Eurographics/ IEEE-VGTC Symposium on Visualization 2009
H.-C. Hege, I. Hotz, and T. Munzner
(Guest Editors)

Volume 28 (2009), Number 3

Direct Visualization of Deformation in Volumes

Stef Busking, Charl P. Botha and Frits H. Post

Delft University of Technology, the Netherlands

Abstract

Deformation is a topic of interest in many disciplines. In particular in medical research, deformations of surfaces and even entire volumetric structures are of interest. Clear visualization of such deformations can lead to important insight into growth processes and progression of disease.

We present new techniques for direct focus+context visualization of deformation fields representing transformations between pairs of volumetric datasets. Typically, such fields are computed by performing a non-rigid registration between two data volumes. Our visualization is based on direct volume rendering and uses the GPU to compute and interactively visualize features of these deformation fields in real-time. We integrate visualization of the deformation field with visualization of the scalar volume affected by the deformations. Furthermore, we present a novel use of texturing in volume rendered visualizations to show additional properties of the vector field on surfaces in the volume.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing I.4.7 [Image Processing and Computer Vision]: Feature Measurement—Feature Representation

1. Introduction

Comparisons play an important role in many scientific areas. When comparing one object to another, differences between the two can be interpreted as deformations which transform one object to the other. In a medical context, such deformations often correlate directly to growth processes or the progression of diseases. For this reason, analysis of deformations has become an important technique for medical researchers to understand these processes.

The analysis of deformation-fields with the purpose of studying morphological changes is called deformation-based morphometry, or DBM [AHF*98]. DBM is primarily promoted in brain-imaging research, where it is especially popular as it can be used to detect morphological differences over an entire brain. This is used, for example, to analyze all differences between subject brains and a standard brain in order to determine image-based characteristics that are associated with schizophrenia [GNB*01].

Another example illustrating the importance of volumetric changes in medical data is the study of rheumatoid arthritis (RA) and ostheoarthritis (OA). These are joint diseases that affect bone and cartilage in different ways. 3D MRI is increasingly being used to study the progression of these diseases over time. The diseases inflict progressively more damage on the affected joints in the form of, for example, bone erosions, bony outgrowths called osteophytes and changes in the cartilage. These characteristics are measured in MRI datasets [MSC*99, KKS*07] and are used to track the progression of the disease.

In all cases where 3D deformation fields have been used to study morphological changes (mostly brain imaging), only rudimentary visualization techniques have been applied. In general, aggregative metrics such as Jacobian-based volume change are calculated and used to typify differences. With suitable visualization techniques, differences can be studied in far more detail. In other areas where 3D image-based changes are being studied but deformation-fields are not yet being used, such as RA and OA progression, suitable visualization techniques should stimulate their introduction and hence facilitate these studies as well.

In this paper, we present new techniques for the direct focus+context visualization of 3D deformation fields, aimed at facilitating the study of volumetric change in medical re-

search. More specifically, they have been designed for use during the early exploration and hypothesis-generation stage of the medical research pipeline. Due to the nature of deformation fields, an ideal visualization of such data should provide insight into three things:

- the local behavior of the deformation in specific areas,
- the context of the deformations, in the form of the scalar volume that is being deformed,
- the effects of this deformation on this scalar volume

Our visualization approach fulfills these requirements: It enables the researcher to explore the deformation field in terms of areas with specific characteristics, such as growth or loss of tissue. Our techniques place these features within the context of the anatomical features present in the scalar volume, and also visualize how these features and the volume itself are affected by the deformations. In the next sections, we first discuss existing techniques for analyzing and visualizing deformation fields. In section 3 we present the details of our visualization. Finally, we present and discuss results obtained by applying our visualization to both synthetic and clinical data, and give directions for future research.

2. Related work

2.1. Deformation analysis

The currently accepted standard for analyzing deformation between 3D images (e.g., CT, MRI) uses non-rigid registration (e.g., [ABH*06, RCA*06]). Registration is the process of deriving a transformation under certain constraints, which aligns images to bring them into the same coordinate frame. A rigid registration limits this transformation to, for instance, an affine transformation, while non-rigid registration allows the registration to match free-form deformations. The most commonly used class of non-rigid transformations are B-splines [RSH*99, KSP07].

After registration, the deformation field which represents the transformation from one image's coordinate frame to the other holds important information about the differences between the two images. Traditionally, this deformation field is analyzed using statistical methods (e.g., [PHS*08]).

2.2. Vector field visualization

While direct visualization can provide valuable insight, e.g. for selecting analysis methods and areas of interest, few papers have explored visualization of deformation fields. As deformation fields are vector fields, they are often visualized using techniques from flow visualization. For instance, Tittgemeyer et al. [TWK02] proposed a visualization based on an estimation of critical points within the vector field. Color-mapped surfaces are used to show the local deformations relative to certain important features within the scalar volume. Riddle et al. [RLF*04] used 2D color-mapped Jacobians to show changes in tissue volume on separate slices of the dataset.

Our visualization of deformation fields is based on direct volume rendering (DVR). DVR is commonly used to visualize scalar volumes. Volume rendering has been used in combination with methods such as LIC [HA04] in order to visualize vector fields. However, such methods require a pre-processing stage to generate the LIC volumes. Similarly, Crawfis et al. [CMBC93] applied LIC, volumetric extensions of streamlines called flow volumes and other techniques to create scalar "density" volumes based on the vector fields. These were then rendered using traditional methods.

Direct rendering of the vector fields has not received much attention. In earlier work, Crawfis and Max [CM92] presented an approach where the vector image is filtered with an oriented kernel during rendering, thereby visualizing the vector direction as an oriented line segment. Frühauf [Frü93] visualized vector fields with DVR by mapping vector direction and magnitude directly to colors and opacities. Both approaches suffer from problems due to occlusion. In fact, dense 3D vector field visualization is considered an unsolved problem. Additionally, traditional flow and vector field visualization techniques are often not suitable for deformation visualization, as they do not provide enough context information to understand the deformations.

Compared to general vector fields, deformation fields such as those encountered in medical research have several properties which could be taken advantage of to improve their visualization. Depending on the nature of the data and problem under consideration, certain types of deformation may be considered more important than others. For example, changes in the amount of tissue are likely to be of greater interest to a researcher than movement of the patient between scans. Additionally, such features are likely to be local (e.g., growth of tumors or development of lesions). Because of this, visualizations of these features are less likely to suffer from occlusion problems.

Other techniques for analysis and visualization of deformation have been presented based on mechanical modeling and stress/strain tensor fields [HFH*04]. Computation of such fields operates under certain assumptions and implies an extra level of abstraction. Our visualization does not require pre-processing beyond the initial computation of the deformation field. Furthermore, we use techniques from importance-based visualization and non-photorealistic rendering to create a sparse but meaningful visualization of the field and its context.

3. Visualizing deformation

In this section, we present a focus+context visualization of deformation fields. While the process of obtaining the deformation field is outside the scope of this paper, this usually involves applying registration twice. Here, the first registration is used to remove differences from the data which are not considered relevant to the problem being studied. For ex-

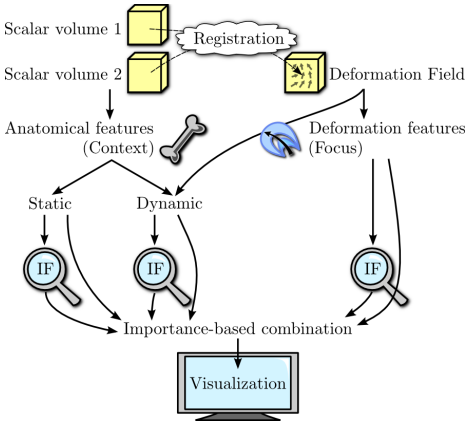

Figure 1: *Visualizing deformation fields*

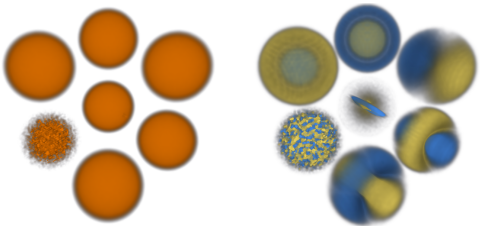

(a) Magnitude (b) Growth

Figure 2: *Visualizations of a synthetic* 64^3 *vector field. The left figure only shows vector magnitude (in orange), while the right figure shows our growth measure (yellow for negative values, blue for positive), giving a more detailed representation of deformation.*

ample, an affine registration can remove a difference in orientation of the patient between two scans. The second, non-rigid registration captures the remaining differences between the images in a deformation field. We share the assumption made in DBM that the deformation field for a large part reflects the actual physiological changes in the patient.

Figure 1 shows the components making up our visualization. The focus of the visualization consists of a direct rendering of features derived from the deformation field. We define measures that allow the user to highlight areas of the volume where the deformation has certain characteristics, e.g., growth or loss of tissue. Context is provided by simultaneous sparse volume rendering of the scalar volume. We distinguish between static and dynamic context in order to show both the surrounding anatomical features as well as the effects of the deformation. We present both interactive and static enhancements to the context visualization that show these effects.

The focus and context visualizations are integrated into a single visualization, such that the resulting combination presents relevant deformation features in an easily identifyable way *within* the context of the deforming volume. For this purpose, we define *interest functions*, which assign importance to features in each part of the visualization. Finally, we discuss the use of texture in direct volume rendering as a channel to provide additional information about the deformation vector field.

3.1. Deformation measures

The deformation fields our visualization deals with are given as vector-valued volumes. Traditional volume rendering uses a transfer function to map values in a volume to colors and opacities. Independently, however, the components of the vector field mean little. Furthermore, the design of transfer functions (let alone general three-dimensional transfer functions) is highly non-trivial. When examining deformations, a researcher is often interested in areas where the deformation has certain characteristics. Therefore, we define scalar valued measures that capture these characteristics. By visualizing these measures using volume rendering, we are able to highlight areas of the volume corresponding to the characteristics that are of interest to the user, without requiring explicit segmentation.

The most straightforward example of this is areas where the magnitude of the deformation exceeds a certain threshold. A direct volume rendering of vector magnitude in a synthetic deformation field is shown in figure 2(a). While this clearly shows areas of significant deformation, the nature of the deformation is not visible.

In analyzing vector fields, the Jacobian matrix J of the field is often used to gain insight into the local behavior of the field. The Jacobian is a matrix consisting of the three first-order partial derivatives of the field. The absolute value of its determinant, $\|J\|$, indicates the local volume change, and is therefore especially useful for analyzing deformations in medical data, as this directly corresponds to growth or loss of tissue. Based on the Jacobian determinant, we define the *growth* measure g:

$$g = \begin{cases} \|J\| - 1 & \text{for} \quad \|J\| < 1 \\ 1 - \frac{1}{\|J\|} & \text{otherwise} \end{cases} \tag{1}$$

Negative growth indicates loss of tissue. The measure is symmetric, and in the range $[-1,1]$, with $g = 0$ meaning no change in volume, $g = -1$ meaning total loss of volume, and $g = 1$ meaning infinite growth. Its absolute value represents local tissue change. We therefore visualize the growth measure by using this absolute value to determine opacity. Features are color coded to indicate whether they represent growth or shrinkage. Figure 2(b) shows a direct volume rendering of growth features in a synthetic dataset.

Because the Jacobian is a derivative of the field, we can

introduce the concept of scale [Wit83]. Approximation of derivatives is usually sensitive to noise in the data. By computing the Jacobian (and the measures derived from it) at a higher scale, the impact of such noise can be removed. Moreover, manipulation of this scale allows the user to easily filter the features in the dataset based on their scale. For this reason, computation of the Jacobian is done on the GPU, allowing for near-interactive manipulation of the scale parameter. We implemented derivative approximation using central differencing, where scale determines the spacing between samples. We also implemented derivative approximation based on convolution with a Gaussian derivative kernel. This yields higher quality approximations of the derivative with a smoothly adjustable scale, at the cost of performance.

While the growth measure is effective in showing areas where change in volume occurs, it fails to highlight areas of possibly significant deformation where the volume remains the same (e.g. translations). The simpler magnitude visualization shown in figure 2(a) does show these areas. However, these often occlude areas where the growth measure is significantly large, as well as potentially useful context structures (see section 3.2). To solve this problem, we use our importance-based filtering techniques presented in section 3.3 to combine the two visualizations, assigning higher importance to the growth features. This essentially highlights areas where deformation occurs without volume change. To avoid occluding growth features, these areas are not visualized directly. Instead, the magnitude measure is used in section 3.3 to determine areas where dynamic context information should be shown, which is sparse enough not to cause occlusion problems.

It should be noted that most non-rigid registration techniques result in a deformation field with vectors which for each point in a so-called *fixed image* point to the corresponding point in the *moving image*. Applying the growth measure to such a field assumes the semantics of going from the fixed image to the moving image. That is, interpreted directly, the field describes a transformation which "deforms" the fixed image. Determining the inverse transformation is often impossible, however, due to the symmetry of the growth measure, its sign can simply be reversed to reverse these semantics.

3.2. Providing context

Deformation features such as those visualized in the previous section are meaningless to a medical researcher unless they can be related to anatomical features. In our visualization, such context information is provided by a scalar volume. Often, this will be one of the volumes used to compute the deformation field. Assuming the registration is perfect, the second volume can be derived from the first by applying the deformation. Therefore, we only need a single volume to provide context information.

In visualizing the context volume, we distinguish between two kinds of structures. *Static context* consists of structures not significantly affected by the deformation. *Dynamic context* consists of those structures wich are deformed significantly, as well as the effects of the deformation on the volume itself. We use separate *modes* to visualize both types of context.

In order to minimize occlusion issues, we propose a sparse visualization of the context volume. As our visualization is aimed at medical researchers, we can assume advanced knowledge about the anatomical structures present in the volume. Therefore, it is not as important to show all information present in the scalar volume; such information can be better obtained using traditional volume visualization techniques. Because the deformations affect the shape of (anatomical) features within the volume, we focus on showing just these shapes.

Because values in CT or MRI scans are relatively uniform within separate tissues, high gradients often indicate boundaries between different types of tissue. Furthermore, the gradient can be used as an estimation of the normal of the boundary surface. To provide static context information, we only show parts of the boundary surface where this gradient is near-perpendicular to the viewing direction. This is achieved by volume rendering a *contourness* measure c, which is defined as follows:

$$c = s\,|\vec{n}| \left(1 - \max\left(1, \alpha \left\| \frac{\vec{n}}{|\vec{n}|} \cdot \vec{e} \right\| \right)\right) \tag{2}$$

where s is the scalar value at the current position, $\vec{n}$ is the scalar volume gradient, $\vec{e}$ is the unit vector pointing towards the camera and $\alpha \geq 1$ is a parameter controlling the sharpness of the contours. This results in a silhouette representation of the surfaces within the volume, shown in figure 3(b).

Visualizing dynamic context consists of two parts. First, we show the effects of the deformation on the structures shown in the static context mode. To achieve this, we can simply use the deformation field multiplied with a user-controllable value to offset sampling positions. This allows the user to interactively morph between the volume before and after deformation. We exaggerate deformations by allowing values outside of the range $[0,1]$. The resulting caricaturistic visualization [RVG06] can make small changes in the volume more apparent. For example, the silhouette contours in figure 3(b) can be smoothly morphed between the two datasets visualized, or even show exaggerated deformations, thus emphasizing their effects.

While morphing works well in an interactive setting, it requires the use of animation, which is not always available when communicating results. For this reason, we also explored static methods for visualizing the effects of the deformation. As a static alternative to morphing we overlay the contours from the deformed volume on those from the origi-

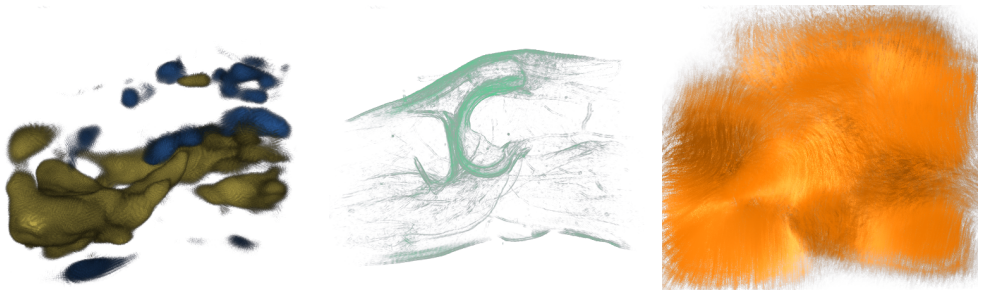

(a) focus (b) static context (c) dynamic context

Figure 3: *Deformation visualization elements*

nal volume. We use different colors for both sets of contours, so they can easily be distinguished.

The second part of our dynamic context visualization is a visualization of the effects of the deformation on the volume itself. For this, we use a dense vector field visualization technique similar to spot noise. For each position $\mathbf{p}$ we sample a precomputed volume containing uniform noise at multiple positions $\mathbf{p} + \delta\vec{v}$, where $\vec{v}$ is the deformation vector at $\mathbf{p}$ and δ ranges from 0 to a user-configurable maximum. This essentially smears out the noise volume along the vectors in the deformation field (see figure 3(c)).

3.3. Integrated visualization

The visualization described so far has three *modes*, examples of which are shown in figure 3:

Focus: visualization of features in the deformation field.
Static context: visualization of parts of the scalar volume unaffected by the deformation.
Dynamic context: visualization of the effects of the deformation on the volume.

These three modes could be used separately to visualize a given dataset. However, in that case relations between features seen in the different modes are not always apparent. We therefore integrate the modes into a single visualization. Our approach is to render all three modes simultaneously and combine results for each step through the volumes.

For each position in the volume, we select the mode that is most relevant for showing that particular point. We do this by defining *interest functions*, which assign importance to each point in the volume. The domain of these functions can be any value computed for this point in any of the modes. However, here we define interest functions which create a meaningful visualization, in which it is easy to distinguish areas with specific flow characteristics.

For the focus mode, the objects of interest are the areas where growth or shrinkage is high. We therefore use the absolute value of growth (equation 1) as the domain of the interest function. Defining the interest function similar to the opacity transfer function for this mode leads to good visibility of the features in the visualization. The interest function could be extended slightly outside the range of opacities to create an empty shell around growth features. This makes them stand out more in the resulting image.

As our static context consists of thin and very sparse contours, we can afford to assign these high importance. This way, they will always be visible, even if other features occupy the same space. Importance is again assigned similar to opacity, based on the measure defined in equation 2. A distinction between dynamic and static context can easily be made based on the local magnitude of the deformation field. We therefore use magnitude as the domain for the dynamic context interest function. We assign dynamic context features lower importance than focus features. This way, we highlight areas where deformation occurs without a change in the amount of volume / tissue.

During rendering we compute the values of these three interest functions for each step through the volume. The mode with the highest importance is selected and applied, resulting in a color and opacity. Finally, these colors and opacities are composited as in traditional direct volume rendering. In our prototype implementation, the interest functions consist of simple linear interpolations from 0 to 1 between double thresholds on the function's domain. These thresholds can be manipulated interactively by the user (as well as the normal transfer functions for each mode), allowing for easy manipulation of the relative importance of features in the visualization (see figure 4)

One important issue to consider when combining multiple modes in a single visualization is that it may not be clear to which mode a given feature belongs. To solve this issue, we use color only to distinguish between features with different semantics. Furthermore, the different modes each have their own visual style, to make them easily distinguishable.

3.4. GPU-based raycasting of scalar and vector fields

We base our techniques on our GPU-accelerated multi-volume raycaster. This raycaster uses a similar technique to

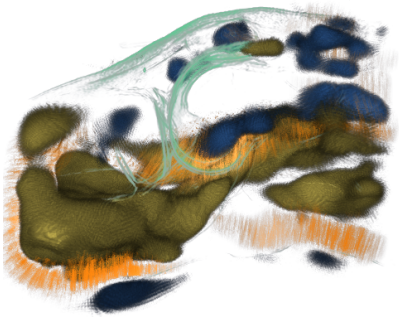

Figure 4: *Combined visualization of focus and context features in a* $512 \times 512 \times 80$ *knee dataset. Jacobians were computed with a kernel radius of* $4 \times 4 \times 0.6$ *voxels. Growth features are shown for* $|g| > 0.14$.

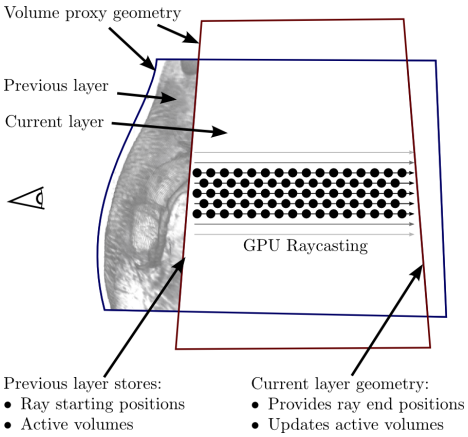

Figure 5: *Using a layered approach for GPU-based multi-volume raycasting*

that presented by Krüger and Westermann [KW03]. In short, we use a depth peeling approach to render proxy geometry for each volume one layer of geometry at a time. A fragment shader is used to perform raycasting between the layers. We keep track of active volumes for each layer, thereby enabling simultaneous multi-volume raycasting. The complete rendering process for two volumes is shown in figure 5.

Rays are cast for each layer from the positions on the previous layer to those on the current one, using a fixed step size (see figure 5). Active volumes are determined, loaded as textures, and sampled for each position along the ray by transforming the current position into their local coordinate frame. Both the measures described in section 3.1 and the visualizations described in section 3.2 are computed in real-time from the vector and scalar volumes on the GPU. The resulting colors and opacities are combined using the importance filtering techniques described in section 3.3, and then composited to form the resulting color and opacity for the layer. To reduce the visual impact of artifacts caused by the fixed step size, we offset the sampling positions of the rays in neighboring pixels by fractions of the step size. While the result has a dithered appearance, it leads to better visibility of structures within the volume, even for larger step sizes, which in turn lead to better performance.

Optionally, gradients can be computed in real-time using central differencing. These can then be used for lighting, allowing for better perception of shapes in the resulting visualization. Due to the computational complexity, however, computing gradients of growth features may be too expensive for real-time rendering. A simple alternative, used in the figures in this paper, is to use a limb darkening effect similar to that used in [HA04]. Limb darkening can be achieved by using a softer transfer function to create a darker halo around surfaces, providing the effect of shading without needing to compute surface normals.

A different technique we explored is using streamline shading [Frü93] to visualize deformation direction. However, this technique only works well on thin, line-like objects oriented in the direction of the deformation. As the features in our focus mode generally form blob-like surfaces, streamline lighting creates a confusing appearance and does not provide intuitive insight into the deformation direction. For this reason, we apply this form of lighting only to the smears in our dynamic context visualization, which do have a clear linear structure.

3.5. Using texture

While texture hardware is often used for volume rendering, actual texturing is not commonly used in volume visualization due to the complexities of mapping two-dimensional images on arbitrary surfaces. However, due to the simplistic and (at least at larger scales) relatively smooth nature of the surfaces in our visualization, and due to our choice of single colors for different features, texture forms an available channel for presenting additional information to the user.

We use texture to indicate the local direction of the deformation vector field. Because this direction is obviously three-dimensional in nature, we split the vector in a surface-tangential and a perpendicular component. The tangential direction is visualized using an oriented stripe texture. The angle between the vector and the surface is visualized in the color of the stripes. Similarly, color can be used to visualize vector magnitude. Figure 6 shows the same feature both without and with several variations of our texturing technique.

Because the surfaces of features might be noisy at lower scales, computing the texture directly on the tangent plane results in a noisy image where the direction may not be clearly visible. Instead, we project the deformation vector

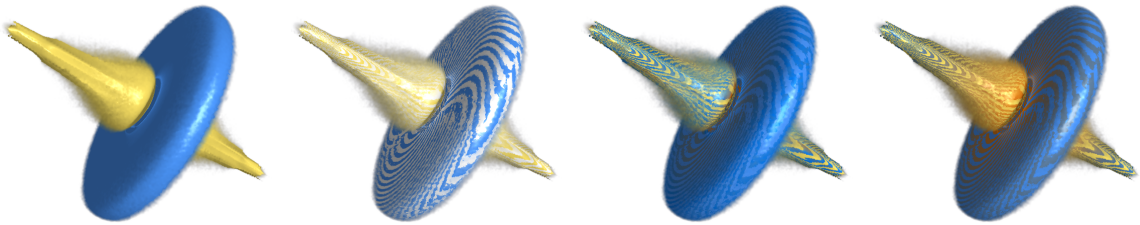

(a) No texture (b) Direction (c) Perpendicularity (d) Magnitude

Figure 6: *Enhancing the visualization with texture reveals various properties of the deformation field, such as a rotational component in this synthetic saddle point deformation (128^3). Colors in the last two figures go from blue (low) to orange (high).*

to the axis-aligned plane most similar to the tangent plane, and compute the texture based on this vector. Next, a striping pattern is applied along the vector perpendicular to the projected deformation vector in the plane. The width of the stripes can be changed by the user to adapt the texture to vector fields with different smoothness.

4. Results

The algorithms described in this paper were implemented in C++ using OpenGL and the GLSL shading language for GPU programming. The visualization runs with interactive performance on Nvidia GeForce GTX 280 hardware, achieving an average framerate of 10 fps on the visualization shown in figure 7. Applying texture or lighting to the focus features is more computationally intensive, as this requires estimation of the surface normal for these features. Given current trends in GPU capabilities, however, performance of these techniques is likely to improve significantly with next generation hardware. The current implementation is limited to volumes and deformation fields that fit in GPU memory. Out-of-core techniques could be used to remove this limitation.

Results from applying our visualizations to artificially generated deformation fields are shown throughout section 3. These vector fields were generated by additive combination of simple vector fields representing rotations, saddles, translations, sources and/or sinks [WH91]. While we applied our visualization to datasets from an OA study (shown in figure 3 and figure 4), these images are intended only to illustrate our techniques as we have yet to work with domain experts to validate these results. In addition to these datasets, we used a pair of MRI images, taken 4 months apart, of the brain of a patient suffering from Multiple Sclerosis. MS leads to the appearance of lesions in the white matter of the brain. The data shows a high number of such lesions, some of which grow, shrink, appear or disappear between the two scans. The images had previously been normalized. As is common in deformation analysis research, we used elastix [KSP07] to apply a standard B-spline non-rigid registration between the two images in order to obtain a deformation field. We then applied our techniques to visualize this field together with one of the original MRI images.

Our visualization, shown in figure 7, shows that changes in the brain clearly appear as growth and shrinkage features. Most of these features are yellow, which corresponds to contraction or loss of tissue. Experts hypothesize that the increasing number of lesions leads to a reduction in the volume of the brain ventricles. The yellow features in the center of the brain seem to be in line with this hypothesis, although further analysis is required to confirm this. The visualization also shows a large deformation outside the brain in the lower left corner, which is probably due to misalignment. The yellow feature on the far left is an artifact, resulting from padding the dataset.

5. Conclusions and future work

We presented new techniques for the focus+context visualization of deformation in volumes, useful in studying development and disease progression in medicine. Our techniques directly visualize deformation fields in terms of the nature and effects of the deformations. In summary, we identify the following contributions:

- direct visualization of features derived from deformation fields, allowing exploration of the fields on a higher semantic level than traditional vector visualizations,
- direct volume rendering of the underlying scalar volume, to show both *context* and *effect* of the deformations,
- importance-based integration of both visualization techniques into a single image,
- the use of texture to present additional information about the deformation field in relation to the features shown in the visualization.

While our focus in this research was on medical datasets, the techniques presented in this paper can be applied in other areas as well. In future work, we intend to refine these techniques and work with domain experts, most notably in OA research, to apply our approach to real data. For this, new types of measures may need to be developed to capture characteristics of interest to these experts. Additionally, we aim

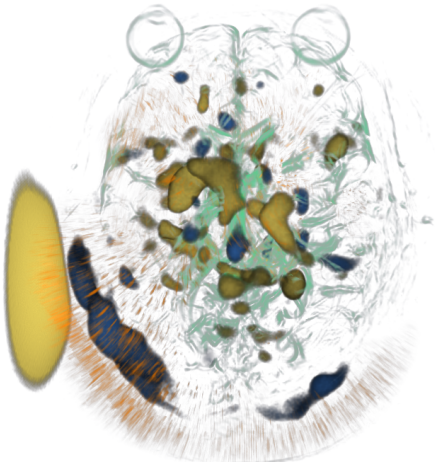

Figure 7: *Visualization of deformation due to changing white matter lesions in a brain MRI dataset of* $181 \times 217 \times 179$ *voxels. Jacobians were computed with a kernel radius of 2 voxels. Growth features are shown for* $|g| > 0.09$.

to extend our filtering techniques to allow a user to select features based on various criteria other than scale. Finally, we intend to extend our techniques to allow comparison of two or more deformation fields.

Acknowledgements

We are grateful to Dr. G. Kloppenburg and Dr. I. Watt of respectively the Reumatology and Radiology departments of the Leiden University Medical Center for their valuable input and OA research datasets. This research is supported by the Netherlands Organisation for Scientific Research (NWO), project number 643.100.503 "Multi-Field Medical Visualisation".

References

[ABH*06] ALJABAR P., BHATIA K. K., HAJNAL J. V., BOARDMAN J. R., SRINIVASAN L., RUTHERFORD M. A., DYET L. E., EDWARDS A. D., RUECKERT D.: Analysis of growth in the developing brain using non-rigid registration. In *IEEE Biomedical Imaging: Nano to Macro, Proceedings* (2006), pp. 201–204.

[AHF*98] ASHBURNER J., HUTTON C., FRACKOWIAK R., JOHNSRUDE I., PRICE C., FRISTON K.: Identifying global anatomical differences: Deformation-based morphometry. *Human Brain Mapping 6*, 5-6 (1998), 348–357.

[CM92] CRAWFIS R., MAX N.: Direct volume visualization of three-dimensional vector fields. In *Symposium on Volume Visualization, Proceedings* (1992), pp. 55–60.

[CMBC93] CRAWFIS R., MAX N., BECKER B., CABRAL B.: Volume rendering of 3D scalar and vector fields at LLNL. In *Supercomputing '93, Proceedings* (1993), pp. 570–576.

[Frü93] FRÜHAUF T.: Raycasting vector fields. In *IEEE Visualization, Proceedings* (1993), pp. 115–120.

[GNB*01] GASER C., NENADIC I., BUCHSBAUM B. R., HAZLETT E. A., BUCHSBAUM M. S.: Deformation-based morphometry and its relation to conventional volumetry of brain lateral ventricles in mri. *NeuroImage 13*, 6 (2001), 1140 – 1145.

[HA04] HELGELAND A., ANDREASSEN O.: Visualization of vector fields using seed LIC and volume rendering. *IEEE Transactions on Visualization and Computer Graphics 10*, 6 (2004), 673–682.

[HFH*04] HOTZ I., FENG L., HAGEN H., HAMANN B., JOY K., JEREMIC B.: Physically based methods for tensor field visualization. In *IEEE Visualization, Proceedings* (2004), pp. 123–130.

[KKS*07] KORNAAT P., KLOPPENBURG M., SHARMA R., BOTHA-SCHEEPERS S., LE GRAVERAND M., COENE L., BLOEM J., WATT I.: Bone marrow edema-like lesions change in volume in the majority of patients with osteoarthritis; associations with clinical features. *European Radiology 17*, 12 (2007), 3073–3078.

[KSP07] KLEIN S., STARING M., PLUIM J. P. W.: Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. *IEEE Transactions on Image Processing 16*, 12 (2007), 2879–2890.

[KW03] KRÜGER J., WESTERMANN R.: Acceleration techniques for GPU-based volume rendering. In *IEEE Visualization, Proceedings* (2003).

[MSC*99] MCQUEEN F. M., STEWART N., CRABBE J., ROBINSON E., YEOMAN S., TAN P. L. J., MCLEAN L.: Magnetic resonance imaging of the wrist in early rheumatoid arthritis reveals progression of erosions despite clinical improvement. *Ann Rheum Dis 58*, 3 (1999), 156–163.

[PHS*08] PIEPERHOFF P., HÖMKE L., SCHNEIDER F., HABEL U., SHAH N. J., ZILLES K., AMUNTS K.: Deformation field morphometry reveals age-related structural differences between the brains of adults up to 51 years. *Journal of Neuroscience 28*, 4 (2008), 828–842.

[RCA*06] RUECKERT D., CHANDRASHEKARA R., ALJABAR P., BHATIA K. K., BOARDMAN J. P., SRINIVASAN L., RUTHERFORD M. A., DYET L. E., EDWARDS A. D., HAJNAL J. V., MOHIADDIN R.: Quantification of growth and motion using non-rigid registration. *Lecture Notes in Computer Science 4241* (2006), 49–60.

[RLF*04] RIDDLE W. R., LI R., FITZPATRICK J. M., DONLEVY S. C., DAWANT B. M., PRICE R. R.: Characterizing changes in MR images with color-coded jacobians. *Magnetic Resonance Imaging 22*, 6 (2004), 769 – 777.

[RSH*99] RUECKERT D., SONODA L. I., HAYES C., HILL D. L. G., LEACH M. O., HAWKES D. J.: Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging 18*, 8 (1999), 712–721.

[RVG06] RAUTEK P., VIOLA I., GRÖLLER M. E.: Caricaturistic visualization. In *IEEE Visualization, Proceedings* (2006), vol. 12, pp. 1085–1092.

[TWK02] TITTGEMEYER M., WOLLNY G., KRUGGEL F.: Visualising deformation fields computed by non-linear image registration. *Computing and Visualization in Science 5*, 1 (2002), 45–51.

[WH91] WEJCHERT J., HAUMANN D.: Animation aerodynamics. In *SIGGRAPH, Proceedings* (1991), pp. 19–22.

[Wit83] WITKIN A. P.: Scale-space filtering. In *International Joint Conference of Artificial Intelligence, Proceedings* (1983), pp. 1019–1021.