University of Groningen

Interactive obstruction-free lensing for volumetric data visualization

Traore, Michael; Hurter, Christophe; Telea, Alexandru

Published in:
IEEE Transactions on Visualization and Computer Graphics

DOI:
[10.1109/TVCG.2018.2864690](https://doi.org/10.1109/TVCG.2018.2864690)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2019

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Traore, M., Hurter, C., & Telea, A. (2019). Interactive obstruction-free lensing for volumetric data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1), 1029-1039.
<https://doi.org/10.1109/TVCG.2018.2864690>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Interactive obstruction-free lensing for volumetric data visualization

Michael Traoré, Christophe Hurter, and Alexandru Telea

Abstract—Occlusion is an issue in volumetric visualization as it prevents direct visualization of the region of interest. While many techniques such as transfer functions, volume segmentation or view distortion have been developed to address this, there is still room for improvement to better support the understanding of objects' vicinity. However, most existing Focus+Context fail to solve partial occlusion in datasets where the target and the occluder are very similar density-wise. For these reasons, we investigate a new technique which maintains the general structure of the investigated volumetric dataset while addressing occlusion issues. With our technique, the user interactively defines an area of interest where an occluded region or object is partially visible. Then our lens starts pushing at its border occluding objects, thus revealing hidden volumetric data. Next, the lens is modified with an extended field of view (fish-eye deformation) to better see the vicinity of the selected region. Finally, the user can freely explore the surroundings of the area under investigation within the lens. To provide real-time exploration, we implemented our lens using a GPU accelerated ray-casting framework to handle ray deformations, local lighting, and local viewpoint manipulation. We illustrate our technique with five application scenarios in baggage inspection, 3D fluid flow visualization, chest radiology, air traffic planning, and DTI fiber exploration.

Index Terms—Interaction techniques, focus + context, volume visualization, volume rendering, raycasting

1 INTRODUCTION

Direct volume rendering (DVR) is a pervasive visualization technique for displaying 3D scalar fields with applications in engineering, material sciences, and medical imaging sciences. However widely adopted, and able to handle large datasets at interactive rates, DVR inherently suffers from the problem of *occlusion*: Structures of interest located deep in the volume, called next *targets*, can be hard to spot and/or explore.

To aid with this, various techniques have been designed including transfer functions, segmentation, selection, and clipping. Yet, all such techniques have limitations. *Global* mechanisms, like transfer function editing, can remove both occluders and targets if these have similar densities. In certain applications, carefully designed transfer functions exist and should be used without (significant) modifications to facilitate understanding and user training [51]. *Local* mechanisms like segmentation, selection, or clipping are more effective in manipulating data confined to a given spatial region. Yet, many such mechanisms assume that one can easily and accurately select targets to remove them (occluders) or keep them (occluded). This is hard to do when *e.g.* one does not have direct access to the targets, or when significant 3D interaction is required to select occluder(s).

A different way to handle occlusion is to use *lenses*. These are flexible lightweight tools which enable local and temporary modifications of the DVR to reveal targets while keeping the global visualization context [7, 42, 43]. However, efficiently selecting the target and removing all in-between occluders is still challenging. More specifically, most existing occlusion management techniques do not simultaneously meet all following requirements: Rapidly create an unobstructed view of the target (R1), allow a flexible local exploration of the target zone (R2), keep the context in which the target is visually embedded (R3), and handle datasets where the target and occluders cannot be separated by transfer function manipulations (R4).

In this paper, we increase the flexibility of lenses for DVR exploration to jointly cover all above requirements. We propose a focus-and-context (F+C) lens that combines a distortion technique, which pushes aside the occluding objects, with a fish-eye field of view, to provide

a better perspective on targets. We specifically target the use-case of *partially occluded* objects, where the user has a glimpse of an interesting structure, buried deep in the data, but only slightly visible from a given viewpoint and transfer-function setting. We allow the user to 'open up' the volume without changing these settings, and reveal the target, by simple point, click, and scroll operations. Next, we provide several F+C modifications of the lighting parameters, transfer function, and geometry in the focus area to better understand the target. Our technique, implemented using a CUDA-based approach, can be easily incorporated in any generic DVR system.

The paper is structured as follows. Section 2 presents related work in occlusion management, lenses, and deformations for DVR visualization. Section 3 introduces the principle of our lens. Section 4 introduces implementation details. Section 5 presents five application scenarios for our lens in baggage inspection, 3D fluid flow visualization, chest radiology, air traffic planning, and DTI fiber exploration. Section 6 discusses our proposal. Finally, Section 7 concludes the paper.

2 RELATED WORK

To start with, let us refine requirements R1, ..., R4.

R1: The technique should rapidly create an unobstructed view of the target, *i.e.*, such a view should be created at interactive frame rates (10 fps or more) with no special pre-processing of the input volume (*e.g.*, segmentation), and with minimal user input (*e.g.*, using simple mouse and/or keyboard-modifier events). All above are needed to ensure that one can freely *explore* the volume by activating the lens anywhere with minimal effort and seeing its effect in real-time.

R2: Allowing a flexible exploration of the target zone means that one can manipulate the zone in the lens in various ways to see how the target is actually embedded in its surrounding context.

R3: Keeping the context means that the visualization around the lens does not change significantly from what would be shown there if the lens were not activated. This is needed to maintain the user's mental map before *vs* after activating the lens.

R4: The lens should enable the exploration of datasets where targets cannot be easily separated (isolated) from their surrounding context simply by manipulating parameters of the transfer functions (TF). One such issue is when targets and surrounding zones have similar densities; in this case, using a single global opacity TF could either render everything opaque (thus, it will be hard to visually isolate the target) or highly transparent (thus, the zone around the target will be transparent but so will be the target too).

Previous work on handling occlusion when exploring volume data can be divided into occlusion management techniques and lenses-and-deformation techniques. We next discuss these and also point out limitations from the perspective of R1, ..., R4.

• Michael Traoré and Christophe Hurter are with ENAC, French Civil Aviation University. E-mail: traore.s.michael@enac.fr; christophe.hurter@enac.fr.

• Alexandru Telea is with Institute Johan Bernoulli, University of Groningen. E-mail: a.c.telea@rug.nl.

Manuscript received 31 Mar. 2018; accepted 1 Aug. 2018.

Date of publication 16 Aug. 2018; date of current version 21 Oct. 2018.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2018.2864690

Fig. 1. (a-c) A baggage scan is viewed from different angles. In view (c), a suspicious sharp object is spotted between a set of mugs. (d-f) Filtering densities using a classical 1D opacity transfer function removes progressively more of the occluders (mugs), but also the target. (g) The user applies the lens on the target object (double-click). An animation starts opening the lens, rays are gathered to pass through occluders. Halfway through the animation, the object is magnified, but only the area close to the lens is visible. (h) The fish-eye field of view at the end of the animation scatters rays to fully show the target. (i) The lens is increased to magnify the target (mouse scroll).

2.1 Occlusion management

Many occlusion management approaches have been proposed [13]. Multiple viewports show the data from different perspectives [49]. This does not help when the target is strongly occluded from *all* possible viewpoints (R4). Virtual X-ray methods make targets visible by turning occluders (half-)transparent [6]. Kruger et al. [27] proposed ClearView, which interactively focuses on specific areas in the data while preserving context without visual clutter by modulating transparency. Correa and Ma [10] proposed visibility-driven transfer functions (TFs) to maximize the visibility of selected data-intervals. Designing good TFs is still challenging (R2) in general: For instance, in baggage inspection, a dissimulation strategy is to hide a threat among same-density objects, so TF editing cannot easily remove occluders but keep the target (R4) [46]. De-occluding a tumor from surrounding similar-density tissue in medical scans is a similar situation [35]. Rezk-Salama and Kolb [36] also considered the voxels' occurrence on the cast ray, besides their densities and positions. Hurter et al. [21, 22] removed occluders in 2D images, volumes, and trail-sets by deforming (pushing them away) in a focus area. Occluders are selected based on data value-ranges, thus we have here the same limitation as ClearView (R4). Li et al. [30] proposed luggage virtual unpacking where targets are cleared by interactively moving away occluders. This, however, alters the *context* (relative position, connectivity) where occluders occur (R3). Recently, an interactive volumetric data exploration via direct voxel manipulation was proposed [46]. Extending such approaches in a DVR setting to more complex deformations or changes of the data-in-focus is computationally challenging (R1).

2.2 Lenses and deformations

An interactive lens is a lightweight tool to solve localized visualization problems by temporarily altering a selected part of the visual data representation [42], and hence provides focus-and-context (F+C) solutions to volumetric data occlusion (R2). Parametrizable lens properties include position, shape, appearance, size, orientation, and selection of included data (focus). The lens *shape* is usually chosen to meet the requirements of the application and is linked to the lens function. Most lenses are circular [41] (a design we also choose), rectangular [26], or adapt their shape to the data-in-focus [34, 39]. The lens *position* and *size* can be changed manually by the user, or automatically to guide users toward interesting events in the data [40] or along a path of interesting events [1]. Our lens updates automatically its properties once a target has been selected. This allows a smooth transition towards an unobstructed and magnified area of interest.

Lenses for DVR face spatial selection and occlusion challenges. Magic Lens [48] addresses these by rendering occlusions with lower opacity and magnifying pre-computed features interactively or automatically in a pre-segmented dataset. This, however, does not provide an (interactive) way to deal with similar-density occluders and tar-

gets (R4) and does not propose local exploration of the target context (R2). GlyphLens [45] removes occluding glyphs by pulling them aside through animation. However, this covers only glyph-based and not DVR visualizations.

Lenses can create discontinuities between their inside and outside areas. Deformations can be a solution to this. Hsu et al. [17] create flexible deformations by non-linearly sampled rays to smoothly project objects at multiple levels of detail. However, this is computationally expensive and far from real-time (R1). Exploded views are used to partition a volume into several segments [5, 37]. This strongly reduces occlusion but distorts the context (R3) and requires some sort of data pre-segmentation which takes time to compute (R1). Correa et al. [8, 9] allow one to manipulate the geometry of a data object. McGuffin et al. [31] performed deformations using peeling to see hidden parts of the data. Such techniques remove potentially important contextual information surrounding the target, which makes it hard to see how the target is embedded in its context (R2).

Deformations can reveal specific data features by using a precomputed segmentation. Tong et al. proposed a deforming lens which moves streamlines to observe the inner part of streamline bundles [44]. Other techniques performed deformations using surgical metaphors [8, 24] to show hidden parts of a volume. Such techniques do not offer tools for local manipulation of the viewpoint that allows seeing a target under multiple perspectives (R2) while keeping the global context.

Using non-straight-line rays is another way to avoid occluders. Cui et al. [11] propose curved (Bézier) rays to support this. Wu et al. [50] refine this further to create multi-perspective views. These approaches require a quite careful ray-path planning in advance, so they are not directly aimed at a lens effect (R2). Also, the target should be accessible from the viewpoint via a (possibly curved) path (R4). Also, from the presented examples, although two examples of DVR models are given, it seems these techniques mainly address de-occlusion in large 3D polygonal scenes such as terrain and city models.

Table 1 summarizes the main advantages and limitations of a set of volumetric exploration techniques selected from the ones mentioned above which are close to our proposal. As visible, none of the techniques scores high on all requirements. The table also lists the number of different datasets and/or use-cases these techniques were tested on. We will validate our proposal on a similar number of use-cases (Sec. 5).

2.3 Detailed contributions

Summarizing the above discussion on the requirements and related work on occlusion management, we propose a new technique which combines high-quality DVR with a fast, versatile, and easy to use, lens to support the interactive exploration of occluded data in volumes. In the classification of view deformations by Carpendale et al. [7], we use a nonlinear radial distortion through an interactive lens to remove

Technique	R1	R2	R3	R4	Use-cases
McGuffin <i>et al.</i> [31]	●	○	●	●	1
Sonnet <i>et al.</i> [37]	●	○	●	●	5
Wang <i>et al.</i> [48]	●	○	●	●	6
Bruckner <i>et al.</i> [5]	●	○	○	●	6
Correa <i>et al.</i> [8,9]	●	○	●	●	8
Cui <i>et al.</i> [11]	●	○	●	○	5
Hsu <i>et al.</i> [17]	○	●	●	●	4
Hurter <i>et al.</i> [21]	○	●	●	○	6
Wu <i>et al.</i> [50]	●	○	●	○	5

Table 1. Related work selection vs requirements R1..R4 (●: good, ○: average, ○: limited) and the number of use-cases (datasets) used in the respective studies to demonstrate these methods.

occluding items and keep the global context while magnifying a partially occluded item. Related to volumetric lens techniques, we frame our contribution as follows: We propose an interactive deforming lens that magnifies and pushes aside occluding objects located in front of a designated focal point which meets the four requirements; the combination of flexible and real-time interactive changing of the focal point, custom bent rays used for DVR, lens deformation, and shading and transfer function in the focus area allow us to provide *on the fly* a range of perspectives of the targets, without having to change the viewpoint or manipulate complex parameters in multiple linked views.

3 PRINCIPLE

Consider the typical DVR algorithm: Given a scalar volume $V \subset \mathbf{R}^3 \rightarrow \mathbf{R}$, each pixel $\mathbf{x} \in I$ in the DVR image $I \subset \mathbf{R}^2$ thereof corresponds to the compositing of sampled data along a ray that passes through V and ends at $\mathbf{x}$. In classical DVR (Fig. 2-a), such rays are defined by the eye position $\mathbf{e}$ and a ray direction unit vector $\mathbf{d} = (\mathbf{x} - \mathbf{e}) / \|\mathbf{x} - \mathbf{e}\|$. Consider now a focus point $\mathbf{f} \in I$ (the lens center) and a lens radius $R > 0$. We modify all rays passing through the lens (focus/0 area $D = \{\mathbf{x} \in I \mid \|\mathbf{x} - \mathbf{f}\| \leq R\}$ in order to de-occlude, magnify, and emphasize a target object. Our ray behavior is divided into three steps: (1) Provide a clear view of the target by moving closer to it and by pushing occluders aside. (2) Set a wide field-of-view (fisheye) to better see the target. (3) Modify the parameters of the lens, lighting, and opacity TF in real time to better explore the target. These steps are detailed next.

3.1 Creating an unobstructed view

The scenario we address is as follows: Given a volume V , users produce a DVR thereof, using whatever suitable TFs and other parameters are applicable. When examining V from various viewpoints, (at least) one viewpoint $(\mathbf{e}, \mathbf{d})$ is found from which some intriguing structure is *partially* visible in I . We call this structure the *target*. Users next want to quickly and easily unravel the target. For this, we proceed as follows: We first *gather* all rays passing through the lens pixels (focus area D) to follow the lens' axis vector $\mathbf{a} = (\mathbf{f} - \mathbf{e}) / \|\mathbf{f} - \mathbf{e}\|$. As explained above, at the location $\mathbf{f}$ of the lens center, we do see an interesting partially occluded target. Hence, by definition, the gathered rays pass *through* occluders to hit this target, otherwise we would not see it. We control gathering by setting the ray direction passing through $\mathbf{x} \in D$ to

$$\mathbf{r}(\mathbf{x}) = (1 - \alpha)\mathbf{a} + \alpha\mathbf{d}, \quad (1)$$

with $\alpha \in [0, 1]$. When $\alpha = 0$ (default), all rays follow the lens axis $\mathbf{a}$, thus, can best pass through obstacles. When $\alpha = 1$, rays follow a straight path. Changing α with the mouse wheel smoothly navigates between the lens effect, *i.e.* opening up a 'hole' in the volume to see the target, and classical DVR.

3.2 Setting a wide field of view

Once the rays pass obstacles (Sec. 3.1), we want to *scatter* them so as to best sample the target. Consider that this target is at some depth $t_{\text{target}} > 0$ within V . After the rays pass the occluders, but before they hit the target, *i.e.*, travel past a distance $t_{\text{min}} < t_{\text{target}}$ through V , we deflect (scatter) them so as to best sample the target. For this, we set

the parametric position of a ray point to

$$\mathbf{p}(\mathbf{x}, t) = \mathbf{r}(\mathbf{x})t + \beta(\mathbf{x} - \mathbf{f})(t - t_{\text{min}}) \quad (2)$$

for any pixel $\mathbf{x} \in D$ and any $t \geq t_{\text{min}}$. Here, $\beta \geq 0$, adjusted via the mouse scroll wheel while pressing the Shift key (Fig. 2-c), controls the ray scattering: Small values magnify a small volume area close to the ray $\mathbf{r}(\mathbf{x})$; larger values sample more of the volume behind the lens. Intuitively, this is as if we moved a magnifying lens to a depth t_{min} inside V . Summarizing, after the user finds an interesting but partially occluded target using *standard* DVR, our lens squeezes rays to pass between occluders and next fans them out to explore the target.

3.3 Interactive exploration of the target

To achieve a more effective exploration, we can interactively modify several parameters of the DVR and the lens, as follows.

Lens radius: The radius R , telling how big is the 'hole' to open up in the volume to see the target, is set via the mouse wheel. The parameters α and β (affecting the gathering and scattering of rays respectively) are set by the mouse wheel and modifier keys. The value t_{min} (depth from which scattering starts) is set using the arrow keys.

Lens axis: Users can rotate the lens axis $\mathbf{a}$ using a virtual trackball activated by the right mouse button. Changing $\mathbf{a}$ effectively samples the target from many viewpoints, allowing the user to look 'around' it to see parts which are not visible from the current viewpoint, but *without* actually changing the viewpoint. This is of high added value, since changing the viewpoint can bring us to a view where the target is fully invisible, so we do not know where precisely to activate the lens anymore. Figure 4 shows three such local rotations for the baggage dataset introduced in Fig. 1. From these, we see that the star-shaped target is relatively thick.

Lighting: We modify the volumetric Phong lighting parameters to better explore the target, as follows. Let $\mathbf{c} = \mathbf{e} + t_{\text{min}}\mathbf{a}$ be a point at depth t_{min} along the lens axis, and let $B(\mathbf{c}, R)$ be a sphere of radius R around this point (Fig. 2b). We call voxels in this sphere 'in focus', and all other voxels in V 'out of focus'. Let ϕ be the specular term coefficient, set to a high value (default: one). First, for all voxels $\mathbf{x} \in B(\mathbf{c}, R)$, we use a specular coefficient $\phi(\mathbf{x}) = \phi(1 - d)$, where $d = \|\mathbf{x} - \mathbf{c}\|/R$. For all voxels outside $B(\mathbf{c}, R)$, we use $\phi(\mathbf{x}) = 0$. Hence, voxels close to the focus point $\mathbf{c}$ appear highly specular; further away from $\mathbf{c}$, voxels get less specular, and voxels out of focus are purely diffuse. Secondly, we allow the user to locally rotate the light vector using the same trackball mechanism as for the lens axis rotation. Let $\mathbf{l}^{\text{lens}}$ be this vector, and let $\mathbf{l}^{\text{global}}$ be the global light vector used by standard DVR. For all voxels in focus, we use a light vector $\mathbf{l}(\mathbf{x}) = (1 - d)\mathbf{l}^{\text{lens}} + d\mathbf{l}^{\text{global}}$. As the user rotates $\mathbf{l}^{\text{lens}}$, the light direction will visibly change in the middle of the lens, stay constant outside it, and smoothly change in between.

The above two mechanisms combined yield the effect of a moving flashlight turning around a shiny target embedded in a constantly-lit diffuse scene. Figure 5 shows these mechanisms for a chest CT dataset containing a deeply buried tumor (the dataset and use-case are described in detail in Sec. 5.3). We see how turning the light highlights small-scale details on the target surface (tumor) without changing the viewpoint or lens location. Moreover, the high specularity in the lens attracts the user's attention to this area; the diffuse lighting outside the lens put less emphasis on the context area.

Opacity: We modify the opacity transfer function along a similar idea as for lighting (Fig. 3). Let $TF_o^{\text{global}} : \mathbf{R} \rightarrow [0, 1]$ be the user-chosen opacity function used globally for the volume. Let Γ be a Gaussian pulse of unit height centered at the average density value $\bar{\rho}$ in $B(\mathbf{c}, R)$ and with standard deviation σ . We estimate $\bar{\rho}$ and σ by considering the density ρ at 150 points randomly sampled inside $B(\mathbf{c}, R)$. Higher values for the sample count yield visually very similar results for our tested volumes of up to 500^3 voxels, while requiring (slightly) more computation costs. Then, for voxels in $B(\mathbf{c}, R)$, we use an effective opacity transfer function $TF_o = TF_o^{\text{global}} + (1 - d)\Gamma$. For voxels outside B , we use TF_o^{global} (standard DVR). This is useful when the user sets

Fig. 2. Obstruction-free lens working. A target is (mostly) hidden by occluders in front of it. (a) Classic DVR shows a small part of the target. (b) Our lens gathers rays to avoid occluders (Sec. 3.1). Once close to the target, rays follow again their initial paths. Yet, only a small part of the target is visible. (c) Scattering rays makes the full target visible (Sec. 3.2). Finally, we adjust the local viewing and lighting directions $\mathbf{a}$, $\mathbf{l}^{lens}$ (Sec. 3.3).

TF_o^{global} to make most out-of-lens voxels relatively transparent. In that case, TF_o will still make voxels in B opaque, thus allowing to see the in-focus structures better. Figure 5 has been generated this way.

Fig. 3. Construction of local transfer function TF_o . See Sec. 3.3.

3.4 Smooth transitions

If we bend rays passing through the lens pixels D (Eqns. 1 and 2) and trace rays starting at pixels in $I \setminus D$ as straight lines, discontinuities appear at the lens borders. We solve this as follows. Let $\mathbf{p}(\mathbf{x}, t)$ be the voxels along a lens ray starting at screen pixel $\mathbf{x}$, as computed by Eqn. 2. Let $\mathbf{p}^{line}(\mathbf{x}, t)$ be the voxels along a straight-line ray starting at $\mathbf{x}$, i.e., computed using $\alpha = 1$ and $\beta = 0$ in Eqns. 1 and 2 respectively. For every value t along every such ray, we compute the interpolated ray $\bar{\mathbf{p}}(\mathbf{x}, t) = (1 - f(d))\mathbf{p}(\mathbf{x}, t) + f(d)\mathbf{p}^{line}(\mathbf{x}, t)$, where d is the distance of $\mathbf{x}$ to the lens axis (normalized to unit by dividing it by R) and $f: [0, 1] \rightarrow [0, 1]$ is an interpolation function. Next, we use the rays $\bar{\mathbf{p}}(\mathbf{x}, t)$ to compute the DVR by standard composition. This way, rays effectively vary smoothly from their bent versions (close to the lens axis) to straight lines (outside the lens). Setting $f(d) = d^2$ keeps the interpolation transitions close to the lens border, so most of the lens is dedicated to show the desired fisheye effect.

Separately, we use a slow-in/slow-out animation [12] to introduce the lens effect. When activating the lens, we vary α and β from their defaults ($\alpha = 1$, $\beta = 0$, i.e. straight-line classical DVR) to their actual user-set values, compute the volume rendering on-the-fly, and display the resulting images. The effect resembles gradually opening a hole in the volume – see the associated video. The speed increase at the start of the animation helps one to quickly see what is revealed in the lens; the decreasing speed at the end helps seeing where the pushed-away occluders actually go. This also gives some semantic to the moving shapes, allowing one to interpret the motion as a magnification of a target, and to keep the focus on visual entities during this transition. When deactivating the lens, we play back the animation in the opposite sense, which suggests closing the opened hole in the volume.

4 IMPLEMENTATION

We implemented our occlusion-free lens by modifying a standard DVR ray caster, publicly available in NVIDIA CUDA's SDK [33]. We modified this ray caster to incorporate the new ray definition (Eqns. 1 and 2), the lens effect, and the local per-voxel Phong lighting parameters,

Fig. 4. Performing local rotations in the lens improves visibility of the shape and the thickness of the partially occluded target object (ninja star).

all controlled via keyboard and mouse. On a PC with 16 GB RAM and a GeForce GTX TITAN X card, we achieve 15 frames per second for volumes up to 512^3 voxels at a 1900×1200 pixels screen resolution. All in all, adding our lens to an existing ray caster should pose no significant implementation problems.

5 APPLICATION SCENARIOS

We next demonstrate our obstruction-free lens via five use-cases considering scalar density volumes from baggage inspection, 3D flow simulation, radiology, air traffic planning, and diffusion tensor imaging.

5.1 Baggage inspection: An unusual blunt object

In airports, security agents deal with volumetric data exploration during baggage inspections. While automatic systems can detect densities of harmful substances such as C-4, TNT, and nitroglycerin, or prohibited articles (threats) like classical firearms and knives, unusual threats are hard to find. Four main concealment strategies exist [46]:

Superposition: A threat may be sheltered among dense materials. While possible to see through such a ‘shield’ using high penetration (enhanced X-ray power) or image processing (contrast improvement), such techniques are not universally available and also require fine-tuning many parameters, which slows down inspection.

Location: Objects located in the corners, edges, or in the luggage’s frame are very hard to spot.

Dissociation: One can conceal a threat by spreading its parts in the luggage, e.g. by disassembling a weapon and scattering its parts.

Lure: A minor threat (lure) like small scissors is clearly visible and catch the security agent’s attention who can miss the real threat.

Fig. 5. Changing lighting parameters in the lens. (a) Constant global specular coefficient. (b) Specular coefficient high in the lens and low outside. (c-f) Changing the in-lens light vector yields the effect of a flashlight rotating around the target. The ball icons illustrate the local light vector direction.

Baggage labeled as suspicious by human inspection or automated scan heuristics must be checked by human agents. Besides time-consuming physical unpacking, one can use ‘virtual unpacking’ tools that segment the 3D scan by a density-based confidence measure and next move the segmented objects away by animation to reduce occlusion [30]. Such systems have been patented and used in production [29]. However, when the automatic segmentation is not optimal, the user must manually change its parameters, repeat the segmentation and animation, which goes back to being time-consuming.

Consider the baggage scan in Fig. 1 ($283 \times 189 \times 344$ voxels, dataset obtained from an actual airport scan). Automatic baggage inspection systems will not detect anything suspect here. However, while visually exploring this baggage from different angles (Fig. 1a-c), we see an object hidden between a set of mugs. To reduce occlusion, a common solution in baggage inspection is to filter materials by density in order to show or hide subsets of the volume. However, for our dataset, the suspect target has almost the same density as the surrounding mugs, so removing the latter also removes the target (Fig. 1d-f). Using the obstruction-free fish-eye lens helps here: Clicking on the sharp detail visible in Fig. 1c first gathers rays so they pass through the low-density zone between the mugs (Fig. 1f). The animation that opens the lens (Fig. 1e-g) reveals an unobstructed view of the target. However, this shows only a small part of the target. Scattering rays next fully reveals the target (Fig. 1h). Adjusting the lens size shows a more detailed view of the target (Fig. 1i). Next, locally turning the viewpoint around the target (Fig. 4) allows the agent to decide that the target is a shuriken (Japanese ninja star weapon). Since the object is very thick and blunt (see Fig. 4), it is not an actual weapon, thus not a threat.

We evaluated our lens for this use-case by a user study. Eight airport security specialists were recruited (ages 23 to 43; experience in baggage scanning 8 months to 20 years, average familiarity with 3D tools from no experience to expert, see Fig. 6). All attended a 20-minute global demo of the lens operation. Next, they were given each a personal training session for using the tool (5 minutes), in which they were instructed on the mouse and keyboard controls. After this, they were asked to work in pairs to examine the above-mentioned baggage CT dataset to form a decision on the nature of the ninja star possible threat (20 minutes of tool usage per person, after which the pair was changed). The idea behind this is that one person operates the tool while the other poses questions or suggest explorations, much like typical airport security operators work with a scanner. In the end, they all separately filled in a web questionnaire covering several questions and also provided open feedback (questionnaire available in the supplementary material).

Figure 6 shows the answers. The first question-set (S1) regarded how easy-to-use, generally effective, and effective *vs* other known tools our lens is for *untargeted* inspection, *i.e.*, when no suspect target is partially visible. Here and next, other tools denote classical 2D X-ray or 3D CT scans used in baggage scanning that the subjects

know. As Fig. 6c-e shows, the answers (on a 5-point Likert scale) were predominantly positive: The tool is easy to use, is useful, and is actually more useful than known tools for untargeted exploration. The second question-set (S2) regarded how good our tool is to examine *specific* targets which are partially visible. Here again, the answers were predominantly positive (Fig. 6f-h). The main appreciated features of our tool are listed in Fig. 6i.

We next summarize the received open feedback. According to the subjects, our tool can provide them a better perception of the items inside the baggage as compared to the classical 2D single-viewpoint X-ray machinery they routinely use. Quotes from the open feedback: “clear added-value compared to all systems I know”; “this tool is a real gain for examining luggage with uniform and/or high densities”; “definitely better than known tools for examining threats I am not familiar with / I have not seen before”. However, our tool should not be used for the typical carry-on baggage inspection which has a very small allowed inspection time (15 to 20 seconds). Our tool is much more interesting for inspecting checked-in baggage, where inspection time-windows are up to 3 minutes. The perceived added value for this use-case is also higher: Opening up checked-in baggage for manual inspection is much more complicated and time-consuming than for carry-on baggage. Moreover, the only system for inspecting checked-in baggage that the subjects knew of is a scanner that aims to *automatically* detect threats via X-ray imagery; this system suffers from false positives, so a manual examination tool like ours could quickly eliminate such false positives, and thus the delays of opening up checked-in baggage. Finally, several subjects suggested that adding a function to display a classical 2D slice view (activated by a key press and aligned with the focal point) would be useful since this would show additional detail.

5.2 Fluid flow: A deep-buried spherical vortex

Flow visualization using streamlines has a long history [3, 32]. For 3D datasets, a key challenge is to balance the streamline density. Low values allow seeing inner regions in the data but can subsample (miss) patterns. High values show more data but create too much occlusion. We next show how our lens can be used to alleviate problems in the latter case. The dataset [16] captures the simulation of water flow in a basin computed on a grid of $128 \times 85 \times 42$ cells using 4595 streamlines with 183K sample points traced by pseudo-random seeding. We convert this set of 3D curves (polylines) to a scalar volume by using GPU-accelerated kernel density estimation (KDE) [28]. Similar techniques have been used to compute density maps of 2D trail-sets [18, 20, 47].

We first explore the density volume (500^3 voxels) using standard DVR (Fig. 7). Note that, given KDE’s smoothing effect, streamlines appear as finite-thickness tubes rather than pixel-thin curves. After turning the viewpoint a bit, we notice a dense spherical item deep in the data (Fig. 7a). To see its shape better, we increase opacity; however, this immediately increases occlusion so the item becomes invisible.

Fig. 6. Evaluation of lens-based baggage inspection (Sec. 5.1).

Conversely, decreasing opacity to reduce occlusion makes the item almost transparent. Our lens solves the problem: In the initial view (Fig. 7a), we point at the target and turn on the lens. This pushes away the occluding stream bundles, and shows that our item is a set of densely-packed, low-speed, tightly-turning streamlines that create a ball-like vortex (Fig. 7b). To make sure our target is spherical, we view it in the lens from different directions, by interactively changing the ray directions in the lens (Fig. 7c). Finally, we can close the lens but keep the target magnified (Fig. 7d). Finding the details of this vortex cannot be done using standard DVR. Interestingly, this vortex has also not been discovered by any of the visualization techniques that used this dataset (according to our knowledge) [15, 16, 28, 38].

5.3 Chest scan: A hard to see tumor

In our third use-case, we consider a contrast chest CT scan ($512 \times 512 \times 110$ voxels) of an elderly patient with a sizeable lung tumor. The tumor was detected in a CT scan performed after the patient reported acute chest pain. Typical examination of these scans by the pulmonologist and radiologist in charge involves slice-based views. Figures 8a-c and 8d-f show two such slice sets (axial, coronal, and sagittal views), produced using typical lung, respectively mediastinal, contrast presets. Although the tumor is visible in all these views, its exact shape, morphology, and connection to the lung walls are hard to assess. Finding such details on the tumor is essential, explained the doctors in charge, to determine the TNM score [4] and also plan treatment. Using standard DVR makes the tumor and its 3D position partially visible (Fig. 8f). Yet, occlusion from the rib cage and other tissues is still present. Using both TF presets and manually changing the TFs in the 3D Slicer tool [25] used to create the DVR could not help de-occluding the tumor without making it partly transparent. The slice images in Fig. 8a-f confirm this by showing that the gray values for the tumor and surrounding skin-and-muscle tissue are very similar. This is due to the fact that the tumor had grown rapidly and started necrotizing, which filled it with fluids, making its density very similar to that of the obstructing (skin and muscle) tissue, explained the pulmonologist. Hence, one cannot remove such occluding tissue in a classical DVR setting by opacity TF manipulation without also removing the tumor. This makes examining this specific tumor harder than for regular cases.

We next used our lens to examine the tumor. Fig. 5 shows several sample snapshots. We see that the tumor is significantly more visible when using the lens than when using standard DVR (Fig. 8d), both in terms of removing the occluding tissue and in terms of the tumor's opacity – compare the inset in Fig. 8d with the images in Fig. 5. Secondly, relighting the tumor from various directions allows one to see small-scale morphological details such as the tumor's surface shape and its connection via protuberances and veins with the lung walls.

We asked the two medical specialists (pulmonologist and radiologist) in charge to state the potential advantages and/or limitations of our lens as compared to standard slicing and DVR techniques, after a 20-minute usage of the tool. Both specialists have over 10 years of medical experience in treating lung cancer, and routinely use several slicing and DVR tools. They work in a private hospital in Belgium and are not actively associated with medical imaging research. Our identities were hidden from them during the lens evaluation. The provided input can be summarized as follows: The occlusion-free lens is definitely easier and faster to use than classical DVR and/or slicing techniques. It is especially more effective than these to get a quick, first impression of a deep buried anatomical detail. Changing the lens' parameters by direct interaction is as simple as changing window/level functions in a typical slice-based tool, and is definitely simpler than tuning typical DVR parameters to obtain similar results. This 'entices' the user to explore, which is a good aspect. The fact that the lens minimizes viewpoint change (volume rotation), *i.e.*, after a suitable viewpoint was found from which a (small) part of the target is visible, one doesn't need to change this viewpoint, is a strong feature, as 3D viewpoint changes are disruptive and cost time. This is important in a cost-aware environment where specialists have very limited time (about 20 minutes) to assess a CT scan. However, the lens should not *replace* classical slice-based exploration, which shows small-scale details better. In the context of

Fig. 7. Flow volume exploration with two different opacity transfer functions (top and bottom rows). In viewpoint (a), we notice a small high-density spherical item. (b) We apply the lens at that location (double click). (c) The directions of rays in the lens are changed to see the whole target in the lens (right click + mouse drag change direction). (d) The lens is gradually closed while keeping the focus area magnified (shift + scroll).

the current dataset (Fig. 8), the lens was useful to both confirm the TNM score (T3 grade tumor, 6.5 cm in size) found via the 2D slices, but much more so for understanding how and where the tumor is connected to surrounding tissue, which is very hard to do using only 2D slices.

5.4 Aircraft trajectories: Outliers in the French sky

We next consider a task from air traffic planning – detecting and studying outliers in large-scale datasets containing tens of thousands of 3D (latitude, longitude, height) trails of aircraft over a given spatio-temporal region [19]. Such datasets are typically viewed using 2D (latitude, longitude) plots where opacity encodes the spatial density of flights – see Fig. 9a, which shows one day of recorded aircraft trajectories over the French airspace. Fig. 9(b) shows a detail zoom-in, where we can see an abnormal – that is, far from straight or slightly curved – aircraft trail: A tanker aircraft performed an eight-shaped loop as it was waiting to refuel other aircraft. Revealing such patterns using 2D techniques, e.g. [23], is very hard. In particular, it is hard to de-occlude these patterns from the overall context of criss-crossing aircraft trails, even when one knows their 2D spatial location.

Our lens can help with this task, as follows. We first convert the set of 3D trails to a 500^3 density volume, using KDE as for the streamline use-case (Sec. 5.2). Examining this volume via standard DVR shows an outlier trail at some point in space, see curved patterns in Fig. 10a. Activating the lens on this area and interactively tuning the target depth t_{min} (since we don't know the trail's height) brings the outlier trajectory in focus and pushes away occluding trails (Fig. 10a). Like in the other examples presented so far, we can quickly change the magnification factor and view direction to better study this trail in context (Fig. 10b-d). From these images, we easily see that the outlier trail has an eight shape. Revealing this outlier trail using standard 2D visualization techniques [23] costs several minutes. Doing the same using our lens costs under one minute. Also, comparing Figs. 9b and 10b-d, we argue that the eight-shape of the outlier trail is much more prominent, and thus recognizable, in the latter images (made using our lens) than in the former ones. Last but not least, the 3D DVR approach that our lens enhances explicitly encodes flight height information, so our lens can use it by interactively tuning the depth value t_{min} where the lens is focused. This cannot be done with 2D techniques which ignore the depth dimension.

We validated our findings with an air traffic data scientist with more than 10 years of experience in air traffic control and planning. She confirmed that this specific eight-shape trail in Fig. 9(b) is an actual

aircraft which performed waiting loops and acted as a fuel supplier for military aircraft. Other comments included the following: Compared to standard 2D visualization techniques, our tool makes detecting outliers easy since there is no need for complex manipulation to reveal such outlier trails. Also, the user does not have to deal with color and alpha mapping parameter-tuning to make specific outliers emerge. Separately, trail visualization easily creates many occlusions leading to either fully opaque areas or too much local overlap, which both hinder seeing and examining specific trails. Our lens does help such cases by distorting the space to locally remove such occlusions. All in all, in the studied dataset (Fig. 9), the lens was specifically useful since, for high transparency, one would not detect the outlier trail, while for low transparency, one would get a hint of the outlier's existence, but not see it in detail due to too much occlusion; the lens allows using low transparency, but removes the clutter caused by it to reveal the outlier.

5.5 Brain fibers: Uncluttering the bridge

Our last use-case considers the exploration of fiber tracts visualized as streamlines of the major eigenvector of a diffusion tensor imaging (DTI) field. Such datasets have a spatially complex structure which makes them hard to explore [2]. In particular, fiber tracts are spread volumetrically over the entire extent of the brain, and create tangled patterns inside which it is hard to see much. DVR techniques are often used to render such tracts, one of the advantages being that close fibers get visually 'merged' to reveal spatially coherent structures, an effect which is not possible when fibers are rendered as polylines. However, DVR methods also create more occlusion, thus difficulties in seeing structures deep within the volume.

We consider an $128 \times 128 \times 51$ DTI volume (same dataset as in [14]). We traced 150352 fibers seeded in, and going over, regions of high fractional anisotropy in this volume. We filtered out fibers shorter than 2mm, yielding a total of 120593 fibers to display (6.4M sample points). Next, we converted this fiber-set to a 512^3 density volume, using KDE with a 3D isotropic kernel of radius 15 voxels, like for the streamline use-case (Sec. 5.2). Figure 11a shows the result, rendered with DVR, with opacity function mapping the fiber density. While terminal fibers are well visible, we cannot see anything inside the volume. Activating the lens in the middle of the volume opens a hole through which a small part of the *corpus callosum*, the fiber bundle wrapping the bridge that connects the two hemispheres, becomes visible. By slightly decreasing opacity (Fig. 11c), the *corpus callosum* gets clearly visible, appearing as a compact structure, due to the KDE blending of neighbor fibers.

Fig. 8. Lung tumor visualization using slices (a-c) and standard DVR (d). Annotations are manually added by the examiner to delineate the tumor location. Images constructed using the 3D Slicer tool [25].

Obtaining such a view of the *corpus callosum* only using DVR would be very hard, since transfer functions would either render separated (non-merged) fibers, or else make the fibers surrounding the structure of interest too thick and occluding.

This scenario has the main difference compared to all previous ones. In all earlier cases, the standard DVR of the data (that is, without the lens) showed us a partial small cue of the structure of interest within the volume, and we used the screen-space location of this structure as the focus point where to activate the lens. In this last scenario, there is no point in the original DVR image (Fig. 11a) from which the *corpus callosum* is even partially visible, due to the high opacity given by the used transfer function. Hence, the user can activate the lens at *any* desired point to peek inside, and towards the center of, the volume. Given the nature of the data, the structure of interest is quite easily visible from most such viewpoints (see lens inset in Fig. 11b). Once its presence is revealed, the user can next adjust the viewpoint and/or the opacity transfer function to get an optimal view on the target, such as the one shown in Fig. 11c. Summarizing, we can use our lens also in cases when no partial view of a target is available.

6 DISCUSSION

Several points of our lens proposal are worth discussing, as follows.

Lens activation: Our lens can support two types of explorations. First, when the user perceives a *part* of a target of interest in a classical DVR image, the lens can be used to reveal the target in full detail. This *directed* exploration supports the task ‘show me more information about *this* item’. The use-cases in Secs. 5.1-5.4 are of this type. Secondly, the user can open up a DVR volume at a 2D location from which no partial detail is visible. This is useful when we know that there *is* an interesting target buried in the volume even without seeing it (*corpus callosum* use-case in Sec. 5.5), thus supporting the task ‘show me the data I *know* it is somewhere in there’, or for free exploration to find unknown patterns

Fig. 9. Visualizing one day of aircraft trajectories over France [23]. (a) Overview of all trails. (b) Zoom, filtering, and color mapping techniques are used to highlight an outlier trajectory of an aircraft performing an eight-shaped loop. Revealing this outlier costs significant user effort.

in a volume, *i.e.* for the task ‘show me what this volume *may* hide in it’. In the first exploration type (target not fully occluded), our lens is simple and rapid to use – point, click, and optionally rotate light or viewpoint. In the second exploration type (target fully occluded or not even sure whether an interesting target exists in the data), the lens is equally simple to use, but several tries to select a suitable focus point and lens depth are needed.

Lens shape: Occluders are pushed away, and deformed, isotropically (Secs. 3.2, 3.4). This simple lens model requires a single parameter, the lens radius R , which makes its usage easy. The deformations evolve smoothly from the lens center (maximal) to outside the lens (no deformation), see Sec. 3.4, which effectively blends the local (in lens) focus with the global (out of lens) context (R3). However, this strongly compresses the deformed occluders close to the lens border, making them hardly visible when the lens is fully active. A possible refinement would be to reduce the deformation of the pushed-away occluders while still pushing them away, thereby improving the F+C effect (R3). However, this would occlude areas outside the lens, basically moving occlusion from *inside* the lens to *outside* and close to it. Finding an optimal balance between minimal deformation (so one can recognize the pushed-away occluders) and minimal clutter (so these occluders do not destroy the lens context) is a topic for future work. Separately, deformed rays may intersect with straight rays, thereby sampling the same voxel(s) to different image pixels. We did not observe in our usage any artifacts that can be ascribed to this issue, nor did the other users of our tool. This can be explained by the fact that such ray intersections are relatively few and we use a compositing transfer function, akin to a low-pass filter.

Parameter setting: Our lens depends on several parameters: the 2D lens center $\mathbf{f}$, lens radius R , lens axis direction $\mathbf{a}$, local light direction $\mathbf{l}^{lens}$, scattering start-distance t_{min} , and gathering and scattering parameters α and β . All these parameters are controlled via a mouse-driven virtual trackball, key modifiers, and the arrow keys (Sec. 3). As the lens works at 15 frames per second, the user can quickly tune the parameters and see their effect (R1). Moreover, all parameters start with good preset values (Sec. 3). A possible refinement would be to pre-segment the target, based on user-given values for $\mathbf{f}$, R , and t_{min} , thereby determining β automatically. However, we believe that manual control of the scattering β is important to allow users to choose their most suitable field-of-view angle. In fact, this flexibility allows a better exploration of the local context (R2).

Implementation: We implement our lens by modifying the ray trajectories constructed in the inner loop (per-pixel raycasting) of a public DVR raycaster [33]. Apart from this, we change the per-voxel lighting and transfer function based on the voxel location in the lens and the parameters given by user interaction (Sec. 3.3). Such changes are limited and easily applicable to any (parallel) raycaster.

Fig. 10. Inspecting an abnormal aircraft trail. (a) The abnormal trail is spotted in an all-trails view as it is highly curved while all other trails are relatively straight. Activating the lens at the outlier location (b) and changing the magnification factor (c) reveals the trail's eight-shape. (d) Rotating the viewpoint provides spatial insight on the embedding of the outlier in the surrounding trails.

Fig. 11. Revealing the *corpus callosum* in a DVR of a set of DTI tracts. (a) The *corpus callosum* is not even partially visible in the original DVR image. (b) The structure of interest is quite easily visible from most such viewpoints. (c) The viewpoint and/or the opacity transfer function can be adjusted to get an optimal view on the target.

Limitations: As explained, de-occluding a target requires either a small fragment thereof to be visible (if so, de-occlusion is very simple and fast), or requires the user to choose the lens focus and target depth based on other insights (which, as explained, requires more trial-and-error). At a higher level, many lens mechanisms exist in the literature, as discussed in Sec. 2. While we have argued that, to our knowledge, none of them simultaneously supports requirements R1...R4, comparing such mechanisms with our lens for specific use-cases and datasets is an important test for the *end-to-end* effectiveness of our proposal. We have not covered this point as obtaining (or replicating) implementations of such lenses is very challenging. This remains an important open point for future work – both for our proposal but also for all other volumetric lens proposals in the literature. In particular, none of the techniques in Tab. 1 were compared side-by-side against other techniques. We

have performed three user evaluations involving specialists in airport baggage security (Sec. 5.1), pulmonology (Sec. 5.3), and air traffic control (Sec. 5.4). In all cases, users were not involved in this work, nor with other work of the authors. However, the set-up of these evaluations stays at the level of formative user experiments. To confirm and refine the obtained (positive) findings, more formal user studies are needed, which we plan to cover next.

7 CONCLUSIONS

In this paper, we presented a new fish-eye-like context-and-focus lens that addresses the occlusion problems inherent in scalar volume rendering. The principle of our lens consists in first gathering (squeezing) rays so that they easily pass through occluding densities (given a user-specified opacity transfer function) and next scattering (fanning out) rays to best sample the target of interest. Our lens can be directly applied to any DVR raycaster and scalar volume dataset. Its main constraint is that the user should be able to find a viewpoint from which the target of interest, deep buried in the data, is at least slightly visible. We also present several modifications of the local rendering parameters within the lens (view direction, lighting parameters, opacity transfer function) that aim to both better separate the focus (lens) from the context (volume) and also allow more detailed examining of the target. Our lens is easy to use – all its parameters are controlled via direct mouse-and-keyboard interaction – and can be efficiently implemented atop of a standard GPU ray caster. Our lens is especially useful for highlighting structures of interest which are both deeply embedded in volumetric data and cannot be revealed by standard transfer function manipulations due to similar densities in the occluders and target. We demonstrate these points using five use-cases involving datasets from baggage detection, fluid visualization, air traffic control, and chest radiology, and DTI fiber tracts.

Several improvements to our proposal are possible, as follows. First and foremost, heuristics can be sought to link all our free parameters (lens size, focus depth, interpolation between focus and context) directly to the volume data, so the user interaction is minimized and therefore exploration efficiency is increased. Secondly, our lens could be extended to different types of volumetric datasets, such as multivariate (vector, tensor) fields. Last but not least, a formal wider-scale evaluation of how the lens addresses more specific tasks, and how it compares to existing tools for these tasks, such as other lens types, is a goal we aim to pursue next.

ACKNOWLEDGMENTS

The authors acknowledge the support of the French National Agency for Research (Agence Nationale de la Recherche ANR) under the grant ANR-14-CE24-0006-01 project TERANOVA and the SESAR Research and Innovation Action Horizon 2020 under project MOTO (The embodied reMote Tower). We also thank dr. Van de Perre and dr. Brozici from the pulmonology and radiology departments at the Heilig Hart Ziekenhuis, Mol, Belgium for their invaluable help with the chest tumor use-case described in Sec. 5.3.

REFERENCES

- [1] J. Alvina, C. Appert, O. Chapuis, and E. Pietriga. RouteLens: Easy route following for map applications. In *Proc. Intl. Working Conference on Advanced Visual Interfaces (AVI)*, pp. 125–128. ACM, 2014. doi: 10.1145/2598153.2598200
- [2] Y. Assaf and O. Pasternak. Diffusion tensor imaging (DTI)-based white matter mapping in brain research: A review. *Journal of Molecular Neuroscience*, 34(1):51–61, 2008. doi: 10.1007/s12031-007-0029-0
- [3] A. Brambilla, R. Carneeky, R. Peikert, I. Viola, and H. Hauser. Illustrative flow visualization: State of the art, trends and challenges. In *Proc. EuroVis – State of the Art Reports*. Eurographics, 2012. doi: 10.2312/conf/EG2012/stars/075-094
- [4] J. Brierley, M. Gospodarowicz, and C. Wittekind. *TNM classification of malignant tumours*. Wiley-Blackwell, 2017. doi: 10.1016/S1470-2045(17)30438-2
- [5] S. Bruckner and M. E. Gröller. Exploded views for volume data. *IEEE TVCG*, 12(5):1077–1084, 2006. doi: 10.1109/TVCG.2006.140
- [6] M. Burns and A. Finkelstein. Adaptive cutaways for comprehensible rendering of polygonal scenes. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 27(5):154:1–154:7, 2008. doi: 10.1145/1457515.1409107
- [7] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. Extending distortion viewing from 2D to 3D. *IEEE Computer Graphics and Applications*, 17(4):42–51, 1997. doi: 10.1109/38.595268
- [8] C. Correa, D. Silver, and M. Chen. Feature aligned volume manipulation for illustration and visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1069–1076, Sept. 2006. doi: 10.1109/Transactions on Visualization and Computer Graphics.2006.144
- [9] C. Correa, D. Silver, and M. Chen. Illustrative deformation for data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1320–1327, Nov. 2007. doi: 10.1109/Transactions on Visualization and Computer Graphics.2007.70565
- [10] C. D. Correa and K. L. Ma. Visibility histograms and visibility-driven transfer functions. *IEEE Transactions on Visualization and Computer Graphics*, 17(2):192–204, 2011. doi: 10.1109/Transactions on Visualization and Computer Graphics.2010.35
- [11] J. Cui, P. Rosen, V. Popescu, and C. Hoffmann. A curved ray camera for handling occlusions through continuous multiperspective visualization. *IEEE TVCG*, 16(6):1235–1242, 2010. doi: 10.1109/TVCG.2010.127
- [12] P. Dragicevic, A. Bezierianos, W. Javed, N. Elmqvist, and J.-D. Fekete. Temporal distortion for animated transitions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 2009–2018. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979233
- [13] N. Elmqvist and P. Tsigas. A taxonomy of 3D occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008. doi: 10.1109/Transactions on Visualization and Computer Graphics.2008.59
- [14] M. H. Everts, E. Begue, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg. Exploration of the brain's white matter structure through visual abstraction and multi-scale local fiber tract contraction. *IEEE Transactions on Visualization and Computer Graphics*, 21(7):808–821, 2015. doi: 10.1109/TVCG.2015.2403323
- [15] M. H. Everts, H. Bekker, J. Roerdink, and T. Isenberg. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, 2009. doi: 10.1109/TVCG.2009.138
- [16] M. Griebel, T. Preusser, M. Rumpf, M. A. Schweitzer, and A. Telea. Flow field clustering via algebraic multigrid. In *Proc. IEEE Visualization*, pp. 35–42, 2004. doi: 10.1109/VISUAL.2004.32
- [17] W.-H. Hsu, K.-L. Ma, and C. Correa. A rendering framework for multi-scale views of 3D models. *ACM Trans. Graph.*, 30(6):131:1–131:10, Dec. 2011. doi: 10.1145/2070781.2024165
- [18] C. Hurter. Image-based visualization: Interactive multidimensional data exploration. *Synthesis Lectures on Visualization*, 3(2):1–127, 2015.
- [19] C. Hurter, S. Conversy, D. Gianazza, and A. Telea. Interactive image-based information visualization for aircraft trajectory analysis. *Transportation Research Part C: Emerging Technologies*, 47:207–227, 2014.
- [20] C. Hurter, O. Ersoy, and A. Telea. Graph bundling by kernel density estimation. *Computer Graphics Forum*, 31(3pt1):865–874, 2012. doi: 10.1111/j.1467-8659.2012.03079.x
- [21] C. Hurter, R. Taylor, S. Carpendale, and A. Telea. Color tunneling: Interactive exploration and selection in volumetric datasets. In *Proc. IEEE PacificVis*, pp. 225–232, March 2014. doi: 10.1109/PacificVis.2014.61
- [22] C. Hurter, A. Telea, and O. Ersoy. MoleView: An attribute and structure-based semantic lens for large element-based plots. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2600–2609, 2011. doi: 10.1109/TVCG.2011.223
- [23] C. Hurter, B. Tissoires, and S. Conversy. FromDaDy: Spreading data across views to support iterative exploration of aircraft trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1017–1024, 2009. doi: 10.1109/TVCG.2009.145
- [24] S. Islam, D. Silver, and M. Chen. Volume splitting and its applications. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):193–203, March 2007. doi: 10.1109/Transactions on Visualization and Computer Graphics.2007.48
- [25] R. Kikinis, S. D. Pieper, and K. Vosburgh. 3D Slicer: A platform for subject-specific image analysis, visualization, and clinical support. *Intra-operative Imaging Image-Guided Therapy*, 3(19):277–289, 2014. doi: 10.1007/978-1-4614-7657-3_19
- [26] R. Kincaid. SignalLens: Focus+context applied to electronic time series. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):900–907, 2010. doi: 10.1109/Transactions on Visualization and Computer Graphics.2010.193
- [27] J. Kruger, J. Schneider, and R. Westermann. ClearView: An interactive context preserving hotspot visualization technique. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):941–948, 2006. doi: 10.1109/TVCG.2006.124
- [28] A. Lhuillier, C. Hurter, and A. Telea. FTEB: Edge bundling of huge graphs by the Fast Fourier Transform. In *Proc. IEEE PacificVis*, 2017. doi: 10.1109/PACIFICVIS.2017.8031594
- [29] W. Li and G. Paladini. Luggage visualization and virtual unpacking, 2013. WO patent WO2013142072A3, Siemens Corp.
- [30] W. Li, G. Paladini, L. Grady, T. Kohlberger, V. K. Singh, and C. Bahlmann. Luggage visualization and virtual unpacking. In *Proc. WASA SIGGRAPH Asia Workshop*, pp. 161–168. ACM, 2012. doi: 10.1145/2425296.2425325
- [31] M. J. McGuffin, L. Tancu, and R. Balakrishnan. Using deformations for browsing volumetric data. In *Proc. IEEE Visualization*, pp. 401–408, 2003. doi: 10.1109/VISUAL.2003.1250400
- [32] W. Merzkirch. *Flow visualization*. Elsevier, 2012.
- [33] NVIDIA, Inc. CUDA SDK code samples, 2018. http://www.nvidia.com/object/cuda_get_samples_3.html.
- [34] C. Pindat, E. Pietriga, O. Chapuis, and C. Puech. JellyLens: Content-aware adaptive lenses. In *Proc. ACM Symp. on User Interface Software and Technology (UIST)*, pp. 261–270, 2012. doi: 10.1145/2380116.2380150
- [35] B. Preim, A. Baer, D. Cunningham, T. Isenberg, and T. Ropinski. A survey of perceptually motivated 3d visualization of medical image data. *Computer Graphics Forum*, 35(3):501–525, 2016. doi: 10.1111/cgf.12927
- [36] C. Rezk-Salama and A. Kolb. Opacity peeling for direct volume rendering. *Computer Graphics Forum*, 25(3):597–606, 2006. doi: 10.1111/j.1467-8659.2006.00979.x
- [37] H. Sonnet, S. Carpendale, and T. Strothotte. Integrating expanding annotations with a 3D explosion probe. In *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '04*, pp. 63–70. ACM, New York, NY, USA, 2004. doi: 10.1145/989863.989871
- [38] A. Telea and J. J. van Wijk. Simplified representation of vector fields. In *Proc. IEEE Visualization*, pp. 263–270, 1999. doi: 10.1109/VISUAL.1999.809865
- [39] C. Thiede, G. Fuchs, and H. Schumann. Smart lenses. In *Proc. Intl. Symp. on Smart Graphics*, pp. 178–189. Springer, 2008. doi: 10.1007/978-3-540-85412-8_16
- [40] C. Tominski. Event-based concepts for user-driven visualization. *Information Visualization*, 10(1):65–81, 2011. doi: 10.1057/ivs.2009.32
- [41] C. Tominski, J. Abello, F. van Ham, and H. Schumann. Fisheye tree views and lenses for graph visualization. In *Proc. Intl. Conf. on Information Visualisation (IV)*, pp. 17–24, 2006. doi: 10.1109/IV.2006.54
- [42] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. Interactive lenses for visualization: An extended survey. *Computer Graphics Forum*, 36(3):173–200, 2017. doi: 10.1111/cgf.12871
- [43] C. Tominski, H. Schumann, G. Andrienko, and N. Andrienko. Stacking-based visualization of trajectory attribute data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2565–2574, 2012. doi: 10.1109/Transactions on Visualization and Computer Graphics.2012.265
- [44] X. Tong, J. Edwards, C. M. Chen, H. W. Shen, C. R. Johnson, and P. C. Wong. View-dependent streamline deformation and exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(7):1788–1801,

- July 2016. doi: 10.1109/Transactions on Visualization and Computer Graphics.2015.2502583
- [45] X. Tong, C. Li, and H. W. Shen. Glyphlens: View-dependent occlusion management in the interactive glyph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):891–900, 2017. doi: 10.1109/Transactions on Visualization and Computer Graphics.2016.2599049
- [46] M. Traoré and C. Hurter. Interactive exploration of 3d scanned baggage. *IEEE Computer Graphics and Applications*, 37(1):27–33, Jan 2017. doi: 10.1109/MCG.2017.10
- [47] M. van der Zwan, V. Codreanu, and A. Telea. CUBu: Universal real-time bundling for large graphs. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2550–2563, 2016. doi: 10.1109/TVCG.2016.2515611
- [48] L. Wang, Y. Zhao, K. Mueller, and A. Kaufman. The magic volume lens: an interactive focus+context technique for volume rendering. In *Proc. IEEE Visualization*, pp. 367–374, 2005. doi: 10.1109/VISUAL.2005.1532818
- [49] M. Q. Wang Baldonado, A. Woodruff, and A. Kuchinsky. Guidelines for using multiple views in information visualization. In *Proc. Working Conf. on Advanced Visual Interfaces (AVI)*, pp. 110–119. ACM, 2000. doi: 10.1145/345513.345271
- [50] M. L. Wu and V. Popescu. Multiperspective focus+context visualization. *IEEE TVCG*, 22(5):1555–1567, 2016. doi: 10.1109/TVCG.2015.2443804
- [51] Y. Wu and H. Qu. Interactive transfer function design based on editing direct volume rendered images. *IEEE Transactions on Visualization and Computer Graphics*, 13(5):1027–1040, 2007. doi: 10.1109/TVCG.2007.1051