Exploring Volumetric Data Using Interactive Statistical Views

Christian Kehl[1,2,3], Dongfang Qu[1,2], James R. Mullins[4], Sophie Viseur[3], Simon J. Buckley[1,2], John A. Howell[4], Robert L. Gawthorpe[2]

[1] UniResearch AS – CIPR, P.O.Box 7810, N-5020 Bergen
[2] University of Bergen – Dept. Earth Science, P.O.Box 7803, N-5020 Bergen, Christian.Kehl@geo.uib.no
[3] Université Aix-Marseille, 3 Place Victor Hugo, F-13331 Marseille
[4] University of Aberdeen – Dept. Geology & Petroleum Geology, AB24 3UE Aberdeen, UK

Abstract. Volumetric data for scientific purposes and applications are made increasingly available with novel modes of acquisition and modelling, which need to be visualised in order to understand the data and derive structured insight. Various methods for volume visualisation exist, while a persisting challenge is the interaction and delineation of object of interest within the data. Transfer functions, commonly used in medical visualisation, are the established means of interaction, which are cognitively challenging to setup and understand. This article presents new approaches for the coupled visualisation of volumetric data and their statistical derivatives to support the interactive data exploration. The presented techniques improve the means of volumetric data exploration in scientific disciplines and application cases for which established transfer function techniques are inadequate, such as structural- and petroleum geology.

1 Introduction

Volume visualization received increasing academic attention over the last years with the advent and availability of volumetric data through novel acquisition, modelling and computation techniques. The rendering of large and complex volumetric datasets, which was previously infeasible to be done interactively and in real-time (see original papers by Levoy[1] and Sabella[2]), is facilitated by advances in graphics hardware and GPU computing. An increasing number of scientific disciples and applications are able to use and analyse volumetric 3D/4D data due to this data availability and the rendering technology progress. Volume visualisation is an established technology within the medical sciences that found its way into practice. Apart from the medical sciences, volume visualisation receives increasing attention in geophysics, replacing previous 2D-slice workflows[3], and approaches for volumetric sketching in geology have been recently explored (see Patel[4], Lidal[5] and Natali[6] for details).

Transfer functions are the most established and recognized technique for interacting with volume visualisations of scalar gridded fields. A transfer function defines a mapping of colour and opacity for a selected value range, often via histograms. The specification of transfer functions for complex datasets and novel

application scenarios is challenging for novice- and adept users[7]. In these cases, the tool limits the analytical progress as the mapping definition is important for structuring the data. This challenge is a potential reason for the slow adoption of enhanced volume visualisation methods in disciplines other than the medical sciences, for example the geosciences or astronomy (see the comparison of current workflow tools in Fig. 1). Initial literature on the volume visualisation was published nearly simultaneously (see Sabella[2] for geosciences and Levoy[1] for anatomical use cases), but since then transfer functions appear to be not equally appropriate for both application cases.

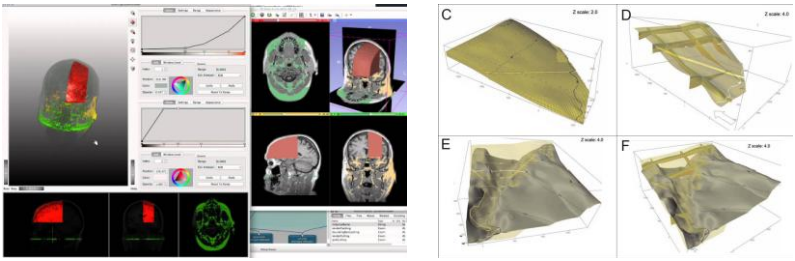

Fig. 1 Comparison of current visualization tools in medical sciences (left, ©MeVisLab) and geosciences (right, Roxar RMS Suite, from Enge et al. 2007).

Within the approach presented in this article, we address the problem by introducing auxiliary, statistical methods to interact with a volume visualisation. The coupled, interactive statistical views are designed to highlight the relation between statistical properties and the volumetric data. The coupled views provide scientists and practitioners with a new alternative to explore complex volume datasets. Geoscientists are the particular target audience of this approach due to their common background in geostatistics. The feasibility and interaction improvement of this approach is assessed on an anatomical dataset and a fault damage zone model from the domain of structural geology.

2 Related Work

The related body of literature for this work splits into the volume rendering methods and the related means of interaction (particularly statistics) to define colour-opacity mappings to specific scalar ranges of volume data.

The original approach for directly rendering volumetric data is “Cuberilles”[8], where voxel cubes are reconstructed using their boundary quadrilles. The rendering of semi-transparent opacity mappings (common to volume visualisation) via Cuberilles still poses challenges originating from the object-order-based rendering approaches that result in opacity confusers. Original papers considered the rendering of fully-opaque structures[9]. Only few up-to-date publications (such as Botha and Post[10]) discuss the issue of generic, object-ordered transparent surface rendering. The lacking attention of this topic is due to the rise of direct volume rendering (DVR) and its inherent order-independence. Direct volume rendering was introduced by Levoy[1], based on raycasting a stacked 3D volume texture. Most direct volume visualisation techniques are commonly controlled using

transfer functions. A challenge similar to each object's scalar mapping in a transfer function is the definition of its isocontour (such as for marching cubes[11]).

The transfer function controls a scalar volume visualisation by mapping opacity *A* and colour *C* to a scalar value range. In object-ordered rendering schemes, this translates to blitting *M* quadrilles in image space, where each pixel's final value p_i is assigned a final colour according to equation 1:

$$p_i = \sum_{i=1}^{M}(C_i \cdot A_i) + \big(C_{i-1} \cdot (1 - A_i)\big);\ C_0 = (0|C_{user}) \tag{1}$$

In raycasting-based DVR, the each pixel's colour-opacity composition is split into 'classification' and 'accumulation'. In contrast to object-ordered rendering, extinction and absorption effects are taken into account during ray traversal. Hence, equations 2 and 3 being solved along a ray of length *t* (i.e. front-to-back composition):

$$C'_i = \sum_{i=0}^{t/\Delta t} C'_{i+1} + (1 - A'_{i+1}) \cdot C_i;\ C_t = (0|C_{user}) \tag{2}$$

$$A'_i = A'_{i+1} + (1 - A'_{i+1}) \cdot A_i \tag{3}$$

A-*C* mappings are commonly created by area definition within a scalar value-ordered histogram. Extensions have been proposed for using multidimensional transfer functions to easier delineate object interfaces[12], but these techniques introduce a cognitive overhead for the scalar interface definition. An issue arises if the range of scalar values includes fuzzy structure boundaries, as it is common with measured data of real objects (e.g. using CT or ultrasound) or statistical models.

For this purpose, statistical means of colour-opacity mapping definitions have been explored. Tenginaki et al. presented the notion of statistical signatures, where the scalar value range is plotted against the local-neighbourhood derivative values in a scatterplot to easier define structural interfaces[13]. Bruckner & Möller[14] proposed a transfer function-complementary technique called 'isosurface similarity maps', which is based on a mutual information framework. An automated, local approximation of scalar-valued statistics has been proposed by Heidacher et al.[15]. Also, momentum curves define mappings based on statistical momentum parameters of local voxel neighbourhoods[16]. Lindholm et al. studied the statistical distribution of data-specific material attributes to simplify the transfer function definition[17]. Furthermore, advantages[18] and distadvantages[19][20] of statistical transfer function definitions have been discussed in the literature.

3 Method

Our approach is technically based on a cuberilles implementation using modern OpenGL 4.0 technology. Cuberilles is chosen for its simplicity and ease of integration with established graphics pipelines. Major drawbacks of the technique are a lack of performance and the incorrect presentation of semi-transparent structures. The technique has a call-bound performance, as for each voxel 6 quads with 8 reoccurring vertices (in a regular lattice) are issued to the GPU, which creates a bottleneck for high-resolution datasets. Our approach resolves the issue by utilizing instanced rendering, which only issues 1 vertex per voxel to the GPU. These vertices are combined with attributes for controlling the visualisation in subsequent stages. On the GPU, the geometry shader is used to create the cubes and assign attributes and clipping values to each voxel, which also reduces the

memory transfer overhead over PCIe. In order to visualise semi-transparent structures coherently, the vertices are re-ordered with respect to their view distance on view-related render updates. Our experiments have shown that parallel sorting technology is mandatory in order preserve interactive framerates. The per-vertex view distance computation and the vertex array sorting are performed in parallel within this approach.

Reconstructing the cubes on the GPU during rendering allows using a larger collection of attributes for the visualisation without impeding the rendering performance. We define attributes for voxel visibility and gaussian statistics (i.e. normalisation, mean, standard deviation) to control the visualisation. The gaussian parameters are derived from user-selected scalar ranges. The final composition amounts to equations 4 to 6 (N_{gauss} being the number of bell curves available):

$$y_{gauss} = e^{-\frac{(x_{gauss}-\mu)^2}{2\sigma^2}} \tag{4}$$

$$A = \begin{cases} y_{gauss} & \leftarrow N_{gauss} > 0 \\ visibility & \leftarrow N_{gauss} = 0 \end{cases} \tag{5}$$

$$p_i = \begin{pmatrix} C_i \cdot \max(A, A_{min}) \\ A \cdot A_i \end{pmatrix} \tag{6}$$

The free parameters of visibility and gaussian attributes are set by statistical views that are directly linked to the volume visualisation (see case study 4b, Fig. 4). The visibility attribute is a combined expression, which is modified when users interact with the statistics as well as probing the volume. The colour attribute C_i is further modified by simple shading, based on a locally-evaluated phong illumination. The overall system further provides traditional sliced views and average intensity projections along the data's major axes, which are adapted in equal manners by statistical interaction.

The statistical views, being the major mean of interaction, are graph plots of histograms, intersection profiles and semivariogram graphs. In addition, a 3D probe calculates the voxel intersections along user-defined ray using fast ray-cube intersection methods[21]. The histogram allows the strict range selection of scalar values, expressed via the visibility attribute. Its related interaction stays fixed regardless of the data dimensionality (e.g. 3D volume, 2D slice or 1D ray). Further, the user can rebase the histogram on selection, resulting in a differential visualisation of the values with respect to the selection. The semivariogram is a measure of distance-based (*h*) value variability *y(h)* according to the following equation (see[22] for more details):

$$y(h) = \frac{1}{2} E\{[Z(x+h) - Z(x)]^2\} \tag{7}$$

, where *E* is the expectation value, *Z(x)* are the sampled points of the ray and *Z(h+x)* are the values separated by a given distance *h*. As a guideline, this sampling provides a smoothly increasing curve for gaussian-distributed values where close-by samples are interdependent and further-away samples show a larger variability. Semivariograms are here restricted to the 1D probing operation, because high-dimension semivariograms are difficult to understand and interact with. Another way of statistical interaction is the modelling of gaussian bell curve functions from the data. According to a selected value range, a polynomial solver is employed to

approximate the parameters of the respective bell curve. In this scheme, numerous curves can possibly overlap depending on the user-defined modelling purpose. The cuberilles volume visualisation adapts on selection to a picked curve approximation (as given in eq. 4), allowing for a fuzzy selection and visualisation of data.

4 Experimental Results

In general, the presented approach of statistical volume visualisation is well received across application disciplines as it offers tangible means to adapt the visual presentation. The cuberilles rendering provides a dense volume rendering, despite the visually inferior quality compared to DVR. The largest drawback of the solution is the interaction performance in high-resolution when navigating the data due to recomputing the distance sorting. The reason is the involved computational complexity: the distance calculation (O(n)) scales well on parallel systems while sorting (O(n·log(n))) has a weak scaling behaviour. Initial experiments on improving the performance using a GPU implementation of radix sort failed. Despite this, uniform-grid volumetric data are comparably simple and efficient to subsample for creating low-resolution previews during the interaction. The system hence offers the user multiple resolutions to operate with a given dataset.

a) Medical case study

The medical case study uses an open access CT dataset of a human male head. The dataset spans an area of 512 x 512 x 743 mm³, stored as a 256x256x166 cellular grid with a resolution of 2 mm x 2 mm x 4.4737 mm. The medical case study is used as a baseline for the usability assessment, because techniques such as histogram analysis, volume intensity profile and intensity-based visual manipulation are inspired by existing medical visualisation approaches. Fig. 2 shows the dataset with a locally evaluated phong illumination.

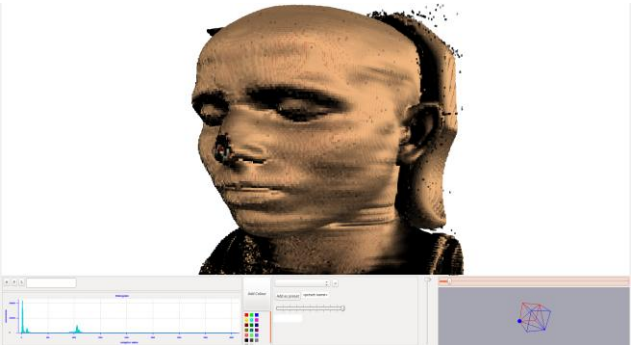

Fig. 2 Example rendering of a human head CT scan with enable phong illumination model evaluation within the presented prototype application.

The use of semivariograms is new to the medical domain as it is usually considered a geostatistical measure. By applying the new metric, we also observe

statistical interdependence of local cells (see green box, Fig. 3) in medical data. In CT samples, this can be caused by the gamma ray absorption of surrounding materials. Despite this, CT samples further away than the semivariogram's turning point (also called *lag*[22]) show an erratic variance response (see the semivariogram curve behaviour past the lag, Fig. 3). This is due to the structured, anatomical nature of the data, contrasting the statistical, often gaussian value distribution in volumetric geomodels. As illustrated in Fig. 3, the 1D ray intersection is highlighted while preserving its global context information.

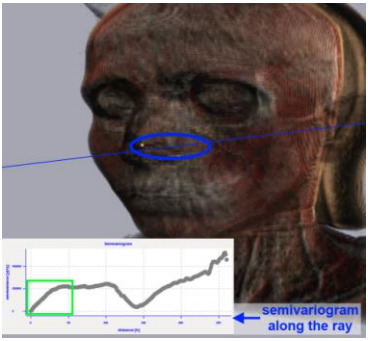

Fig. 3 Probing tool with 1D statistics, marking volume intersections (blue ellipse).

This context-preserving selection is also applied when value-filtering a histogram, allowing a better understanding of the context close to the filtered data (see Fig. 5). A more fuzzy selection is achieved by modelling the scalar values using gaussian curves (see Fig. 4). Furthermore, modelling multiple, contrasting curves allows for uncertainty on how to separate the volume structures.

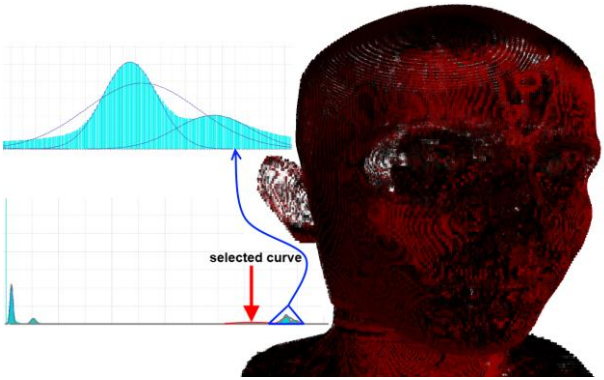

Fig. 4 Contextural selection using gaussian bell curves extracted from the histogram, also allowing overlapped curve models (top-left image). The selection of a curve highlights the relevant regions.

Furthermore, the illumination significantly adds to the structural and geometric perception of the data in comparison to simple value presentations. Fig. 5 shows a comparison between a coloured plane cutting-, a focus-and-context- and a lit

visualisation of the same data, each of which being used for different purposes.

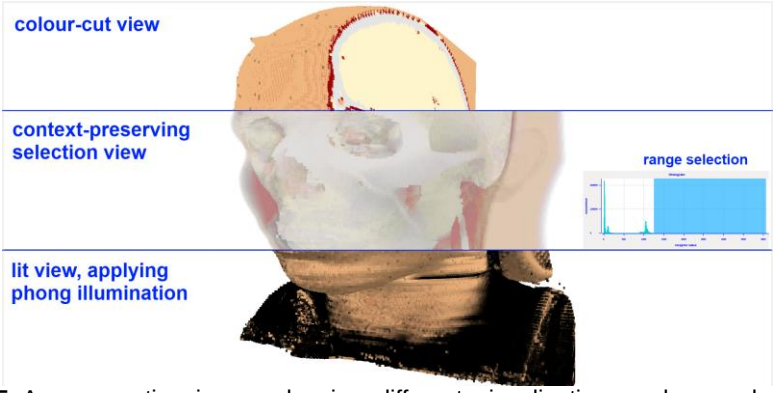

Fig. 5 A comparative image showing different visualisation modes, such as a colour-cut view, a selection view and the lit model view.

b) Fault facies case study of structural geology

The second case study emerges from structural geology, which studies earth crust deformations. Knowledge of rock deformations is essential for many applications (e.g. subsurface hydrocarbon flow and geothermal energy). A novel modelling method that produced the volume model under study is called 'fault facies'[23][24]. The model shown here represents a deformation band-dominated fault damage zone, within which various heterogeneous deformations occur. The model's volume is 100 x 1200 x 600 m^3 (voxel size of 1m x 5m x 5m; grid size of 100 x 240 x 120 elements). The scalar values indicated different degrees of deformation, ranging from 'no deformation' (0) to 'high deformation' (3). The degrees of deformation are modelled within Roxar RMS[1] using truncated gaussian simulation (TGS).

The average intensity projection, creating a view similar to common xray photographs, allows geologists to compare the model with small-scale, physical rock samples that can be scanned in the lab. The 3D visualisation allows quick adaptations of colour- and opacity mappings that highlight different aspects of the data. As visible in Fig. 6a, a standard mapping to fully-opaque greyscale intensity vales gives a coarse scalar overview. A deformation value mapping to opacity allows to visually skip under-deformed areas while highlighting flow-impairing accumulations of high deformations (Fig. 6b). A value mapping to complementary colours provides insight about the composition heterogeneity of the fault damage zone (Fig. 6c).

[1] Roxar IRAP RMS Suite – http://www.emersonprocess.com/Roxar/RMS

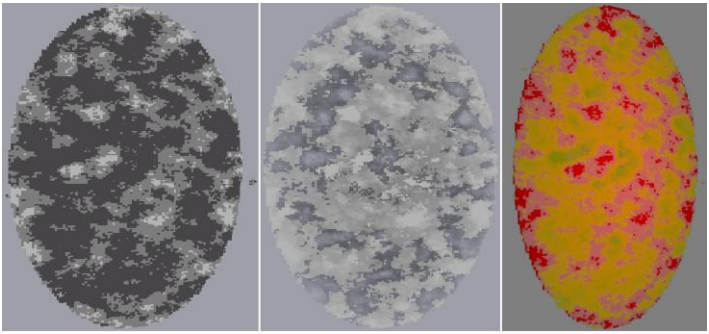

(a) (b) (c)

Fig. 6 Different colour-opacity mappings for deformations.

Fig. 7b illustrates the advantages of applying illumination to the volume rendering. Due to the strong gradients along deformation interfaces, the illumination makes the structure of rock accumulations and 'pockets' visible. Pockets are bounded volumes of low deformations. Together with the focus-and-context selective visualisation, the approach makes potential flow paths through deformed fault damage zones easily visible – an insight that is valuable but rarely discussed in established workflows, where such faults are modelled as impenetrable barriers to flow. This hypothesis is further validated by the semivariograms extracted from the 1D intersections (see Fig. 7a).

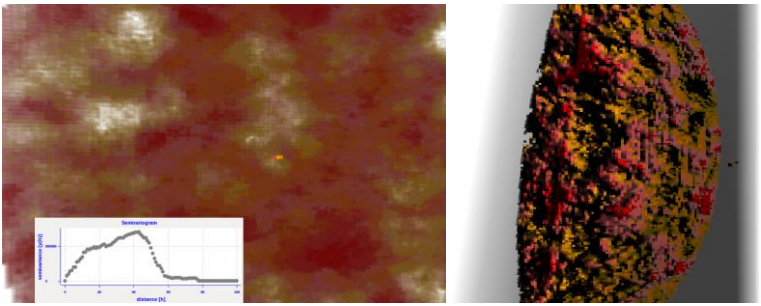

(a) (b)

Fig. 7 The combined visualisation of statistics (a) and illuminated voxels (b) shows deformation pockets in the data.

5 Discussion

This article introduces a new approach for exploring 3D volume datasets by using linked statistical views and a strongly-coupled interaction. The cuberilles volume visualisation with modern OpenGL 4.0 results in satisfying visualisations. Techniques such as focus-and-context, context-preserving filtering, gauss-curve selection, and the local illumination proved useful for exploring both medical and geological datasets. While the presented approach is only an auxiliary tool for

medical volume analysis, it is a large improvement in geology where enhanced volume visualisations are lacking in the available tools. The web-distribution and cross-platform applicability is advantageous for further, external evaluation.

Rendering drawbacks within this approach are connected to the cuberilles technique. A performance drop is observable when interacting with large datasets. While users currently need to select a data subsample, the approach can be extended to automatically select low-resolution models on viewpoint manipulations and to load high-resolution data when a viewpoint is established. This can also be realised within a DVR rendering approach. DVR also provides for higher presentation fidelity of the data.

The statistical approach can be extended for the focus-and-context visualisation and the intersection visualisation. Along an intersection pathline, a semivariogram is computed for which the distance-dependent variance can be used to limit the context volume. Voxels that are further away than the lag distance can be made fully transparent, resulting in a more visible intersection trace. The same use of semivariogram values holds true for a focus-and-context filtering to restrict the rendered context to a selection halo, clipped by the semivariogram lag. On the other hand, both techniques demand the calculation of 3D semivariograms.

From an application's point-of view, another case study from the domain of petroleum geology for hydrocarbon reservoir analogues has been skipped in this article. Generally, the tool's application is similar to the presented geological case. The voxel anisotropy in geological models results in very flat voxels, leading to issues in gradient-based illumination.

As a final remark, domain experts and practitioners from other geoscience domains are encouraged to evaluate the software on further case study.

6 Acknowledgements

This research is part of the VOM2MPS project (code 234111/E30), funded by RCN& FORCE through Petromaks 2. The fault facies model is part of fault facies project and property of UniResearch AS CIPR. The authors further thank Stefan Bruckner, Ivan Viola and Jan Tveranger for guidance and assistance. Roxar is acknowledged for providing an academic RMS license for the fault facies modelling.

References

[1] M. Levoy, "Display of surfaces from volume data," *IEEE CGA,* vol. 8, pp. 29-37, 1988.

[2] P. Sabella, "A Rendering Algorithm for Visualizing 3D Scalar Fields," in *SIGGRAPH*, 1988.

[3] D. Patel, C. Giertsen, J. Thurmond, J. Gjelberg and M. E. Gröller, "The seismic analyzer: Interpreting and Illustrating 2D Seismic Data," *IEEE TVCG,* vol. 14, no. 6, pp. 1571-1578, 2008.

[4] D. Patel, T. Langeland and V. Solteszova, "Geoillustrator - Fast Sketching of Geological Illustrations and Animations," in *EGU General Assembly Conference*, 2014.

[5] E. M. Lidal, H. Hauser and I. Viola, "Design principles for cutaway visualization of geological models," in *Proc. 28th Spring Conf. Comp. Graph.*, 2013.

[6] M. Natali, E. M. Lidal, J. Parulek, I. Viola and D. Patel, "Modeling terrains and subsurface

geology," in *Proc. EuroGraphics*, 2013.

[7] C. P. Botha and F. Post, "New technique for transfer function specification in direct volume rendering using real-time visual feedback," in *Proc. Medical Imaging*, 2002.

[8] G. T. Herman and H. K. Lui, "Three-dimensional display of human organs from computed tomography," *Comp. Graph. & Img. Proc.,* vol. 9, pp. 1-21, 1979.

[9] L.-S. Chen, G. T. Herman, A. R. Reynolds and J. K. Udupa, "Surface shading in the cuberille environment," *IEEE CGA,* p. 1985, 1985.

[10] C. P. Botha and F. Post, "Improved perspective visibility ordering for object-ordered volume rendering," *The Visual Computer,* vol. 21, pp. 887-896, 2005.

[11] W. E. Lorensen and H. E. Cline, "Marching Cubes: A high resolution 3D surface construction algorithm," in *SIGGRAPH*, 1987.

[12] J. Kniss, G. Kindlmann and C. Hansen, "Multidimensional transfer functions for interactive volume rendering," *IEEE TVCG,* vol. 8, pp. 270-285, 2002.

[13] S. Tenginaki, J. Lee and R. Machiraju, "Salient iso-surface detection with model-independent statistical signatures," in *Visualization*, 2001.

[14] S. Bruckner and T. Möller, "Isosurface Similarity Maps," in *Proc. Eurographics / IEEE - VGTC*, 2010.

[15] M. Haidacher, D. Patel, S. Bruckner, A. Kanitsar and M. Gröller, "Volume visualization based on statistical transfer-function spaces," in *PacificVis*, 2010.

[16] D. Patel, M. Haidacher, J.-P. Balabanian and E. M. Gröller, "Momentum Curves," in *PacificVis*, 2009.

[17] S. Lindholm, P. Ljung, C. Lundstrom, A. Persson and A. Ynnerman, "Spatial Conditioning of Transfer Functions Using Local Material Distributions," *IEEE TVCG,* vol. 16, pp. 1301-1310, 2010.

[18] H. Carr, D. Brian and D. Brian, "On Histograms and Isosurface Statistics," *IEEE TVCG,* vol. 12, pp. 1259-1266, 2006.

[19] M. Meyer, C. E. Scheidegger, J. M. Schreiner, B. Duffy, H. Carr and C. Silva, "Revisiting Histograms and Isosurface Statistics," *IEEE TVCG,* vol. 14, pp. 1659-1666, 2008.

[20] B. Duffy, H. Carr and T. Möller, "Integrating Isosurface Statistics and Histograms," *IEEE TVCG,* vol. 19, pp. 263-277, 2013.

[21] M. Eisemann, M. Magnor, T. Grosch and S. Müller, "Fast ray-axis-aligned bounding box overlap tests using ray slopes," *Graphics, GPU and Game Tools,* vol. 12, pp. 35-46, 2007.

[22] P. Ringrose and M. Bentley, Reservoir Model Design: A Practitioner's Guide, The Netherlands: Springer, 2014.

[23] A. Braathen, J. Tveranger, H. Fossen, T. Skar, N. Cardozo, S. E. Semshaug, E. Bastesen and E. Sverdrup, "Fault facies and its application to sandstone reservoirs," *AAPG Bulletin,* vol. 93, no. 7, pp. 891-917, 2009.

[24] D. Qu and J. Tveranger, "Incorporation of deformation band fault damage zones in reservoir models," *AAPG Bulletin,* vol. 100, pp. 423-443, 2016.

[25] H. D. Enge, S. J. Buckley, A. Rotevatn and J. A. Howell, "From outcrop to reservoir simulation model: Workflow and procedures," *Geosphere,* vol. 3, no. 6, pp. 469-490, 2007.