampled volume, I_{orig} the original one, $dimx$, $dimy$ and $dimz$ the image dimensions and $MAXVOX$ the end scale value of the voxel intensity.

The other measure applied is the percentage of largely different voxels (PLDV) that counts how many voxels differs more than a threshold (we set this value to 5% of the maximum value). This measure could be related to the number of voxels that may be misclassified by a simple segmentation algorithm, meaning also that a smart interpolation could obtain a more precise segmentation of anatomical structures. Table 1 shows the results obtained with the described algorithm for the z-upscaling and Table 2 shows the same results for isotropic upscaling. In both cases the proposed method is more effective than the compared kernel-based techniques. The result is sufficiently encouraging to suggest a wider comparison involving different interpolation techniques and a different imaging modalities and acquisition protocols and to test practical application of offline volume supersampling.

4.2. Volume rendering

The good quality of the upscaled volumes appear clearly also in the described volumetric visualization.

Figures 6 shows the renderings, obtained with a ray casting procedure with same sampling step and transfer function of a CT scan (from the public database of University Hospital of Geneva, <http://pubimage.hcuge.ch:8080/>) of the tho-

dataset	Linear	Spline	Our Meth.
MR1	5.73	7.85	8.49
MR2	16.21	18.74	19.65
CT	16.97	19.26	20.07
CT2	21.80	26.89	28.36
AVG	14.98	18.19	19.14

dataset	Linear	Spline	Our Meth.
MR1	20.2	17.5	16.1
MR2	8.0	6.5	5.8
CT	7.5	5.7	5.1
CT2	5.9	3.1	2.2
AVG	10.4	8.2	7.3

Table 1: Top: PSNR(dB) obtained in the comparison of original medical datasets with downscaled versions (collapsed slices), upscaled in the z direction by a factor 2. Bottom: Percentage of largely different voxels (%), lower is better values.

rax, and of an upscaled version of it. With the original data ($512 \times 512 \times 743$), the resulting image is oversmoothed and some details vanish due to the excessive smoothing of the gray levels making correct tissue classification difficult (a). The visualization of the edge adaptive and energy preserving upscaled volume (isotropically enlarged of a $2 \times$ factor), performed with the same rendering parameters appears of a better quality (b). The enhancement of the result appears, however, if we zoom on a small region (right): the improvements in the visualization are particularly evident in presence of poorly contrasted tubular structures. Figure 7 shows similar results on a high resolution contrasted angio-CT scan from the same archive. Improvements are even more visible




Figure 5: The visual quality of corresponding slices of volumes upscaled (2x) with different methods is clearly different. From left to right: Nearest neighbor, Linear interpolation, that appears quite smoothed, Spline interpolation, still a bit smoothed, and proposed method, providing extremely sharp images.

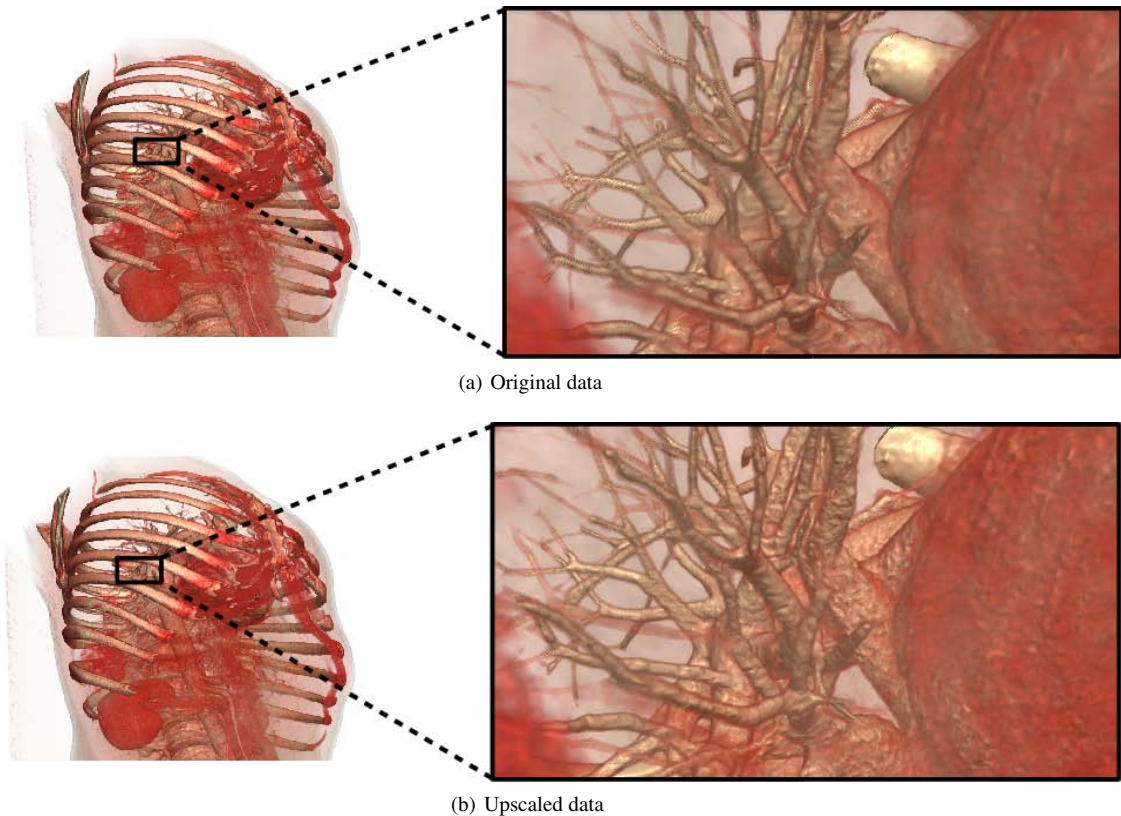


Figure 6: Volume visualization of a contrasted thoracic CT (data from University of Geneva): (a) rendering of the original data and a higher resolution detail. (b) Visualization of the same data set but processed with our upscaling approach. The visualization uses the same transfer function and sampling step. The enhancement in the perception of surfaces and tubular structures is more evident on the zoomed images.

with higher upscaling factors: Fig.8, shows volume rendering results (same transfer function and sampling step) of a contrasted CT acquisition of the iliac bifurcation and of its super-resolved version. Here the original dataset has been upscaled of a factor 4 along the x and y directions and of a

factor 8 along the z direction. Small structures appear clearly sharper and more contrasted after the super-resolution enhancement (right). The major limit of the current implementation of the optimization-based supersampling, is the lim-

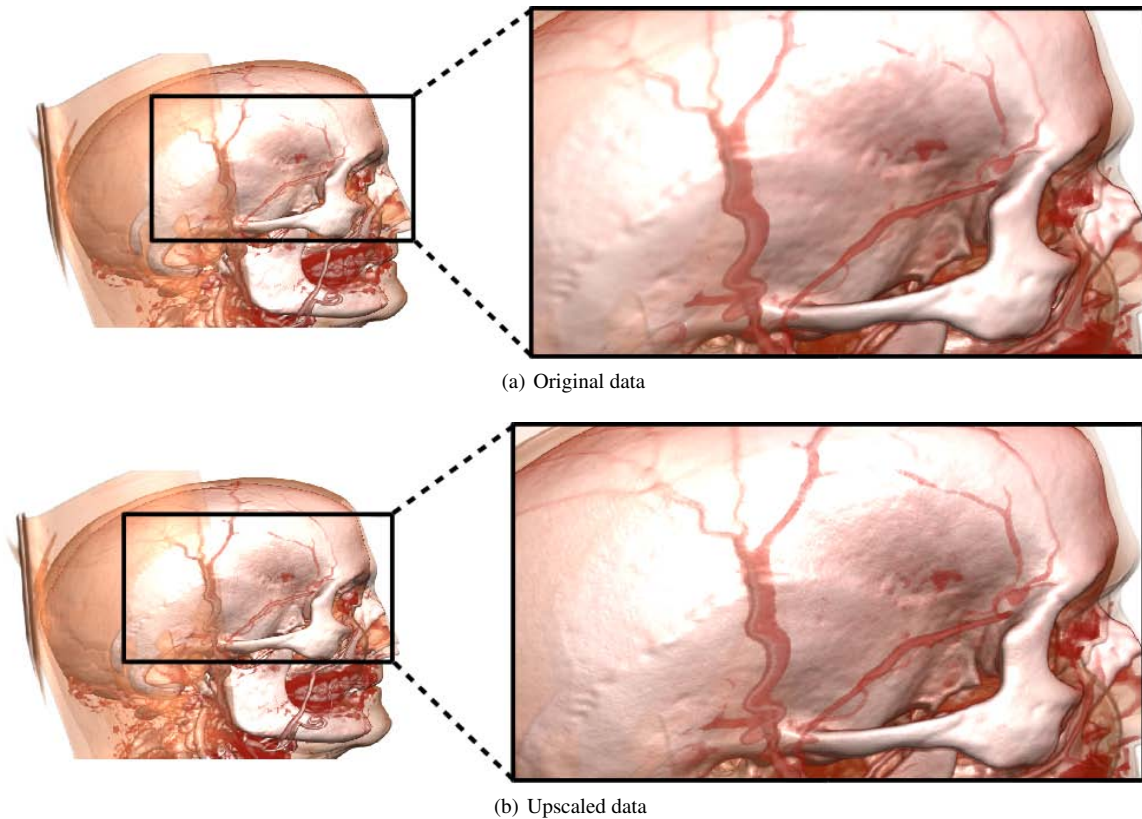


Figure 7: Volume visualization of a contrasted angio-CT of the head (data from University of Geneva): (a) rendering of the original data and a higher resolution detail. (b) Visualization of the same data set but processed with our upscaling approach. The visualization uses the same transfer function and sampling step. Notice how finer structures and vessels are easier to follow.

ited reduction of the voxelization artifacts, that could be, however, handled with a specific post processing.

If data are upscaled of a very large factor, it is still possible to visualize it interactively using the adaptive techniques previously described. The method used for the visualization builds on the ability to rapidly traverse an octree structure, and is based on the stackless ray traversal method for kd-trees [HBS98, PGSS07] and recently adapted to GPU volume rendering [GMI08, CNLE09]. We employ an extended version [IGM10] to support preintegration employing a CUDA algorithm which takes advantage of visibility information gathered during the octree traversal to avoid loading occluded data. Spatio-temporal coherence can be exploited taking advantage of the scatter memory write capability of CUDA threads to mark the visibility of all rendered octree blocks.

5. Conclusions

We presented new algorithms for edge-directed upscaling of medical datasets, based on voxel splitting and on an opti-

mization procedure smoothing first and second order derivatives of the gray level while keeping the intensity inside each splitted voxel constant. The algorithm provides results that appear very sharp, even if does not completely remove voxelization. However, objective tests measuring the similarity between the original high resolution data and corresponding subsampled/upsampled have shown that this similarity is higher than that obtained by kernel-based upscaling method. Furthermore, we have shown a reasonable application of similar algorithms: the enhancement of ray casting volume visualization through the use of the edge-directed upsampling of the visualized data. The method requires the use of a recent adaptive volume rendering technique able to handle the interactive visualization of huge datasets. Volumes are not only pre-processed in order to create a multiscale representation by downsampling original ones, but upscaling is performed to enhance the quality of the visualization of small details.

The perceived higher quality of the upscaled volume rendering shows that the technique can be successfully ap-




Figure 8: The rendering of the adaptively upsampled volume (right) provides a sharper perception of small details (look at thin vessels) with respect to a similar visualization (same transfer function and sampling rate) of the original volume (left)

	Linear	Spline	Our Meth.
MR1	9.55	13.07	14.22
MR2	3.63	6.61	7.47
CT	4.62	7.62	8.82
CT2	15.54	22.01	23.97
AVG	8.33	12.32	13.62

	Linear	Spline	Our Meth.
MR1	14.4	11.2	9.8
MR2	24.3	19.9	17.8
CT	20.3	16.7	14.9
CT2	11.0	5.4	3.3
AVG	17.5	13.3	11.4

Table 2: Top: PSNR(dB) obtained in the comparison of original medical datasets with downsampled versions (collapsed voxels) upsampled isotropically by a factor 2. Bottom: same comparison using PLDV (%), lower is better).

plied in the clinical domain to enhance the perception of 3D anatomy.

Acknowledgments

This work is partially supported by the EU Marie Curie Program under the 3DANATOMICALHUMAN project (MRTNCT-2006-035763).

References

- [CNLE09] CRASSIN C., NEYRET F., LEFEBVRE S., EISEMANN E.: Gigavoxels: Ray-guided streaming for efficient and detailed voxel rendering. In *Proc. I3D* (2009), pp. 15–22. 6
- [Fat07] FATTAL R.: Image upsampling via imposed edge statistics. *ACM Transactions on Graphics* 26, 3 (2007), 95. 1
- [FJP02] FREEMAN W., JONES T., PASZTOR E.: Example-based

super-resolution. *IEEE Computer Graphics and Applications* 22, 2 (2002), 56–65. 1

- [GA08] GIACHETTI A., ASUNI N.: Fast artifact free image interpolation. In *Proc. BMVC 2008* (2008). 1, 2
- [Gia10] GIACHETTI A.: Irradiance preserving image interpolation. In *Proc. IEEE Int. Conf. Pattern Recognition* (2010). 1
- [GMI08] GOBBETTI E., MARTON F., IGLESIAS GUITIÁN J.: A single-pass GPU ray casting framework for interactive out-of-core rendering of massive volumetric datasets. *The Visual Computer* 24, 7-9 (2008), 797–806. Proc. CGI 2008. 3, 6
- [HBS98] HAVRAN V., BITTNER J., SÁRA J.: Ray tracing with rope trees. In *Proc. Spring Conf. Comput. Graph* (1998), pp. 130–140. 6
- [HZB92] HERMAN G. T., ZHENG J., BUCHOLTZ C. A.: Shape-based interpolation. *IEEE Comput. Graph. Appl.* 12, 3 (1992), 69–79. 1
- [IGM10] IGLESIAS GUITIÁN J., GOBBETTI E., MARTON F.: View-dependent exploration of massive volumetric models on large-scale light field displays. *The Visual Computer* 26 (2010), 1037–1047. 10.1007/s00371-010-0453-y. 6
- [KK08] KIM K. I., KWON Y.: Example-based learning for single-image super-resolution. In *Proceedings of the 30th DAGM symp. on Patt. Rec.* (Berlin, Heidelberg, 2008), Springer-Verlag, pp. 456–465. 1
- [KTN*07] KAMIMURA K., TSUMURA N., NAKAGUCHI T., MIYAKE Y., MOTOMURA H.: Video super-resolution using tex-ton substitution. In *ACM SIGGRAPH 2007 posters* (New York, NY, USA, 2007), ACM, p. 63. 1
- [LO01] LI X., ORCHARD M. T.: New edge-directed interpolation. *IEEE Trans. on Image Proc.* 10 (2001), 1521–1527. 1, 2
- [PGSS07] POPOV S., GÜNTHER J., SEIDEL H.-P., SLUSALLEK P.: Stackless kd-tree traversal for high performance GPU ray tracing. *Computer Graphics Forum* 26, 3 (2007), 415–424. 6
- [PSR*04] PENNEY G., SCHNABEL J., RUECKERT D., VIERGEVER M., NIESSEN W.: Registration-based interpolation. *IEEE Transactions on Medical Imaging* 23, 7 (2004), 922–926. 1
- [SXS08] SUN J., XU Z., SHUM H.: Image super-resolution using gradient profile prior. In *CVPR08* (2008), pp. 1–8. 1
- [TMSV09] TORONTO N., MORSE B., SEPPI K., VENTURA D.: Super-resolution via recapture and bayesian effect modeling. In *Proc. IEEE Conference on Comp. Vision and Pattern Recognition* (2009), pp. 2388–2395. 1