

Digitisation and 3D Reconstruction of 30 Year Old Microscopic Sections of Human Embryo, Foetus and Orbit

Joris E. van Zwieten¹, Charl P. Botha¹, Ben Willekens²,
Sander Schutte³, Frits H. Post¹, and Huib J. Simonsz⁴

¹ Data Visualisation Group, Delft University of Technology

² The Netherlands Ophthalmic Research Institute, Amsterdam

³ Biomechanical Engineering, Delft University of Technology

⁴ Ophthalmology Department, Erasmus Medical Centre, Rotterdam

Abstract. A collection of 2200 microscopic sections was recently recovered at the Netherlands Ophthalmic Research Institute and the Department of Anatomy and Embryology of the Academic Medical Centre in Amsterdam. The sections were created thirty years ago and constitute the largest and most detailed study of human orbital anatomy to date. In order to preserve the collection, it was digitised. This paper documents a practical approach to the automatic reconstruction of a 3-D representation of the original objects from the digitised sections. To illustrate the results of our approach, we show a multi-planar reconstruction and a 3-D direct volume rendering of a reconstructed foetal head.

1 Introduction

Recently, a collection of 2200 microscopic sections from 2 human embryos, 7 human foetuses, and 3 adult *orbits* was recovered at the Netherlands' Ophthalmic Research Institute and the Department of Anatomy and Embryology of the Academic Medical Centre in Amsterdam.

The *orbit* is the bony cavity that contains the eye and its appendages, i.e. the eye socket. A *microscopic section* is a very thin slice of tissue, laid flat on a glass slide, stained, mounted in a medium of proper refractive index and covered with a very thin piece of glass called a coverslip [1]. Figure 1 shows an example of two microscopic sections of a 79 mm (crown to rump length) human fetus from the recovered collection.

The collection, henceforth referred to as the *orbita collection*, forms the heritage of the late J.A. Los, L. Koornneef, and A.B. de Haan. The sections were created between 1972 and 1986. Koornneef studied the development and the adult configuration of previously unknown connective tissue *septa*, or thin partitions, found within the orbital fat. De Haan studied the development of the bones that compose the orbit. Especially the work by Los and Koornneef has become well known among anatomists worldwide. The collection of sections itself can be considered one of the largest and most detailed studies of orbital anatomy to date.

When the sections were recovered it was noticed that the dyes used for staining had started to fade. To preserve the collection, all the sections were digitised at an optical

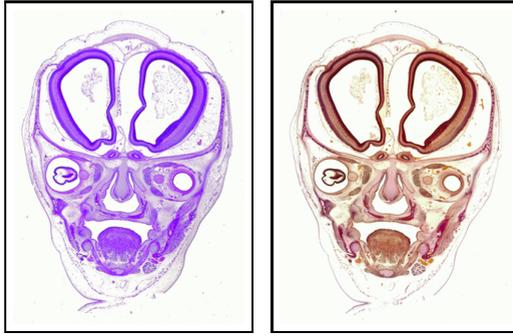


Fig. 1. Two adjacent sections from series S73345 that illustrate the different dye combinations used alternatingly in this series: on the left haematoxylin-eosin and on the right haematoxylin-Van Gieson's picro-fuchsin

resolution of 2500 dpi. Subsequently, 3-D datasets were constructed from the digitised sections. Thirty years after the microscopic sections were created, 3-D representations of the original specimen could now be studied. Both the digitised sections and the 3-D reconstructions will be made publicly available via the Internet.

In this paper, we detail the preservation and reconstruction of this historically important dataset. Our contributions are:

- We present a practical and complete procedure to reconstruct a 3-D representation of the original objects from a collection of thirty year old microscopic sections.
- Most registration similarity metrics are based on scalar data. In order to make use of these metrics, the colour sections have to be converted to a scalar representations. We compare traditional grey-scaling and a more refined PCA-based method with regards to their suitability for this task.
- A number of the datasets from the orbita collection consist of sections that have been stained using different dye combinations alternatingly (see figure 1). We investigate the effect of these alternating dye combinations on the similarity metric and show that the scalar conversion contributes to the colour invariance of our approach.

2 Related Work

To our knowledge, the most recent survey of literature on the reconstruction from microscopic sections dates back to 1986 [2]. Since then several other surveys on registration were published [3,4], but none of these deal specifically with microscopic sections.

An interesting “new” paradigm for the reconstruction from microscopic sections is *episcopic-diascopic* reconstruction [5,6,7]. Before a section is cut, a photograph is taken of the embedding block. This photograph is called the *episcopic*, or see-on, image. After cutting, a section is stained, mounted, and digitised. The result is called the *diascopic*, or see-through, image. Thus, each section is represented by a pair of images. Of course, episcopic images are free of the deformation that is introduced by histological processing and therefore represent a valid reference of shape. Non-rigid registration

of each diascopic image onto its corresponding episcopic image guarantees that the anatomical differences between adjacent sections are preserved. This is an important advantage over methods that apply non-rigid registration to pairs of adjacent diascopic sections. Additionally, registration error only affects a single section and does not propagate through the stack.

Unfortunately, no episcopic images were recorded when the orbita collection was created. In this case, the canonical approach is to register pairs of adjacent sections [8,9,10]. The authors of [10] cluster the displacement vector field and estimate an affine transform for each cluster, which are subsequently embedded into one non-rigid transform. Their approach is to tailor the transform space to the observed deformation pattern. This helps to avoid transforms that are known to be wrong from the outset. Instead of performing registration on pairs of adjacent sections, some authors register a section with respect to both adjacent sections [11,12]. This is an interesting method, because it encourages transforms that yield smooth structures. If the original structures are known to be smooth with respect to the section thickness, one can assume that rapid changes in the direction of structures are a consequence of histological processing and should therefore be suppressed.

3 Preparation of Microscopic Sections

To study a specimen the size of a human orbit, embryo, or foetus under a light microscope it has to be *sectioned*, i.e. cut into thin sections. The field of *microtechnique* is concerned with the preparation of biological material for microscopic observation. This section provides a short introduction to the field and its terminology (derived from [13]) as well as a description of the specific preparation of the orbita sections.

The creation of microscopic sections from a biological specimen can be divided into five steps: fixation, embedding, sectioning, staining and mounting. The aim of *fixation* is to terminate life processes quickly and preserve the organisation of cells with as little deformation as possible. The canonical fixative is *formaldehyde*, which was also used to fixate the specimens that compose the orbita collection.

A specimen is *embedded* to support it during sectioning. Three well-known embedding methods are: paraffin embedding, nitrocellulose embedding, and freezing, of which paraffin embedding is the most common. The orbita collection was embedded in nitrocellulose.

Historically, the *sectioning* of a specimen was performed by hand using a razor. Much better results can be obtained by using a *microtome*. Microtomes are machines in which a block of embedded tissue and a knife can be fixed. To cut a section, either the knife is drawn across the embedding block, or the block is drawn over the knife.

The various constituents of tissue (e.g. cytoplasm, connective tissue, and lipids) contained in a section can be discriminated by *staining*. Dyes used for staining are often applied in pairs: a basic dye, which stains constituents with predominantly acid groupings (also termed *stain*), and an acid dye (also termed *counterstain*).

Several combinations of dyes were used to stain the orbita collection sections (see figure 1). Koornneef used haematoxylin-azophloxin, while De Haan used haematoxylin-eosin (H-E). Both used haematoxylin-Van Gieson's picro-fuchsin. We will discuss these

dyes in short. Because azophloxin and eosin yield highly similar results, they can be assumed identical for our purposes.

Alum haematoxylin is a basic dye. This dye stains nuclei deep blue or purplish, hyaline cartilage and bone is stained light to medium blue. Eosin is a red acid dye, which stains cytoplasm, muscle fibres, blood corpuscles and connective tissue. The combination of haematoxylin and eosin displays a general histological overview. Van Gieson's picro-fuchsin is an acid dye, which can be used as a counterstain for haematoxylin. It stains collagenous connective tissue, muscle, and keratin. Applied to embryos it produces a good distinction between bone and cartilage.

A section is *mounted* to preserve it and to enhance the visibility of the structures it contains. A section is laid down on a glass slide, a suitable mounting medium is applied, and a coverglass is placed over the section.

4 Digitisation

Digitisation of the orbita sections was performed at the Netherlands Ophthalmic Research Institute. Each glass slide was placed into a bracket that was custom made to fit an ArtixScanTM 4500t film scanner from Microtek.

This scanner is designed specifically to scan photographic film and slides. It has an optical resolution of 2500 dpi at a bit depth of 14 bits per colour. Scanners are often used to scan objects that are already records of three components of colour information and therefore might use narrow band filters. If a natural object such as a histological section is scanned with such a scanner the reported RGB values will be inaccurate [14]. However, according to the specifications of the scanner, the employed filters are not of the narrow band type.

5 Reconstruction

Reconstruction of 3-D datasets from the digitised microscopic sections consists of a pre-processing phase and a registration phase. In the pre-processing phase, the digitised images are prepared for registration. During the registration phase, pairs of adjacent images are registered onto each other. In other words, for each pair a transformation is sought that fits one image of the pair (the *floating* image) onto the other image (the *reference* image). The result of reconstruction is a 3-D dataset.

5.1 Pre-processing

The pre-processing phase consists of resampling, segmentation, and an RGB to scalar data conversion.

Resampling. Firstly, the resolution of the sections is decreased by *resampling*. Reconstruction of a resampled series takes less time, which facilitates fast prototyping. Reconstruction of the resampled data can be interpreted as the first step of a progressive refinement algorithm: the solution computed for the resampled data can be used as a starting point for reconstruction of the high resolution data. All sections were resampled to a standard width of 512 pixels while maintaining the original aspect ratio.

Resampling was performed with a lanczos kernel. This is a windowed sinc kernel that is often used for high quality image resampling.

Segmentation. Segmentation is the process of separating an object of interest from the background in an image. It is relevant to reconstruction, because a reconstruction should only be based on the object of interest (and/or external fiducials). As a corollary: if the background is largely homogeneous, segmentation may be unnecessary. However, even if segmentation is unnecessary for accurate reconstruction, it may still speed up computation because background pixels can be ignored.

The output of the segmentation procedure is a binary mask of the object. An initial segmentation is performed by exploiting colour differences between the background and the object of interest. Separation was found to be easier in HSV (*Hue, Saturation, Value*) colour space than in RGB (*Red, Green, Blue*) colour space. A colour threshold is derived from the histogram. We experimented with two automatic threshold selection methods: the triangle method and the background symmetry method, see [15, pp.95-96]. The first was used if the histogram contained a single background peak, the latter if it contained both a background and a foreground peak. However, in our experience, manually selection of a fixed threshold yields the best results.

The initial thresholding marks holes within the object as background. However, as these holes are generated by the object they can be considered as features for reconstruction. Therefore, the segmentation procedure was extended to close the holes in the initial mask. First, binary dilation is applied to close the object's outer contour. Next, binary propagation is used to remove unconnected background artifacts. Finally, inversion followed by binary propagation of the background and another inversion closes any remaining holes in the mask. Figure 2 illustrates the various steps of this segmentation procedure.



Fig. 2. From left to right, the five steps of the segmentation procedure: input section, thresholding on hue, binary dilation to close contour, background artifact removal, filling the remaining holes in the mask

Conversion to Scalar Data. An essential component of a registration algorithm is the *similarity metric* which quantifies the *goodness-of-fit* or similarity of two adjacent sections. Most of these metrics are designed for *scalar* data, i.e. one value per position.

Microscopic sections are treated with dyes to differentiate the anatomical structures of interest and therefore colour provides information about the tissue comprising the section. Colour information is inherently *multi-valued* instead of scalar. Existing literature on the registration of microscopic sections describe similarity metrics designed for

scalar data and apply these metrics to intensity images. Of course, intensity by itself does provide information about anatomical structure, which may well be sufficient to attain accurate registration. However, there are examples of applications where colour information can be used effectively to enhance the result, e.g. edge detection.

Instead of extending a similarity metric to account for colour data, we decided to reduce the dimensionality of the data to allow the use of similarity metrics designed for scalar data while minimising the loss of information.

A standard conversion to greyscale can be viewed as a mapping:

$$f(r, g, b) = \frac{r}{3} + \frac{g}{3} + \frac{b}{3} = [r \ g \ b] \cdot \left[\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3} \right]^T$$

A colour is mapped onto the dot product between that colour and the vector $\left[\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3} \right]^T$. In geometrical terms, f projects the RGB colour space onto the line through $\mathbf{0}$ with direction vector $\left[\frac{1}{3} \ \frac{1}{3} \ \frac{1}{3} \right]^T$ (and scales the result with the length of this vector).

Of course, we can project the colour space onto a line in any given direction. In fact, we can use *Principal Component Analysis* or PCA to find the direction that retains the largest part of the total variance compared to all other directions. Table 1 summarises the results of this approach. Notice how the blue component is almost completely ignored for the eosin sections, because it is approximately constant. Also, although Van Gieson's dye displays several hues on the orbita sections, i.e. purple-red-orange-yellow, these colours approximately lie along a line in the RGB colour space. For the Van Gieson sections, both PCA and straightforward conversion to grey scale retain a large part of the total variance. In general, however, mapping colours onto the first principal direction will preserve more variance than conversion to grey scale. Therefore, this method of scalar conversion is recommended.

Table 1. The retained variance as a percentage of the total variance and the axis corresponding to either greyscale or the first principal component. We randomly selected 75 sections of each dye combination from series S73345 as input data.

Stain	Method	Retained variance	Direction	Spherical variance
Eosin	GREY	63.11% ($\sigma = 3.51$)	0.5774 0.5774 0.5774	N/A
Van Gieson	GREY	93.57% ($\sigma = 1.87$)	0.5774 0.5774 0.5774	N/A
Eosin	PCA	95.03% ($\sigma = 1.83$)	0.4473 0.8930 0.0507	0.0011
Van Gieson	PCA	97.61% ($\sigma = 0.59$)	0.3939 0.6669 0.6325	0.0014

The spherical variance¹ of the axes found by PCA is small (see table 1). In other words, variance within a group of sections stained with a single dye combination (e.g. variance in staining time and fading of the sections) does not have much effect on the direction of the axis of largest variance. Therefore, the average direction could be used to convert all the sections of a series that were stained with a single dye combination.

¹ Spherical variance is a measure of the dispersion about a mean direction, which has a range of $[0, 1]$ (see [16, pp.218,219]).

5.2 Registration

In the registration phase, pairs of digitised sections are transformed to fit onto each other. The selected *transformation family* determines the degrees of freedom of the transformation. A *similarity metric* is used to quantify the quality of the fit. An *optimisation method* searches the space of transforms to find a transform that maximises the similarity metric. In the following subsections, we discuss these three components with the registration of the orbita collection in mind.

Transformation Family. As argued by [10, p.2] the selection of a transformation family should be guided by a priori knowledge of the deformation caused by preparation. If the transformation family is too constrained it cannot adequately express the deformation, if it is too general it may introduce solutions that are known to be incorrect.

The sections related to the study of Koornneef were originally created to investigate the spatial relations between newly discovered septa and the bony boundaries of the orbit. However, the preparation of the sections is quite favourable for 3-D reconstruction. Firstly, nitrocellulose was used as embedding medium. Deformation of material embedded in nitrocellulose during sectioning is relatively small compared to commonly used alternatives such as paraffin. Deformation during embedding is also less prominent because the procedure does not require heat. Secondly, decalcification was performed with E.D.T.A., a procedure that while slower causes less damage than commonly used alternatives. Thirdly, relatively thick sections were cut to enhance the visibility of the septa. Thicker sections generally deform less easily. Visual inspection of the sections revealed a nearly circular eyeball, which supports the idea that the sections suffered little overall deformation [17,13]. The sections associated with the study of De Haan were created to investigate the development of the orbit. The material was also embedded in nitrocellulose. Although nitric acid was used instead of E.D.T.A. for decalcification, visual inspection revealed only minor deformations [18].

Because section deformation should be minor given the preparation procedure we decided to use rigid transforms for registration. More general transforms could remove differences between sections that are a consequence of anatomy [9, p.30]. Furthermore, after performing reconstruction based on rigid transforms it is easier to judge the extent of non-rigid deformation in the sections.

Similarity Metric. A similarity metric should be chosen that reflects the physical relation between the image intensities² [19]. However, for the orbita collection this relation is unknown.

To investigate the relation between the image intensities we *manually* registered a small number of selected section pairs from series S73345 and plotted a joint histogram of the overlaying intensities (see figure 3). The sections in series S73345 are alternately stained with haematoxylin-eosin and haematoxylin-Van Gieson's dye. The histogram in figure 3 (left) was created from a pair of sections that are not directly adjacent but one section apart. The figure therefore shows the relation between a pair of sections

² In this paper, the word *intensity* refers to the value obtained by converting colour data to scalar data using PCA.

stained with a single dye combination. We also investigated the relation between section pairs stained with alternating dye combinations, see figure 3 (middle and right). However, as the results were similar, we will focus on the first case.

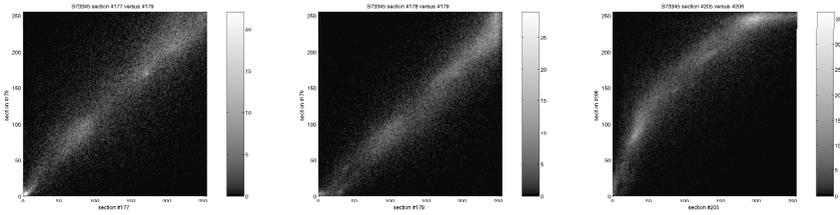


Fig. 3. Joint histogram of intensities after manual registration: (left) a pair of same-dye sections (H-E), (middle) a pair of sections stained with alternating dye combinations (H-E, H-Van Gieson), (right) a pair of sections stained with alternating dye combinations (H-Van Gieson, H-E). Values in the joint histogram that lie further than ten standard deviations from the mean have been clamped to provide a clearer picture.

From figure 3 (left) the intensities seem to approximate a linear relation. The correlation coefficient squared (r^2) equals 0.695 for the haematoxylin-eosin pair and 0.729 for the haematoxylin-Van Gieson's dye pair. In other words, 68.9% and 72.6% of the total variance can be explained by linear regression respectively. We repeated the same procedure on a pair of sections from series S74161, which is stained exclusively with haematoxylin-eosin, and found $r^2 = 0.595$.

The strength of these linear relations is only moderate, especially for the section pair from series S74161. Figure 4 shows the optimisation surfaces generated both by a similarity metric based on *normalised correlation* and one based on *mutual information* for this section pair. The normalised correlation metric implicitly assumes a linear intensity relation, while the mutual information metric assumes a much more general, statistical relation.

Both similarity metrics have a distinct minimum that approximately coincides with the manual alignment. It appears that the linear intensity relation is strong enough to be detected reliably by the normalised correlation metric.

The normalised correlation metric has a larger capture range than the mutual information metric. Also, it yields a 'smoother' optimisation surface, because it is computed on the entire overlap domain while the mutual information metric is computed on a random sample of points from the overlap domain [20]. Because a smooth optimisation surface and a large capture range are both desirable properties for a similarity metric, we decided to use the normalised correlation metric for section pairs stained with a single dye combination.

Figure 1 shows a pair of adjacent sections stained with different dye combinations. Two things are apparent from this figure. First of all, the two dye combinations mostly differentiate the same tissues. Intuitively one would expect this to be advantageous for registration. Secondly, although the two dye combinations produce different hues, the pattern of colour saturation seems roughly the same.

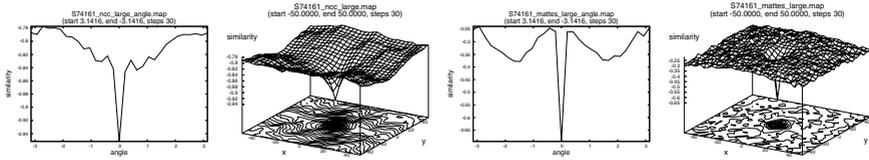


Fig. 4. Optimisation surfaces generated by the normalised correlation metric and by mutual information on a pair of manually aligned sections from series S74161. The two surfaces on the left show similarity over rotation angle and over translation for normalised correlation. The two on the right show similarity over rotation angle and over translation for mutual information.

We plotted a joint histogram for several pairs of sections stained with alternating dye combinations, as described for section pairs stained with a single dye combination. The joint histogram in the middle of figure 3 shows an approximately linear ($r^2 = 0.817$) relation, while the relation shown by the histogram on the right seems non-linear ($r^2 = 0.740$). Although the normalised correlation metric is designed to detect linear intensity relations, it still shows a distinct minimum and a large capture range for these manually aligned section pairs. Therefore, we also used the normalised correlation metric to register section pairs stained with alternating dye combinations. In this case, scalar conversion contributes to the colour invariance of our registration approach.

Optimisation Method. The papers we reviewed mostly featured optimisation algorithms that do not explicitly compute the partial derivatives of the similarity metric with respect to the independent variables. Specifically, the downhill simplex method by Nelder and Mead and Powell’s direction set method seem to be popular. (See [21, pp.408–420] for a description of these algorithms.)

However, it is possible to compute an approximation of the partial derivatives of the normalised correlation metric and the mutual information metric. Algorithms that make use of derivative information often converge in fewer iterations. This does not necessarily mean that they are always faster, because more computation is needed per iteration [21, p.395].

We used a popular adaptive step size gradient descent optimiser. During each iteration, this optimiser takes a step in the direction opposite to that of the local gradient. If the angle between the local gradient at the new position and the local gradient at the old position is larger than 90 degrees, the optimiser halves the step size.

6 Results

We have reconstructed almost all volumes from the orbita collection. For the purpose of this section, we focus on two series that will be referred to as S73345 and S74161. The former consists of 250 sections and the latter of 97 sections. For some of the measurements, we worked with a 100-slice subvolume of S73345. Benchmarks were performed on a pentium 4, 2.66 GHz, with 1 GB of ram.

In an attempt to speed up the reconstruction we used the centres of gravity of the binary *segmentation masks* to estimate an initial registration transformation. The initial



Fig. 5. Cross-sectional view of original and registered microscopic sections at top and bottom respectively. Qualitatively, the axial contours of the original 3-D objects seem to have been recovered.

Table 2. An overview of the number of iterations and total time required to perform reconstruction of various series, with and without using the binary segmentation masks to estimate an initial transform

Series	No. of sections	Masks	No. of iterations	Time
S73345-subvol	100	NO	15935	190 min
S73345-subvol	100	YES	3091	23 min
S73345	250	YES	10190	73 min
S74161	97	NO	26350	187 min
S74161	97	YES	2779	22 min

transformation translates the centre of gravity of the floating image onto the centre of gravity of the reference image. Table 2 shows the performance improvement resulting from this modification. In the two documented cases we observed speed-ups of a factor of 5 and 9 respectively. The centre of gravity calculation is rather sensitive to outliers.



Fig. 6. Direct volume rendering of the head of a human foetus created from a reconstructed dataset. The section-like variations visible are due to small differences in the image intensity that show up due to the specifics of the transfer function.

However, the object of this series (the head of a human foetus) has a very regular shape and therefore an accurate initial transformation could be obtained in most situations.

Figure 5 shows a virtual section along the sagittal plane, perpendicular to the plane of sectioning, before and after reconstruction. Figure 6 is a direct volume rendering of the same reconstruction.

7 Conclusions and Future Work

In this paper, we presented a practical procedure for the 3-D reconstruction of a large set of 30 year old microscopic sections. This procedure consists of a pre-processing and a registration phase. In the pre-processing phase, the data is resampled, segmented, and converted to scalar data. We investigated the application of principal component analysis (PCA) for this conversion. Converting to scalar data allowed the use of standard similarity metrics in the registration phase. In addition, it facilitated the registration of consecutive sections stained with alternating dye combinations.

As part of the registration phase, we discussed the selection of a transformation family based on the characteristics of the data. We also investigated the relation between adjacent sections stained with the same and with different dyes.

Finally, we showed a multi-planar view and a direct volume rendering of a reconstructed foetal head to illustrate the results of our procedure. Visual inspection of the reconstructions indicates a high overall quality.

An important conclusion is that nitrocellulose embedding, although it is an extremely labour-intensive process, yields sections that show only minor deformation even after a period of three decades. As a consequence, rigid registration of this kind of data will generally yield high quality reconstructions.

The entire collection, including the raw digitised sections and our reconstructions, will be made available to the public via the Internet. At the moment, we are developing a web-site and client-server applications for the querying and visualisation of the orbita collection.

References

1. Ham, A., Leeson, T.: *Histology*. 5th edn. Pitman medical publishing (1965)
2. Huijsmans, D., Lamers, W., Los, J., Strackee, J.: Toward computerized morphometric facilities: a review of 58 software packages for computer-aided three-dimensional reconstruction, quantification, and picture generation from parallel serial sections. *The Anatomical Record* **216** (1986) 449–470
3. van den Elsen, P., Pol, E.J., Viergever, M.: Medical image matching - a review with classification. *IEEE Engineering in Medicine and Biology* (1993) 26–39
4. Maintz, J., Viergever, M.: A survey of medical image registration. *Medical Image Analysis* **2**(1) (1998) 1–36
5. Laan, A., Lamers, W., Huijsmans, D., te Kortschot, A., Smith, J., Strackee, J., Los, J.: Deformation-corrected computer-aided three-dimensional reconstruction of immunohistochemically stained organs: application to the rat heart during early organogenesis. *The Anatomical Record* **224** (1989) 443–457

6. Verbeek, F.: Three-dimensional reconstruction of biological objects from serial sections including deformation correction. PhD thesis, Delft Technical University (1995)
7. Kim, B., Boes, J., Frey, K., Meyer, C.: Mutual information for automated unwarping of rat brain autoradiographs. *Neuroimage* **5** (1997) 31–40
8. Rydmark, M., Jansson, T., Berthold, C.H., Gustavsson, T.: Computer-assisted realignment of light micrograph images from consecutive section series of cat cerebral cortex. *Journal of Microscopy* **165** (1992) 29–47
9. Ourselin, S., Roche, A., Subsol, G., Pennec, X., Ayache, N.: Reconstructing a 3d structure from serial histological sections. *Image and vision computing* **19** (2000) 25–31
10. Pitiot, A., Malandain, G., Bardinet, E., Thompson, P.: Piecewise affine registration of biological images. In Gee, J., Maintz, J.A., Vannier, M.W., eds.: *Second International Workshop on Biomedical Image Registration (WBIR2003)*. Volume 2717 of *Lecture Notes in Computer Science*, Springer-Verlag (2003) 91–101
11. Guest, E., Baldock, R.: Automatic reconstruction of serial sections using the finite element method. *Bioimaging* **3** (1995)
12. Krinidis, S., Nikou, C., Pitas, I.: A global energy function for the alignment of serially acquired slices. *IEEE Transactions on Information Technology in Biomedicine* **7**(2) (2003) 108–113
13. Galigher, A., Korloff, E.: *Essentials of practical microtechnique*. Lea and Febiger (1964)
14. Poynton, C.: *Color faq – frequently asked questions about color* (2002)
15. Young, I., Gerbrands, J., van Vliet, L.: *Fundamentals of Image Processing*. Delft University of Technology (1998)
16. Mardia, K.: *Statistics of directional data. Probability and Mathematical Statistics*. Academic Press (1972)
17. Koornneef, L.: *Spatial aspects of orbital musculo-fibrous tissue in man*. PhD thesis, Amsterdam University (1976)
18. de Haan, A.: *De prenatale ontwikkeling van de humane orbita*. PhD thesis (1983-1986) Unpublished.
19. Roche, A., Malandain, G., Ayache, N.: Unifying maximum likelihood approaches in medical image registration. *International Journal of Imaging Systems and Technology: Special Issue on 3D Imaging* **11** (2000) 71–80
20. Ibáñez, L., Schroeder, W., Ng, L., Cates, J.: *The ITK software guide*. 2nd edn. (2005)
21. Press, W., Teukolsky, S., Vetterling, W., Flannery, B.: *Numerical recipes in C, the art of scientific computing*. 2nd edn. Cambridge University Press (1994)