

Dynamic Visualisation of Orbital Fat Deformation using Anatomy-Guided Interaction

P.J. Schaafsma, S. Schutte, H.J. Simonsz, F.H. Post, C.P. Botha

Abstract

The human eye is a biomechanical system. Orbital fat plays an important role in the working of this system, but its behaviour during eye movement is not well understood. To give insight into this behaviour, visualisation is a useful tool. This paper presents a complete pipeline for interactive particle-based visualisation and exploration of orbital fat deformation from MRI data. Sensible 3D particle seeding is important in this type of visualisation. We address that problem with a two-step process: Interactive, anatomy-guided slice positioning, and contour-based region of interest specification. Since the deformation calculation is unlikely to be correct everywhere, we derive and visualise an uncertainty measure based on deformed and original MRI data. We also performed a case study evaluation to investigate the benefits of our approach towards orbital fat deformation visualisation.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Interaction techniques

1. Introduction

The human eye is a biomechanical system in which the eyeball, optic nerve and eye muscles work together to direct the eye and capture images to be sent to the visual cortex. These anatomical structures reside inside the orbita, or eye socket. The orbital fat is the semiliquid tissue that fills the space between these structures and passively deforms according to the shape and position of other tissue. This happens, for instance, as the optic nerve is dragged through it due to eyeball rotation, or when it is being pushed away by the left eye muscle rolling up onto the eyeball when gazing to the right. The eye fat plays many roles, such as keeping the muscles from pulling the eyeball back into the eye socket, lubricating the muscles' movement against the orbita wall, and dampening the rapid and sudden rotation of the eyeball, which can quickly accelerate to upto 1000° per second.

The behaviour of orbital fat is not well understood. Improving this is important because a better understanding leads to more informed decisions during eye surgery. To this end, visualisation is a valuable tool. Its goal is to give insight into complex data. Deformation data is often presented as low-level 3D time-dependent vector data, while human understanding responds better to a higher-level description as patterns and features, such as pushing-away, dragging-along and gliding-through. Visualisation helps to close the gap between these different levels of representation.

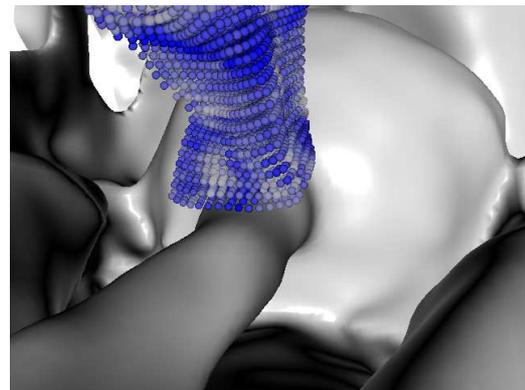


Figure 1: Particle visualisation of fat deformation in the right orbita, posterior sideways view. A regular grid of particles was seeded behind the optic nerve, which in this view is moving downwards. The image shows a deformed grid, indicating fat movement. Particles are coloured according to a deformation uncertainty measure.

In this paper, we present a complete pipeline for the dynamic visualisation and interactive exploration of orbital fat deformation, see figure 1. After MRI data acquisition of an axial skull section containing both orbitae, we use the Lucas and Kanade optic flow algorithm to calculate deforma-

tion data as 3D time-dependent vector fields. An isosurface rendering of relevant anatomical features provides a frame of reference in which further visualisation and exploration take place. Anatomy-guided region of interest specification provides a seeding volume for particle-based visualisation of deformation data. An uncertainty measure, derived from comparison of deformed and original MRI data, is used to qualify the particle visualisation. These techniques are complemented by streamlines and pathlines, also seeded from a specified region of interest. Furthermore, we evaluated our approach through a case study conducted with an eye biomechanics scientist. A video demonstration of our work is available online[†].

With this paper, our main contributions are:

- Intuitive, anatomically relevant slice positioning and region of interest specification
- Focus+Context uncertainty visualisation for deformation data
- Improved visualisation of orbital fat deformation through tightly coupled dynamic visualisation and interactive exploration

The rest of this paper follows standard structure guidelines. After related work (section 2), we focus on the body of our research: the visualisation and interaction pipeline (section 3). Then, the results of a domain expert evaluation are presented (section 4), and we conclude with analysis of our results and mention promising avenues for future research (section 5).

2. Related Work

Visualisation of orbital fat deformation was researched by Botha et al. [BdGS*07]. They use the Demons registration algorithm to calculate deformation fields from a sequence of MR images. While we only focus on general, interactive, dynamic visualisation of this type of data, their work also addressed specific questions on the relative motion of fat tissue to other orbital structures. The 3D Lucas and Kanade algorithm we use to obtain orbital deformation data from MR images was presented in detail by Abramoff et al. [AV02]. While we follow their work in the calculation of deformation data, we use a different approach for the uncertainty calculation that is closer to the data. Instead of deriving deformation data from MR images data, Schutte et al. [SvdBvK*06] worked on development of a finite element model that describes the eye's behaviour. Schoemaker et al. [SHM*06] determined material properties of orbital tissue for use in such an FEM model.

Fuchs et al. [FWH10] use an outline of the region of interest (ROI) on the view plane to create an initial ROI. Interactive operation on vertices and edges allows creation of

complex polyhedra. This makes their technique general, albeit slower than ours. They also discuss and compare region of interest specification techniques in general. Schiffman et al. [SNB*03] draw closed contours a brain surface to directly specify a surface ROI.

Research on data-guided plane placement beyond standard orthogonal slices and freely rotating surfaces is limited. Kreylos et al. [KBK08] fit primitives, such as cylinders and cubes, to a point set. The point set is a result of interactive brushing in point cloud rendering. Being able to fit many types of shapes is useful, although it does require the extra step of selecting the type of shape to fit.

Visualising general uncertainty in time-dependent vector fields was researched for texture-based visualisation by Botchen et al. [BWE05]. They distinguish uncertainty levels by parameterising visualisation style, but do not necessarily guide user attention towards low uncertainty regions.

3. Deformation Visualisation Pipeline

In this section we present in relevant detail the complete pipeline from MRI data acquisition to the visualisation results. The pipeline can be divided into three main sections: data acquisition and processing, interaction, and visualisation, see figure 2. All three sections work together to create an interactive environment for unrestricted exploration of orbital fat deformation.

3.1. Data acquisition and Processing

Our dataset consists of 15 axial MRI volumes containing both orbitae of a volunteer subject. Each 3D scalar volume corresponds to a different gaze direction, ranging from -30° (left) to $+40^\circ$ (right) axial rotation at constant 5° increments. The dimensions of each volume are $256 \times 256 \times 116$ voxels (sagittal \times coronal \times axial), spaced at $0.5469\text{mm} \times 0.5469\text{mm} \times 0.5000\text{mm}$. During acquisition, the subject was asked to bite down on a rigid landmark object, which enabled accurate dataset realignment afterwards. The raw scalar data is used in three subsequent processing steps: segmentation, deformation calculation, and uncertainty calculation. We will discuss the latter two in more detail below.

We assume the rotation in each volume to be exactly the specified, incremental angle, instead of having measured it ourselves after acquisition, for instance by measuring the orientation of the optic nerve and lens, both of which can be accurately identified in the MRI data. We also assume that the deformation of the orbital structures during high speed rotation equals the static deformation of orbital structures at fixed eyeball rotation angle, i.e. dynamic deformation equals static deformation. Under this assumption, the actual deformation depends only on view direction, not on angular velocity of the eyeball. Also, while spatial resolution is dictated

[†] <http://graphics.tudelft.nl/publications/Schaafsma2010-OrbitalFatDeformation.avi>

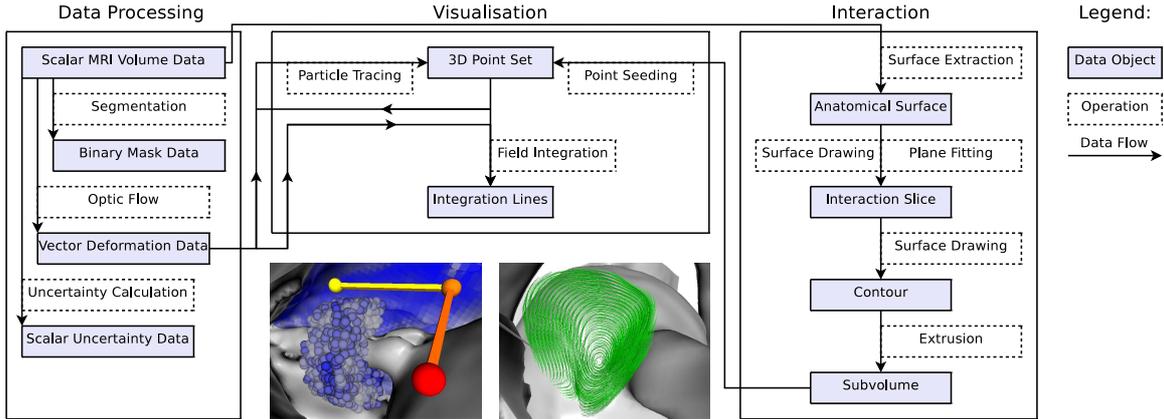


Figure 2: Three main components of our visualisation pipeline for orbital fat deformation. After MRI data acquisition (not shown), data processing, interaction and visualisation work together to make interactive exploration of the dataset possible. Note that the main purpose of the interaction pipeline component is to provide a point set for further visualisation.

by the acquisition process, we define the time frame between consecutive images simply as unit time.

3.1.1. Deformation Calculation

Following Abràmoff [AV02], we used a 3D implementation of the Lucas and Kanade (LK3D) [LK81] optic flow algorithm to calculate dense deformation data for each of the 15 scalar volumes. LK3D produces per voxel deformation data solely based on a least squares estimate from spatial and temporal image derivatives over a small, Gaussian-weighted region of influence. Aside from allowing generation of physically improbable data, it can also produce erratic deformation data. At the cost of sensitivity to small scale features, this can be improved by either smoothing the original scalar data or image derivatives, or by increasing the algorithm's region of influence. The image derivatives are calculated by convolution of the MR images with a 4D Gaussian derivative kernel.

The spatial and temporal image derivative were created from MR image data by convolution with a spatially symmetric 4D Gaussian derivative kernel. The kernel width was 23 voxels with a variance of 1.0mm^2 in each spatial direction, and 3 voxels and 1.0 respectively in time. LK3D weighted least squares minimisation also uses a Gaussian kernel to weight the contributions of individual pixels. We used a kernel width of 13 voxels and variance of 4.0mm^2 in spatial direction. We did not minimise over multiple time steps.

3.1.2. Uncertainty Calculation

The LK3D algorithm does not produce accurate results everywhere, to which several factors contribute. First, noise in MRI data cascades through the derivative images to the deformation data. The sensitivity of LK3D to input noise is

highest in uniform intensity regions. Second, the quantised character of the MRI data, both in time and space, allows only estimates of the true image derivatives. Convolution of the MRI data with a 4D Gaussian derivative kernel addresses both these issues, but also smooths out small-scale image patterns. Its variance determines the balance between sensitivity to detail and noise robustness. Third, the Gaussian variance of the weighted LK3D region of influence introduces a similar source of error.

To calculate uncertainty for a single deformation image we first place a regular, super-sampled grid of virtual particles in the domain. These particles are then advected over unit time using the deformation data. Each particle carries the linearly interpolated intensity value from its start location. The intensity value of the image next in sequence at the advected particle's location is sampled and subtracted from the carried intensity value. A large difference leads to high uncertainty. The uncertainty value is then distributed over eight voxels similar to the intensity value interpolation at initial particle location. Since the last deformation field has neither a next scalar nor deformation field, we duplicate the second to last uncertainty volume. This way we have an uncertainty field for each scalar and vector field.

3.2. Interaction

Trying to visualise too much deformation data at once will lead to many of the problems common to the 3D visualisation challenge, mainly occlusion and cluttering. Specifying a region of interest addresses these problems. We do this as a two-step process. First, the user specifies a slice through the dataset by drawing a set of points on a reference surface. After fine-tuning slice position and orientation, a region of interest is defined by drawing and extruding a contour on the specified slice. Both steps are discussed in detail below.

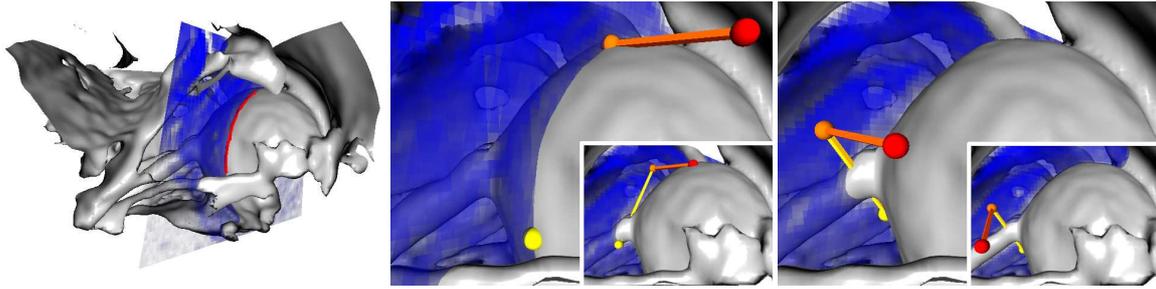


Figure 3: Slice positioning and adjustment in the right orbita, inferior axial view. Prominent features in the grey isosurface are the eyeball (large spherical shape to the right) and optic nerve (cylindrical structure attached on the left side to the eyeball). Left: A slice is determined by least squares fitting a plane to a point set (red) drawn on the eyeball. Centre: First and last drawn point form the basis for a slice transformation widget. Inset shows the slice dragged through the volume along the plane normal. Right: Repositioning of the widget by drawing a line on the slice through the optic nerve. Inset shows the slice after rotation according to the new widget.

3.2.1. Anatomy-guided Slice Positioning

To provide an anatomical frame of reference for slice positioning, we observe that in MRI data, different anatomical structures are represented by different scalar ranges. Although they overlap and MRI data is noisy, the main features of the eye are clearly distinguishable. Isosurface extraction and rendering is a simple yet sufficient technique to start with when placing an anatomically relevant slice inside a volume.

Drawing directly on the isosurface leads to a point set from which we determine the position and orientation of a plane using least squares estimation of the four plane equation parameters. This technique effectively gives the user the opportunity to draw on a visible surface an approximation of part of the prospective slice intersection with the surface, although results are well-defined for almost any drawing pattern. Another way of using this technique in an interactive viewing environment, is to first reposition the camera such that it would coincide with the plane. A single, straight drawing gesture then quickly defines a plane of desired orientation and position. In our experience, this last method is easier and more reliable, especially in highly curved areas of the reference surface.

Final adjustment or predictable repositioning of the slice is required if first results are either unsatisfactory or planned slice position and orientation are otherwise closely related to a prior placement attempt. We have designed a three-component, L-shaped widget that allows normal translation and two degrees of rotational freedom of the slice, see figure 3. The widget is operated by adjusting the position of the three spherical components. The orange corner component allows translation in the plane normal direction. The yellow in-plane component rotates the plane around the axis normal to the widget and through the translation component. The red out-of-plane component rotates the plane around the

in-plane axis of the widget. Initially, the widget's location and orientation are determined by the point set from which the slice itself was derived. The orange and yellow component correspond to the start and end drawing point projected onto the plane. Simply redrawing on the plane repositions the widget according to the new point set.

This interaction style for plane placement strikes a balance between complete freedom of orientation and positioning, and predictable, data-guided results. Note that the initial placement using the second method already often provides good results. Although we demonstrate this technique by using the plane as the basis for region of interest specification, other techniques that rely on slices through in volumetric data, such as Multi-Planar Reconstruction, also benefit from a more intuitive, data-guided approach.

3.2.2. Region of Interest Specification

Once the interaction slice is in place, specifying a region of interest by extrusion of a contour drawn on the slice is straightforward. Contour drawing is easily distinguished from widget repositioning by start and endpoint proximity testing. Having completed a contour, extrusion in the slice normal direction using a widget similar to the translation component of the slice widget defines a closed region of interest volume, see figure 4. This region is used during visualisation as a point seeding volume, to be used in the particle-based visualisation as well as for streamline and pathline generation. Contour drawing on the interaction slice allows specification of custom regions of interest, bound only by the requirement that they be an extrusion of a non-selfintersecting contour of coplanar points.

Particles are initially seeded randomly over the data domain. They are also tagged as invalid and hence not displayed. During tracing, described in section 3.3.3, this is detected and individual particles are reseeded inside the region

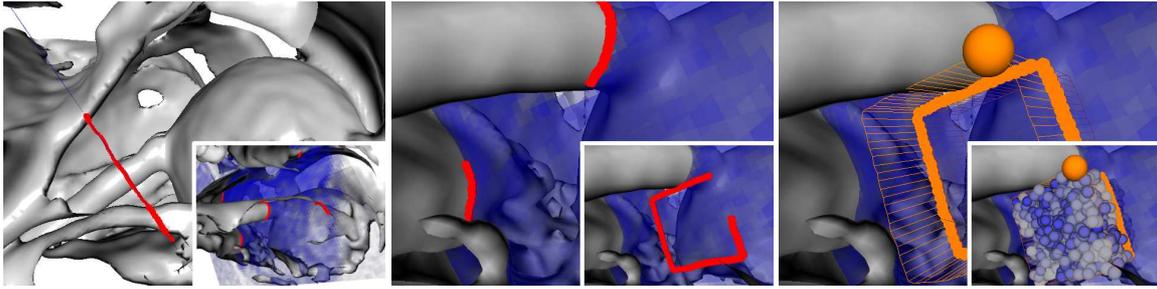


Figure 4: Region of interest specification. Left: Direct placement of a slice through the optic nerve in the view direction. Inset is the result of a camera change towards the bottom left. Centre: A zoomed-in, rotated view shows the drawing point set distribution over the isosurface. Inset shows a contour being drawn. Right: A contour is extruded to form a region of interest just below the optic nerve. Inset shows particles seeded inside the region of interest, just after animation has started.

of interest. For efficiency and speed, this process is split into two phases. First, upon contour completion, initialisation of a texture that separates the contour inside from the contour outside. Second, when a particle is reseeded, remapping of the particle position and a texture lookup to see whether the new location is inside the region of interest. This separation ensures, after initialisation, very fast region of interest inside/outside testing for prospective particle reseed locations. In detail, we take the following steps.

Initialisation of reseeded texture upon contour completion:

- Obtain an approximately minimal in-plane bounding rectangle around the contour using Principal Component Analysis.
- Map the contour points onto the unit square. To do this, express the contour points as a linear combination of the scaled principal components, such that the mapped coordinates are within $[0.0, 1.0]$.
- Use an Advancing Front algorithm [PVMZ87] to create an inside mesh for the mapped contour.
- Render the mesh as a unit quad to a texture of arbitrary resolution. This texture will be available until another contour is drawn.

Particle reseeded during particle tracing:

- Hash the particle coordinates into the unit cube, providing some randomness.
- Do a texture lookup using hashed particle coordinates projected onto the unit square.
- If lookup indicates 'outside', rehash the particle location in a way that prevents repeated failure.
- If lookup indicates 'inside', remap the hashed particle location back to world coordinates, and validate the particle.

3.3. Visualisation

Immediately after specifying a seed volume, the user can start the dynamic part of the visualisation, after which

the chosen deformation field representation and anatomical structures are continuously animated over the available time points. At any time, settings such as the resolution of seeding can be adapted interactively, with immediate visual feedback. In the following subsections, the available visualisation techniques are discussed in more detail.

3.3.1. Anatomical Context Visualisation

The goal of anatomy visualisation as context to the deformation data is little more than providing a frame of reference. This helps the viewer to form a mental image of the eye fat deformation during exploration. It should not attract attention, but be there to give meaning to the deformation visualisation. We use both the original MR images and the isosurfacing results to provide different levels of context.

The interaction plane is presented as an interpolated, semi-transparent slice of the original MRI data. It provides only minimal context, but avoids any occlusion problems caused by the isosurface. Visualised together with the isosurface, the textured, semi-transparent plane does partly occlude the reference surface, while against a plain background, scalar structures are clearly visible.

3.3.2. Uncertainty Visualisation in Deformation

We visualise uncertainty by adjusting the rendering style of individual particles, an example of which is shown in figure 5. The uncertainty is based on two measures: A normalised, linearised value derived from deformation and scalar data, see section 3.1.2, and a binary mask value derived from segmentation indicating whether a particle finds itself in the orbital fat. Borrowing ideas from Focus+Context, we use visibility, colour, and spatial frequency to naturally avert user attention away from high uncertainty regions.

We have designed an uncertainty transfer function that combines both measures into a single function. First, a particle's location is checked against the orbital fat mask, de-

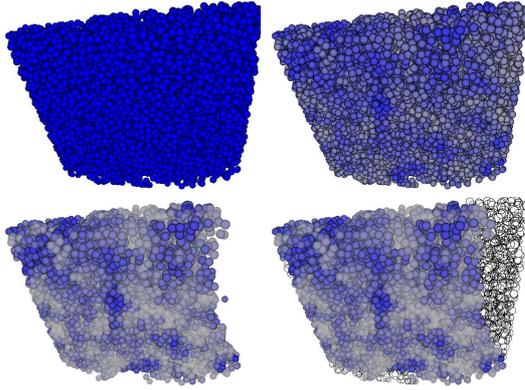


Figure 5: *Uncertainty visualisation style. Top-Left: no uncertainty visualisation, all particles rendered in the same style. Top-right: Color as indication of uncertainty. Higher saturation equals lower uncertainty. High spatial frequency makes all particles attract almost equal attention during animation. Bottom-Left: Fading of sharp edges for high uncertainty data naturally draws attention to other regions. This is especially noticeable during animation. Out-of-mask particles on the right are not shown. Bottom-Right: Silhouette visualisation of out-of-mask particles. Note that due to their number, they are not clearly distinguishable. The rings also attract too much attention due to their crisp edges.*

cluding between out-of-mask and in-mask particle rendering. Second, the uncertainty value for in-mask particles is mapped onto a perceptually linear colour scale. The rendering style for out-of-mask particles had to convey just enough information to show that the particles are there, but not attract user attention away from in-mask particles, however high their uncertainty may be.

3.3.3. Particle-based Deformation Visualisation

Interactive particle-based visualisation can show global deformation patterns as well as small-scale local features, see figure 6. Its insensitivity to noise, due to a few in thousands of particles not detracting from global flow patterns, makes it a good visualisation choice for our type of data. As an exploratory technique for time-dependent deformation data it is useful as well, as it assumes nothing about the data and captures small-scale temporal aspects of deformation. Combined with the particle seeding strategy presented in this paper, this technique delivers good results. Identifying larger-scale temporal aspects or spatial patterns in the deformation data is more difficult, as a particle has no memory.

Interactive adjustment of particle size is useful to match the scale of the visualisation to the region of interest. Note that the size of a particle is a visual attribute only; the particle is still advected over time as a dimensionless point. Using particle density to determine the number of particles to

be seeded inside the region of interest is an improvement over using a fixed number of particles. The advantage is that during exploration, specifying different regions-of-interest of varying size, the particle density is not affected, leading to a more uniform visualisation style between exploration iterations.

3.3.4. Pathlines and Projected Streamlines

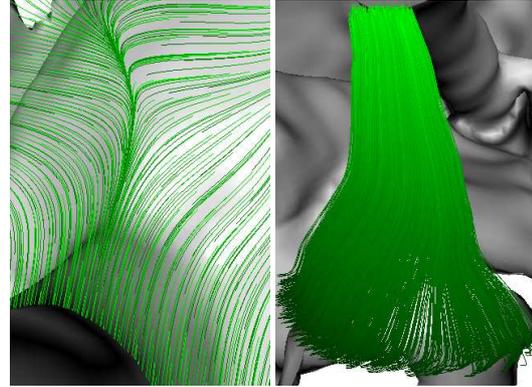


Figure 7: *Streamlines and pathlines. Left: Projected Streamlines behind the opticus (dark grey patch on the bottom) as it moves downwards. The projection of the streamlines onto the interaction plane might give the idea that fat tissue is being compressed. This is not the case, as the projection ignores any motion perpendicular to the interaction plane. Right: Pathlines seeded just below the bottom edge of the opticus show that the fat initially travels downwards with the opticus, but eventually gets pushed to the side. Color value indicates time.*

We also visualise the deformation data using pathlines and streamlines, see figure 7. These techniques show the deformation field in a different way than particle animation. Although from simple pathlines, for instance, it is hard to get a feel of how a deformation field behaves from one instant to the next, when densely seeded they do indicate the volume that the deforming matter in the seeding region has occupied over time. Similarly, streamlines, being instantaneous by nature, are not always suitable to give insight into time-dependent data, but rather serve to inspect the characteristics of the deformation field at a given instance.

4. Evaluation

Based on the principles and terminology set out by Yin for case study research [Yin09], we performed an evaluation, investigating the use of the visualisation tool by the second author of this paper, a published expert in eye biomechanics. Henceforth, we refer to this expert simply as “the user”. The main study question was formulated as “How can the visualisation tool assist biomechanics researchers in studying the deformation of orbital fat under different directions

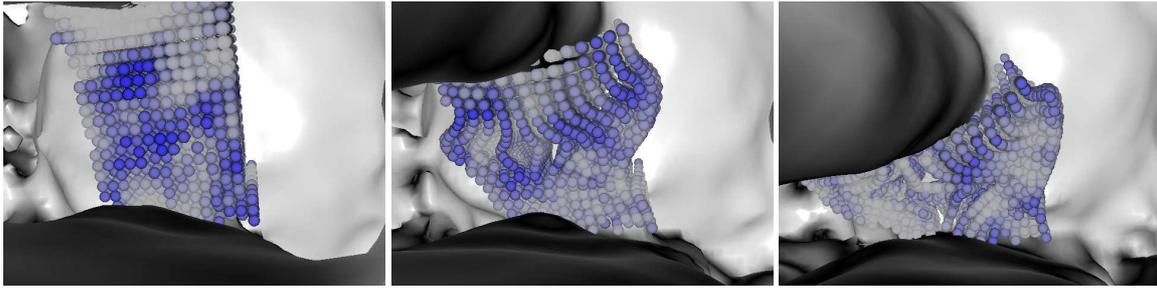


Figure 6: Deformation of orbital fat. Left: Particles are laid out in a grid pattern just below the opticus. The eyeball is in the back. Centre: As the opticus moves downwards, orbital fat (represented by particles) is starting to be pushed downward and to the side, out of the way of the opticus. Right: As the opticus moves close to one of the eye muscles, fat is getting pushed more and more to the side.

of gaze?” and the case was defined as “Use of the tool by an expert in eye biomechanics.” We have chosen for this evaluation method as we wanted to study the application of the complete tool *in vivo*, and as the potential user group is quite small.

The evaluation was performed partly with the user operating the tool by himself, and partly by the first author demonstrating. During the evaluation, we investigated a number of case study propositions that we had previously defined independently of the user. The rest of this section consists of the user’s feedback, structured along the study propositions (printed in italics), followed by our conclusions.

The anatomy-guided plane interaction approach enables the rapid definition of meaningful 2D planar reconstructions of the 3D data. The user, having extensive experience working with CAD software, was in general quite impressed with the anatomy-guided plane placement approach. Besides the general utility of the method, the ability to rotate the plane easily around its own axis was also appreciated. The following points of criticism were discussed: 1) Due to the free plane placement, as opposed to the more traditional orthogonal planes, maintaining a sense of orientation was difficult. Adding orientation feedback, for example a glyph representing the gaze direction and the foot-head direction would alleviate these problems. 2) The plane manipulation spheres can be confusing. These should be replaced with glyphs that clearly show their purpose. 3) The effectivity of the plane placement approach is of course dependent on the quality of the anatomical surfaces.

The anatomy-guided plane interaction and volume of interest specification enables the rapid definition of volumes containing potentially interesting, according to expert expectation, fat deformations. After a few minutes of instruction, the user was able to independently explore the data and rapidly specify anatomy-based regions of interest. It was remarked that the method assumes good knowledge of the anatomy under study.

The anatomy-guided interaction is significantly better (easier / more intuitive / more effective) than other systems with which the domain scientist has experience. The user has experience with traditional CAD-software interfaces, but remarked that our interaction approach is indeed better in the context of anatomy-driven selection.

The dense real-time particle tracing facilitates insight into complex deformation patterns better than region-query based techniques such as advection balls. (animation / dynamic; density) The user agreed strongly with this proposition and added that the dense particle tracing was more intuitive.

Confidence is clearly and intuitively represented. Initially we used outline rendering for particles outside of the orbital fat, and a perceptually linearised yellow-to-blue colour scale to represent confidence or uncertainty in the rest of the deformation field. We asked the user to guess what the meaning was of the different colouring, without giving him any other information, upon which he speculated that the outline rendering meant that there was no data, or that it was an extrapolation of other valid data. When we explained the meaning, they commented that the blue-to-yellow scale was not immediately clear in terms of what was high and low, and that they found the grey-to-blue part of the colour scale more suitable in this regard.

Visual representation of confidence / uncertainty aids in the correct interpretation of the fat deformation, as it indicates where the deformation model results should be more carefully considered. The user remarked that due to all particles moving, it was hard to intuitively associate colour with confidence and integrate its implications during the visual interpretation of the data. He did admit that if he was to discover an interesting phenomenon in a low confidence area, he would double-check the data. The remark was also made that the outline representation was clearly distinguished from the rest of the particles.

The visualisation reflects the expectations of the eye

biomechanics domain scientist. During the evaluation, the user completely independently specified a volume of interest between the lateral eye muscle and the eyeball and managed to create a comprehensive visualisation of the orbital fat evacuating rapidly to above and below as the muscle rolls up on the eye during gaze change. This behaviour is the topic of great interest in the eye biomechanics community. Besides this example, the user agreed that for a number of expected behaviours the visualisation reflects his expectations.

The visualisation could be used as a component in the validation of numerical models of the eye (compare model behaviour globally to measured behaviour). There was emphatic agreement with this proposition, with the user adding that the visualisation gives insight into the differences between simulation and measurement.

The visualisation could improve current research on orbital fat deformation (insight / understanding / hypotheses / practical applications for research results / compare FEM behaviour to this behaviour?) The user agreed strongly with the statement that it would facilitate understanding and commented that research up to now was primarily hypothesis-driven and region-query based, as opposed to explorative.

From these propositions, we conclude that the tool can assist eye biomechanics researchers by offering them a completely new way to visually study the behaviour of orbital fat as it deforms during eye movement, both through the dense real-time particle tracing and through the flexible selection of anatomically relevant regions of interest. In addition, a number of suggestions for improvement can now be acted upon, most important of which the anatomical orientation feedback. The intriguing role of uncertainty visualisation requires further study, although these first evaluation steps are promising.

5. Conclusion and Future Work

In this paper, we have presented a complete pipeline for the interactive visualisation of orbital fat deformation under different directions of gaze. Our work improves on the state of the art in two ways: First, interactive anatomy-guided plane placement and volume of interest specification allows for the flexible exploration of deformation of data in an anatomical context. Second, dense particle-based visualisation, according to a domain expert, significantly improves on visualisation styles employed by earlier work in this area. In addition, we investigated the role our application could play in the study of orbital fat deformation through a case study research evaluation conducted together with an eye biomechanics expert.

We plan to deploy and refine our tool in the further study of orbital fat deformation. During the evaluation, previously elusive aspects such as the vertical fat evacuation during eye movement could be convincingly visualised and have led to two of our domain scientist colleagues requesting a more

in-depth investigation of the clinical implications. Finally, we would like to apply these techniques to other types of anatomical deformation.

References

- [AV02] ABRÀMOFF M. D., VIERGEVER M. A.: Computation and Visualization of Three-Dimensional Soft Tissue Motion in the Orbit. *IEEE Transactions on Medical Imaging 21* (2002), 296–304.
- [BdGS*07] BOTHA C. P., DE GRAAF T., SCHUTTE S., ROOT R., WIELOPOLSKI P., VAN DER HELM F. C. T., SIMONSZ H. J., POST F. H.: MRI-based visualisation of orbital fat deformation during eye motion. In *Visualization in Medicine and Life Sciences (VMLS)* (2007), Linsen L., Hagen H., Hamann B., (Eds.), pp. 221–236.
- [BWE05] BOTCHEN R., WEISKOPF D., ERTL T.: Texture-based visualization of uncertainty in flow fields. In *Proceedings of IEEE Visualization 2005* (2005), Cite-seer, p. 647aA\$654.
- [FWH10] FUCHS R., WELKER V., HORNEGGER J.: Non-convex polyhedral volume of interest selection. *Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society 34*, 2 (2010), 105–13.
- [KBK08] KREYLOS O., BAWDEN G., KELLOGG L.: *Advances in Visual Computing*, vol. 5358 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [LK81] LUCAS B. D., KANADE T.: An Iterative Image Registration Technique with an Application to Stereo Vision, 1981.
- [PVMZ87] PERAIRE J., VAHDATI M., MORGAN K., ZIENKIEWICZ O. C.: Adaptive remeshing for compressible flow computations. *Journal of Computational Physics 72*, 2 (1987).
- [SHM*06] SCHOEMAKER I., HOEFNAGEL P. P. W., MASTENBROEK T. J., KOLFF C. F., SCHUTTE S., VAN DER HELM F. C. T., PICKEN S. J., GERRITSEN A. F. C., WIELOPOLSKI P. A., SPEKREIJSE H., SIMONSZ H. J.: Elasticity, viscosity, and deformation of orbital fat. *Investigative ophthalmology & visual science 47*, 11 (November 2006), 4819–26.
- [SNB*03] SHIFFMAN S., NG Y.-R., BROSNAN T. J., ELIEZ S., LINKS J. M., KELKAR U. V., REISS A. L.: Interactive specification of regions of interest on brain surfaces. *NeuroImage 20*, 3 (November 2003), 1811–6.
- [SvdBvK*06] SCHUTTE S., VAN DEN BEDEM S. P. W., VAN KEULEN F., VAN DER HELM F. C. T., SIMONSZ H. J.: A finite-element analysis model of orbital biomechanics. *Vision research 46*, 11 (2006), 1724–31.
- [Yin09] YIN R. K.: *Case Study Research: Design and Methods*, fourth ed. Sage, 2009.