Anti-Aliased Euclidean Distance Transform on 3D Sampling Lattices

Elisabeth Linnér and Robin Strand

Center for Image Analysis, Uppsala University, Sweden
{elisabeth,robin}@cb.uu.se

Abstract. The Euclidean distance transform (EDT) is used in many essential operations in image processing, such as basic morphology, level sets, registration and path finding. The anti-aliased Euclidean distance transform (AAEDT), previously presented for two-dimensional images, uses the gray-level information in, for example, area sampled images to calculate distances with sub-pixel precision. Here, we extend the studies of AAEDT to three dimensions, and to the Body-Centered Cubic (BCC) and Face-Centered Cubic (FCC) lattices, which are, in many respects, considered the optimal three-dimensional sampling lattices. We compare different ways of converting gray-level information to distance values, and find that the lesser directional dependencies of optimal sampling lattices lead to better approximations of the true Euclidean distance.

1 Introduction

1.1 Supersampling and Coverage

In a binary image, the spatial elements (spels) are classified either as part of an object or of the background. When imaging a continuous scene, this leads to jagged edges, and much of the information on the edge location and length is lost. With inspiration from anti-aliasing in computer graphics, it is shown in [13] that part of this information is preserved if the image is supersampled, and the intensity of a spel is proportional to the part of its Voronoi cell that is inside of an object, as in Figure 1. The intensity $c(\mathbf{p})$ of a spel $\mathbf{p}$ is then referred to as its coverage value. The theory is further developed in [12], and it is shown that some physical properties, such as volume and surface area, can be measured with higher accuracy in supersampled images. These qualities are important in, for example, cancer diagnostics [3, 7].

In 2D imaging devices, the intensity of a spel is usually computed through integration over some environment of the spel center [2], and the result is more or less equivalent to a coverage value. 3D imaging techniques, such as computed tomography (CT), are often based on combining 2D images from different perspectives, and thus we feel comfortable to suggest and apply coverage based image processing methods in 3D as well.

E. Barcucci et al. (Eds.): DGCI 2014, LNCS 8668, pp. 88–98, 2014.

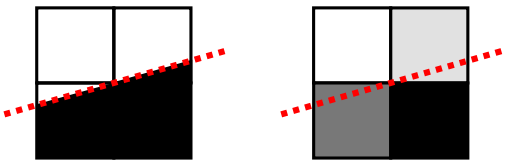

(a) continuous edge (b) coverage values

Fig. 1. Assignment of intensity proportional to spel coverage

1.2 Euclidean Distance Transforms

The Euclidean distance d_E between two points $\mathbf{p}_1$ and $\mathbf{p}_2$ in n-dimensional space is defined as

$$d_E(\mathbf{p}_1, \mathbf{p}_2) = \sqrt{\sum_{i=1}^{n} (\mathbf{p}_1(i) - \mathbf{p}_2(i))^2}.$$

The Euclidean distance transform (EDT) maps every point $\mathbf{p}$ to

$$d_{EDT}(\mathbf{p}) = \inf_{\mathbf{p}_\omega \in \Omega} (d_E(\mathbf{p}, \mathbf{p}_\omega)),$$

where Ω is some object(s) or a set of seed points. This transform is used in many essential operations in image processing, such as basic morphology, level sets, registration and path finding [1]. In 3D, it can be used for, among other things, visualization, modeling and animation [6].

Anti-Aliased Euclidean Distance Transform. The anti-aliased Euclidean distance transform (AAEDT) for two-dimensional images is presented in [5]. It uses coverage information to compute the Euclidean distance from an object with sub-spel precision. The distance d_{AAEDT} is defined as

$$d_{AAEDT}(\mathbf{p}) = \min_{\mathbf{p}_\omega \in \partial\Omega} (|\mathbf{d}_E(\mathbf{p}, \mathbf{p}_\omega)| + d_f(\mathbf{p}_\omega)), \quad (1)$$

where $|\mathbf{d}_E(\mathbf{p}, \mathbf{p}_\omega)|$ is the Euclidean distance between the centers of a background spel $\mathbf{p}$ and an edge spel $\mathbf{p}_\omega \in \partial\Omega$, $\partial\Omega$ being the edge of a binary and areasampled object Ω, and $d_f(\mathbf{p}_\omega)$ is the distance between the edge and the center of the edge spel. This is illustrated in Figure 2. The addition of the term $d_f(\mathbf{p}_\omega)$ is the source of the sub-pixel precision, which improves, for example, the accuracy of level sets and small-scale morphology.

1.3 Three-Dimensional Sampling Lattices

The spels in a digital image represent sample points, which are organized in a so-called sampling lattice. The most common sampling lattice is the Cartesian Cubic (CC) lattice, resulting in square spels in 2D, and cubic spels in 3D.

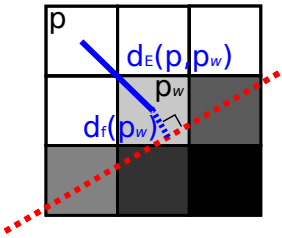

Fig. 2. Illustration of the anti-aliased Euclidean distance d_{AAEDT} between two spels p and p_ω

Unfortunately, the sampling properties of the CC lattice are strongly direction dependent [8, 14]. Consequently, to guarantee some minimum resolution in all directions within the image, some directions must be over-sampled, and redundant data must thus be stored and processed. An alternative is to use the Body-Centered Cubic (BCC) and Face-Centered Cubic (FCC) sampling lattices. Their direction dependence is much weaker, and, for band-limited signals, a minimum resolution in all directions can be obtained using $\sim 30\%$ fewer sample points than if a CC lattice were used [4, 8, 10, 14, 15].

1.4 Scope of This Paper

In AAEDT, $d_f(\mathbf{p}_\omega)$ is computed from the coverage value $c(\mathbf{p}_\omega)$ of the edge spel, usually under the assumption that the object surface intersecting the spel is locally flat, and that $d_f(\mathbf{p}_\omega)$ is measured along the surface normal [5]. However, as the orientation of the surface is unknown, the asymmetry of the Voronoi cell of the spel makes it difficult to construct a mapping between $c(\mathbf{p}_\omega)$ and $d_f(\mathbf{p}_\omega)$. Moreover, for spels not located on the edge, the vector propagation process may introduce a discrepancy between the direction to the nearest edge spel and that to the closest point on the surface. We propose that the lesser direction dependencies of the BCC and FCC sampling lattices, compared to the CC sampling lattice, lead to improved performance of AAEDT.

2 Method

2.1 Implementation

We use the graph-based AAEDT implementation presented in [9], which can be adapted to any dimensionality and sampling lattice by changing the definition of the spel neighborhood. In this way, we ensure fair comparison of the lattices.

2.2 Computation of $d_f(\mathbf{p}_\omega)$

We want to find an expression on the form $d_f(c)$ to approximate the value $d_f(\mathbf{p}_\omega)$ of an edge spel $\mathbf{p}_\omega$ from its coverage value $c(\mathbf{p}_\omega)$. As a reference for

different approximation methods, we simulate $d_f(\mathbf{p}_\omega)$ and $c(\mathbf{p}_\omega)$ for the Voronoi cells of the CC, BCC and FCC lattices: We construct a set of planes, the normals of which are uniformly distributed within the symmetry regions of the Voronoi cells, intersecting the cell center. Using the Monte-Carlo sampling method [11], we approximate the portions of the cell that are located above and below the plane. We repeat the process while moving the plane away from the spel center in small steps, until the entire Voronoi cell is below the plane. The simulation output is shown in Figure 3, with plots of three different approximations $d_f(c)$ of $d_f(\mathbf{p}_\omega)$ from $c(\mathbf{p}_\omega)$, described below.

The implementation in [5] approximates $d_f(\mathbf{p}_\omega)$ within an edge spel using

$$d_f(c) = 0.5 - c(\mathbf{p}_\omega). \tag{2}$$

This is the exact relationship between $d_f(\mathbf{p}_\omega)$ and $c(\mathbf{p}_\omega)$ for for a spel of a CC lattice being intersected by a plane perpendicular to a lattice vector. Although $d_f(c)$ is only computed in the initialization of AAEDT, the computational simplicity of (2) is an attractive property. There is no equally simple formula for BCC and FCC lattices. However, linear regression on the simulation output in Figure 3 shows that (2) is actually an even better approximation of the relationship between $d_f(\mathbf{p}_\omega)$ and $c(\mathbf{p}_\omega)$ on these lattices, than on the CC lattice. Higher order regression on the data in Figure 3 leads to overfitting rather than improvement of the approximation.

As we do not calculate the orientation of the plane that intersects the spel, we want $d_f(c)$ to be orientation independent. As the ideal Voronoi cell of a sample point, with respect to sampling properties, is a ball, which is completely orientation independent, we derive $d_f(c)$ for a ball of unit volume. As our implementation uses 256 gray levels, we tabulate the relationship for 256 coverage values in the interval $[0, 1]$.

For the third approximation method plotted in Figure 3, we simply use the mean value of the Monte-Carlo simulation output. Again, we tabulate the relationship for 256 coverage values in the interval $[0, 1]$. This is the only approximation method in this study that is lattice dependent.

3 Experiments

3.1 Choice of Test Images

We study the behavior of AAEDT applied to images of supersampled binary balls, with 256 gray levels representing degrees of spel coverage, sampled on CC, BCC and FCC lattices. We use the exterior and interior distance from the ball surface as examples of convex and concave surfaces, respectively, of different orientations. By varying the ball radius r_s within some range $r_s \in [r_{min}, r_{max}], r_{min}, r_{max} \gg r_v$, where r_v is the average radius of a spel, we study the impact of surface curvature on the accuracy of AAEDT.

We use the distance from balls where $r_s \approx r_v$ to indicate the behavior of AAEDT for undersampled objects.

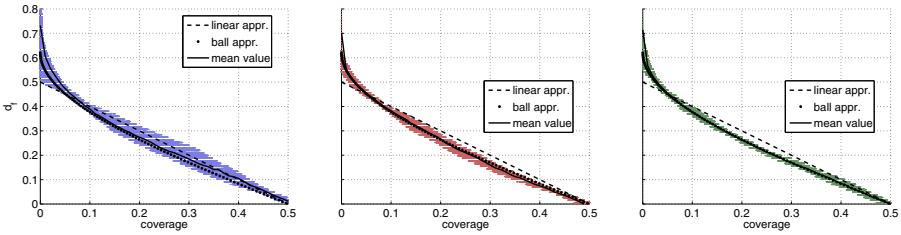

(a) CC lattice (b) BCC lattice (c) FCC lattice

Fig. 3. Approximative mappings between $d_f(\mathbf{p}_\omega)$, the distance between the intersecting plane and spel center, and the $c(\mathbf{p}_\omega)$, the spel coverage value, plotted on top of the output of a Monte-Carlo simulation of the relationship. As $d_f(c)$ behaves as an odd function centered at $c(\mathbf{p}_\omega) = 0.5$, it is only plotted for $c \in [0, 0.5]$.

For every lattice, we use a set of balls where the center points are evenly distributed within the symmetry region of the Voronoi cell of that lattice, so that the ball center is unlikely to coincide with a sample point. The sample density for all lattices is one spel per unit volume.

4 Results

The results are expressed in terms of

$$\epsilon(\mathbf{p}) = d_{EDT}(\mathbf{p}) - d_{AAEDT}(\mathbf{p}),$$

and its mean value $\epsilon(\mathbf{p})_m$, where $d_{EDT}(\mathbf{p})$ is the exact Euclidean distance transform, the unsigned relative error

$$|\epsilon_r(\mathbf{p})| = \frac{|d_{EDT}(\mathbf{p}) - d_{AAEDT}(\mathbf{p})|}{d_{EDT}(\mathbf{p})},$$

and the mean unsigned relative error $|\epsilon_r(\mathbf{p})|_m$. The error is evaluated within a Euclidean distance of 50 units from the ball surface, where the unit distance is defined in relation to the unit volume of the spels.

4.1 Bias Error

The mean error $\epsilon(\mathbf{p})_m$, computed on the CC, BCC and FCC lattices using linear, ball-based and mean value-based approximations of $d_f(c)$, is shown in Figures 4, 5 and 6.

For convex surfaces with $r_s \gg r_v$, and linear approximation of $d_f(c)$, the distance is underestimated on all lattices, although much more so on the CC lattice. The mean error on the CC lattice also exhibits a larger standard deviation than that on the BCC and FCC lattices. On the CC lattice, both the mean error and standard deviation are reduced by ball-based approximation of $d_f(c)$, and

even further by mean value-based approximation. The error on the BCC lattice seems to be almost completely unbiased for ball-based approximation relative to the others, and tends somewhat towards underestimation for mean value-based approximation. On the FCC lattice, we see a tendency towards underestimation for all approximation methods, but it is the least apparent for ball-based approximation.

For concave surfaces with $r_s \gg r_v$, the mean error is very close to being unbiased for small r_s, and tends towards underestimation for less curved surfaces. As the bias increases, the standard deviation grows. As for convex surfaces, the result on the CC lattice is improved by ball-based and mean value-based approximation of $d_f(c)$. On the BCC and FCC lattices, it is very difficult to discern any bias for ball-based approximation, while there is a slight underestimation using mean value-based approximation.

For $r_s \approx r_v$, all lattices and approximations of $d_f(c)$ yield equivalent results, always underestimating the distance.

4.2 Error Range

Figures 7, 8 and 9 show the first (25^{th} percentile), second (median) and third (75^{th} percentile) quartiles of $|\epsilon_r(\mathbf{p})|$.

For convex surfaces with $r_s \gg r_v$, the CC lattice is clearly outperformed by the BCC and FCC lattices. The range of $|\epsilon_r(\mathbf{p})|$, indicated by the 25^{th} and 75^{th} percentiles, is highly concentrated around the median error on BCC and FCC, while it is notably larger on the CC lattice. However, the performance of the CC lattice is very much improved when ball-based or mean value-based approximation of $d_f(c)$ is used.

For concave surfaces, and for $r_s \approx r_v$, the performance is almost equivalent on all lattices. However, the growth of the range of $|\epsilon_r(\mathbf{p})|$, that occurs for small r_s, starts at an earlier stage on the CC lattice than on BCC and FCC.

In Figures 10, 11 and 12, we see $|\epsilon_r(\mathbf{p})|_m$ as a function of the distance d to the ball surface.

The approximation of $d_f(c)$ takes place at $d < 1$, at which state the difference between the lattices is small. In the cases of $r_s \gg r_v$, the BCC and FCC lattices seem to be at a small advantage in making this approximation, compared to the CC lattice, when ball-based or mean value-based approximation is used.

For all curvatures, lattices and approximations of $d_f(c)$, $|\epsilon_r(\mathbf{p})|_m$ decreases rapidly as the distance increases. For $r_s \gg r_v$, this decrease is more rapid on the BCC and FCC lattices than on the CC lattice.

5 Discussion

5.1 Bias Errors

The AAEDT has two possible sources of bias errors: The approximation of $d_f(c)$, and the risk that $|\mathbf{d}_{AAEDT}(\mathbf{p}, \mathbf{p}_\omega)|$ is incorrect due to omitted edge spels, as

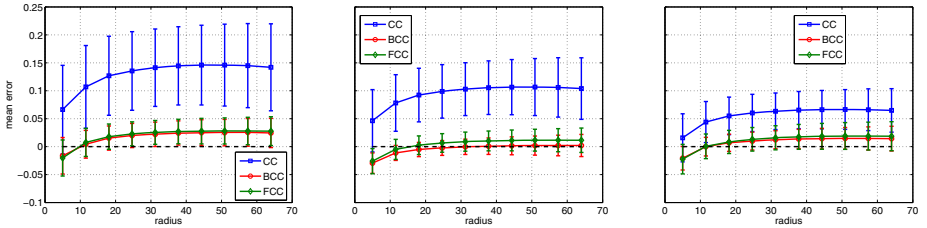

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 4. Mean error with one standard deviation, convex surface, $r_s \gg r_v$

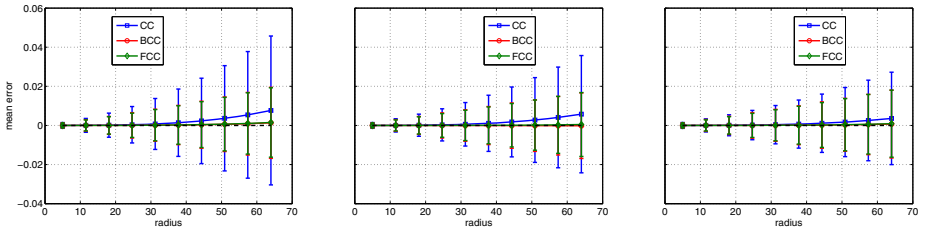

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 5. Mean error with one standard deviation, concave surface, $r_s \gg r_v$

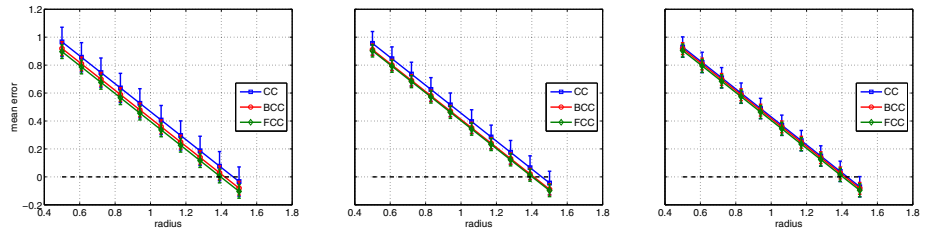

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 6. Mean error with one standard deviation, convex surface, $r_s \approx r_v$

described in [9]. The former may cause both over- and underestimation, while the latter always causes overestimation of the distance.

Underestimation is likely a result from the fact that the distance from an edge to the center of an edge spel $\mathbf{p}_\omega$ depends not only on $c(\mathbf{p}_\omega)$, but also on the edge orientation. This is very prominent on the CC lattice, where the variation is a factor of $\sqrt{2}$, making it very difficult to make a representative model for mapping distance to coverage. In Figure 3, we can see that this causes problems for both linear and mean-value based approximation of $d_f(c)$, as the large variance close to $c = 0$ and $c = 1$ dulls the slope of the curve, leading to a large difference between

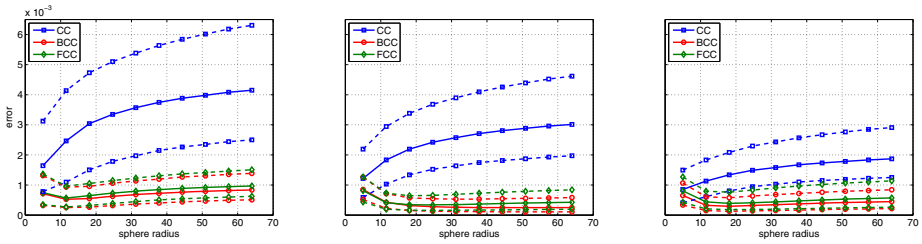

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 7. First, second and third quartiles, convex surface, $r_s \gg r_v$

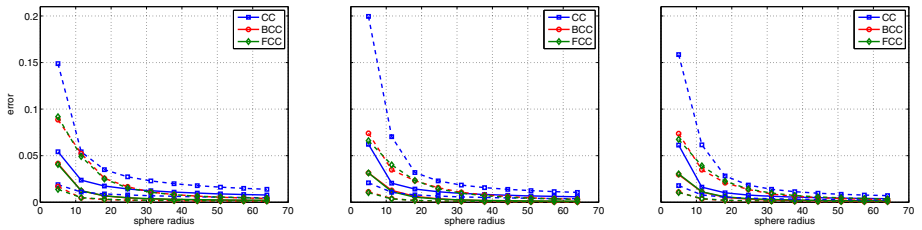

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 8. First, second and third quartiles, concave surface, $r_s \gg r_v$

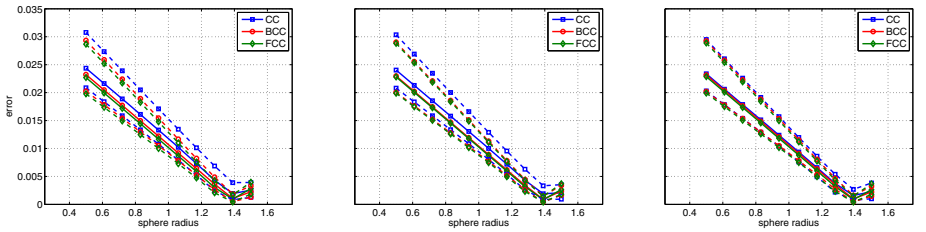

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 9. First, second and third quartiles, convex surface, $r_s \approx r_v$

the largest $d_f(c)$ that can be assigned to an edge spel, and the smallest d_{AAEDT} that can be assigned to a background spel. For a ball, we have $\lim_{c_{ball} \to 0} d_f(c) = r_v$ and $\lim_{c_{ball} \to 1} d_f(c) = -r_v$, with a smooth transition from fully covered (or uncovered) to partly covered. It is possible that this, combined with the low directional dependencies of the BCC and FCC lattices, results in the low bias observed for the ball-based approximation of $d_f(c)$.

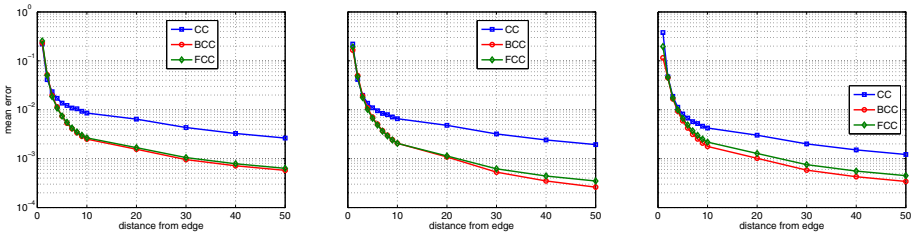

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 10. Mean unsigned relative error vs. convex surface, $r_s \gg r_v$

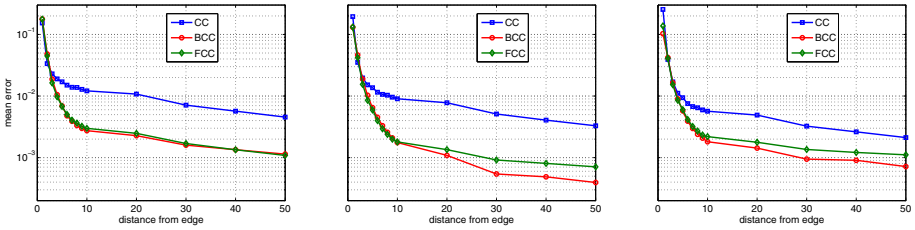

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 11. Mean unsigned relative error vs. concave surface, $r_s \gg r_v$

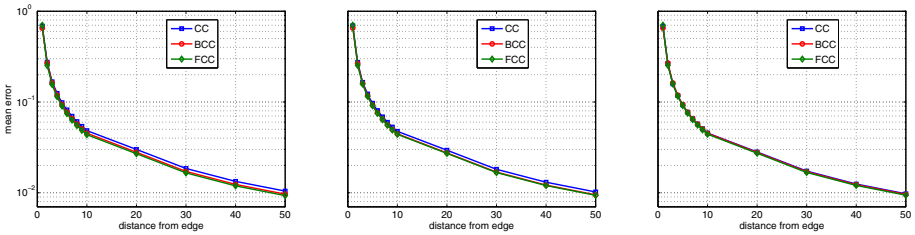

(a) linear $d_f(c)$ (b) ball coverage $d_f(c)$ (c) mean value $d_f(c)$

Fig. 12. Mean unsigned relative error vs. convex surface, $r_s \approx r_v$

5.2 Error Range

Figures 10, 11 and 12 show that $|\epsilon_r(\mathbf{p})|_m$ is smaller farther away from the surface. From this, we draw the conclusion that errors arise mainly from the approximation of $d_f(c)$, and not from the vector propagation.

It seems that the effect of improving the approximation of $d_f(c)$ is the most noticeable on the CC lattice. This is expected, as the cubic spels are more directionally dependent than the truncated octahedra and rhombic dodecahedra

of the BCC and FCC lattices, respectively. Ironically, even the linear approximation used in [5] and [9] is much less accurate on the CC lattice than on BCC and FCC. Actually, although considerably improved when using the mean value-based approximation of $d_f(c)$, the performance of AAEDT on the CC lattice is still not as good as that on the BCC and FCC lattices, using only the linear approximation.

The best performance is achieved by using the ball-based approximation of $d_f(c)$ on the BCC and FCC lattices. As explained above, this behavior is likely to be a consequence of the treatment of edge spels with $c \approx 0$ and $c \approx 1$.

The increase of the median of $|\epsilon_r(\mathbf{p})|$ for large r_s in Figure 7 does not necessarily mean that AAEDT is less accurate for flat surfaces. As we compute the error for $d \leq 50$ for all r_s, the ratio $n_{\text{close}}/n_{\text{far}}$, where n_{close} is the number of spels close to the surface and n_{far} the number of spels far away, is smaller for small balls. Hence, the ratio $n_{\text{large error}}/n_{\text{small error}}$ is also smaller. As the decrease in $|\epsilon_r(\mathbf{p})|$, visible in Figures 10 and 11, seems to be smooth, this affects both the mean value and median of $|\epsilon_r(\mathbf{p})|$.

The bad performance for $r_s \approx r_v$ shows that AAEDT is highly dependent on the sampling density being proportional to image scale, as $d_f(\mathbf{p}_\omega)$ cannot be accurately approximated from surfaces that are not locally flat.

5.3 Conclusions and Future Work

In this paper, we show how the performance of AAEDT in 3D can be significantly improved. We analyze the impact of the approximation of $d_f(c)$, and we explore the advantages of sampling lattices with lesser directional dependencies than the wide-spread CC lattice. Next, we hope to investigate the behavior of AAEDT in the presence of sharp corners and more complex surface curvature. In [5], it is suggested that gradient information is used to estimate the orientation of the surface, which may improve the performance on the BCC and FCC lattices even further.

Acknowledgements. The authors want to thank M.Sc. Max Morén for enabling the production of test images on the BCC and FCC lattices.

References

1. Borgefors, G.: Applications using distance transforms. In: Arcelli, C., Cordella, L.P., Sanniti di Baja, G. (eds.) Aspects of Visual Form Processing: Proceedings of the Second International Workshop on Visual Form, pp. 83–108. World Scientific Publishing (1994)
2. Chen, W., Li, M., Su, X.: Error analysis about ccd sampling in fourier transform profilometry. Optik - International Journal for Light and Electron Optics 120(13), 652–657 (2009)
3. Clifford Chao, K.S., Ozyigit, G., Blanco, A.I., Thorstad, W.L., O Deasy, J., Haughey, B.H., Spector, G.J., Sessions, D.G.: Intensity-modulated radiation therapy for oropharyngeal carcinoma: impact of tumor volume. International Journal of Radiation Oncology*Biology*Physics 59(1), 43–50 (2004)

4. Entezari, A.: Towards computing on non-cartesian lattices. Tech. rep. (2006)
5. Gustavson, S., Strand, R.: Anti-aliased Euclidean distance transform. Pattern Recognition Letters 32(2), 252–257 (2011)
6. Jones, M.W., Bærentsen, J.A., Sramek, M.: 3D distance fields: A survey of techniques and applications. IEEE Transactions on Visualization and Computer Graphics 12(4), 581–599 (2006)
7. Lebioda, A., Żyromska, A., Makarewicz, R., Furtak, J.: Tumour surface area as a prognostic factor in primary and recurrent glioblastoma irradiated with 192ir implantation. Reports of Practical Oncology & Radiotherapy 13(1), 15–22 (2008)
8. Linnér, E., Strand, R.: Aliasing properties of voxels in three-dimensional sampling lattices. In: Lirkov, I., Margenov, S., Waśniewski, J. (eds.) LSSC 2011. LNCS, vol. 7116, pp. 507–514. Springer, Heidelberg (2012)
9. Linnér, E., Strand, R.: A graph-based implementation of the anti-aliased Euclidean distance transform. In: International Conference on Pattern Recognition (August 2014)
10. Meng, T., Smith, B., Entezari, A., Kirkpatrick, A.E., Weiskopf, D., Kalantari, L., Möller, T.: On visual quality of optimal 3D sampling and reconstruction. In: Proceedings of Graphics Interface (2007)
11. Metropolis, N., Ulam, S.: The monte carlo method. Journal of the American Statistical Association 44(247), 335–341 (1949)
12. Sladoje, N., Lindblad, J.: High-precision boundary length estimation by utilizing gray-level information. IEEE Transactions on Pattern Analysis and Machine Intelligence 31(2), 357–363 (2009)
13. Sladoje, N., Nyström, I., Saha, P.K.: Measurements of digitized objects with fuzzy borders in 2D and 3D. Image and Vision Computing 23(2), 123–132 (2005)
14. Strand, R.: Sampling and aliasing properties of three-dimensional point-lattices (2010), `http://www.diva-portal.org/smash/record.jsf?searchId=1&pid=diva2:392445&rvn=3`
15. Theußl, T., Möller, T., Gröller, M.E.: Optimal regular volume sampling. In: Proceedings of the Conference on Visualization (2001)