Linköping Studies in Science and Technology
Dissertations, No. 1862

**Structural properties of the mastoid
using image analysis and visualization**

Olivier Cros

Department of Biomedical Engineering
Linköping University, Sweden

Linköping, June 2017

Volume rendering of a 3D shape analysis where medial balls are fitting the mastoid air cell system of a human temporal bone specimen, overlaid over a manually cropped version of the original data. The different colours represent the different size classes: from very small in dark blue to vary large in purple. The bone is rendered with a modified transfer function so as to reveal the micro-channels.

AALBORG UNIVERSITY HOSPITAL

Copyright © 2017 Olivier Cros, unless otherwise noted.

The content from the same author from the licentiate thesis No. 1730 published in 2015 was reused in this doctoral thesis. Printed in Sweden by LiU-Tryck, Linköping 2017

ISBN 978-91-7685-505-8

ISSN 0345-7524

Abstract

The mastoid, located in the temporal bone, houses an air cell system whose cells have a variation in size that can go far below current conventional clinical CT scanner resolution. Therefore, the mastoid air cell system is only partially represented in a clinical CT scan. Where the conventional clinical CT scanner lacks level of minute details, micro-CT scanning provides an overwhelming amount of fine details. The temporal bone being one of the most complex in the human body, visualization of micro-CT scanning of this bone awakens the curiosity of the experimenter, especially with the correct visualization settings.

This thesis first presents a statistical analysis determining the surface area to volume ratio of the mastoid air cell system of human temporal bone, from micro-CT scanning using methods previously applied for conventional clinical CT scans. The study compared current results with previous studies, with successive downsampling the data down to a resolution found in conventional clinical CT scans. The results from the statistical analysis showed that all the small mastoid air cells, that cannot be detected in conventional clinical CT scans, do heavily contribute to the estimation of the surface area, and in consequence to the estimation of the surface area to volume ratio by a factor of about 2.6. Such a result further strengthens the idea of the mastoid to play an active role in pressure regulation and gas exchange.

Discovery of micro-channels through specific use of a non-traditional transfer function was then reported, where a qualitative and a quantitative pre-analysis were performed and reported. To gain more knowledge about these micro-channels, a local structure tensor analysis was applied where structures are described in terms of planar, tubular, or isotropic structures. The results from this structural tensor analysis suggest these micro-channels to potentially be part of a more complex framework, which hypothetically would provide a separate blood supply for the mucosa lining the mastoid air cell system.

The knowledge gained from analysing the micro-channels as locally providing blood to the mucosa, led to the consideration of how inflammation of the mucosa could impact the pneumatization of the mastoid air cell system. Though very primitive, a 3D shape analysis of the mastoid air cell system was carried out. The mastoid air cell system was first represented in a compact form through a medial axis, from which medial balls could be used. The medial balls, representative of how large the mastoid air cells can be locally, were used in two complementary clustering methods, one based

on the size diameter of the medial balls and one based on their location within the mastoid air cell system. From both quantitative and qualitative statistics, it was possible to map the clusters based on pre-defined regions already described in the literature, which opened the door for new hypotheses concerning the effect of mucosal inflammation on the mastoid pneumatization.

Last but not least, discovery of other structures, previously unreported in the literature, were also visually observed and briefly discussed in this thesis. Further analysis of these unknown structures is needed.

Acknowledgements

Many people have contributed to this thesis, directly or indirectly. First, I am grateful for my employer, the department of Otolaryngology, Head and Neck Surgery, at Aalborg Hospital South in Aalborg, Denmark. Without the financial support and great patience, I would not be here today.

I would like to particularly thank MD & PhD Michael Gaihede, my main clinical supervisor, for your constant support and providing this amazing knowledge you have about the ear both from an anatomical but also from a physiological point of view. Simona, you are not forgotten either, I hope you will keep up the hard work.

I would also like to thank Professor Hans Knutsson, my main technical supervisor, for being the source of many ideas and inspiration. Thank you also Dr. Mats Andersson for endless discussions about other things than work, and motivating me during tough periods from personal matters but also when in doubt with my academic career. Also big thanks to Dr. Anders Eklund, my co-supervisor, for your already endless support and our daily talks.

Thank you Professor Magnus Borga for your early supervision. Also, thanks to my colleagues and the staff at the department of biomedical engineering for always being kind and helpful, especially Göran Salerud. Thank you for having me so long.

I would also like to acknowledge the people at the centre for X-ray Tomography, Department of Physics and Astronomy, University of Ghent, Belgium; especially Elin Pawels and Manuel Dierick and their colleagues. Without you, this PhD would definitely have turned in a different direction.

I also want to greet my friends, especially Anders Eklund, Peter Remjstad, Maria Ewerlöf, Filipe Marreiros, Patrick Bennysson, Rafael Sanchez without whom I would have felt quite alone on the daily basis. You have been a great support directly and indirectly.

All my gratitudes to my family, especially my parents Jacques and Nicole Cros, for always being supportive and engaged in my research. Without you I would not be where I am. Mum and Dad, I hope you are in peace where you are now, and I miss you terribly. You have no idea how much I learned from you two and I still carry your principles in me. Thanks to my

both sisters Nathalie Pesson and Véronique Poupon and their respective families.

Mum, thanks to you, I met my beautiful fiancé and future wife Élodie Debiez and despite the fact that you and dad will never meet her, I am sure you would have adopted her right away. Thank you mum for this last but very important gift in life.

Élodie Debiez, thank you for your moral support and your positivism when I most needed it. Thank you Marie Debiez for integrating me in your sweet family. I know that it has been a very tough period for you and me Élodie, but because we have been there for each others, we got stronger than ever. And looking at the future, I also know we will have beautiful moments. We both lost important people in our life, but we also gained someone in life and maybe more. I have no other words than Je t'aime!

Olivier Cros.

Linköping, June 2017.

Table of Contents

1	Introduction	1
1.1	Foreword	1
1.2	Thesis outline	2
1.3	Publications	4
1.4	Related Publications	5
1.5	Abbreviations	6
1.6	Acronyms	6
2	General Anatomy & Physiology	7
2.1	Introduction	7
2.2	Anatomy of the temporal bone	8
2.2.1	Outer ear	9
2.2.2	Middle ear	10
2.2.3	Inner ear	11
2.2.4	Eustachian tube	13
3	The Mastoid	17
3.1	Introduction	17
3.2	Development process	17
3.3	After birth development	18
3.4	Origin of the mastoid air cells in the newborns	18
3.5	The adult mastoid	21
3.6	Level of pneumatization	23
3.7	Pressure regulation and gas exchange	24
3.8	Inflammation of the mucosa lining the air cells	27
3.9	Otitis media	30
4	Imaging Modalities	31
4.1	Introduction	31
4.2	Clinical X-ray CT scanner	33
4.3	Cone-Beam CT scanner	35
4.4	Histological sections	36
4.5	Micro-CT scanner	39
5	Observations from the data and aims of the thesis	43
5.1	Introduction	43
5.2	Midline of the mastoid air cell system	43

5.3	What is a mastoid air cell in terms of shape?	44
5.4	What is the size of an air cell?	47
5.5	Air-mucosa versus mucosa-bone surface area	47
5.6	Discovery of micro-channels	49
5.7	Aims of the thesis	50
6	Image Processing	53
6.1	Introduction	53
6.1.1	Segmentation by thresholding	53
6.1.2	Morphology on binary images	55
6.1.3	Masking original data over a binary segmentation	59
6.1.4	Measuring surface area and volume	60
6.2	More advanced image processing	62
6.2.1	Filter design	64
6.2.2	The quadrature filter	71
6.2.3	Tensor analysis	75
6.2.4	Extraction of planar, tubular, and isotropic structures	76
6.3	Enhancement through adaptive filtering	79
6.4	Volume rendering	87
6.4.1	Multiple volume renderings	90
7	3D Shape Analysis	97
7.1	Introduction	97
7.2	General definition of a shape	98
7.3	Euclidean distance transform	103
7.4	Skeletonization	105
7.5	Medial surface	109
7.6	Medial axis	112
7.7	Medial balls	113
7.8	K-means clustering	114
7.9	Convex hull	115
7.10	Clustering the air cells	116
7.10.1	Clustering in relation to size	116
7.10.2	Clustering in relation to 3D location	120
8	Summary of Papers	127
8.1	Introduction	127
8.2	Paper I	128
8.3	Paper II	128
8.4	Paper III	129
8.5	Paper IV	129
8.6	Paper V	130
8.7	Paper VI	130

9	Discussion & Conclusion	131
9.1	Fulfillment of the aims or not?	131
9.1.1	Geometrical aspects of mastoid air cell system . . .	131
9.1.2	Interpretation and geometrical description of the dis- covered micro-channels	136
9.2	Future methodological aspects to consider	139
9.2.1	Clinical aspects	139
9.2.2	Technical aspects	143
9.3	Strands and membranes within the MACS	148
9.4	Pores in the septae	150
9.5	Bone spicules	152
9.6	Conclusion	154

Introduction

1.1 Foreword

Similarly to the lungs, the rate of gas exchange carried out by the mastoid process in the temporal bone is determined by the mucosal surface area [64]. This is particularly the case for the mastoid bone with its complex air cell system, with cells of varying sizes and shapes. An impaired gas exchange, besides a malfunctioning Eustachian tube, results in a negative middle ear pressure; i.e. both in the tympanum and in the mastoid air cell system. Negative middle ear pressure leads to different middle ear diseases, such as the otitis media with effusion, cholesteatoma invading all the airspaces in the middle ear, or retraction pockets in the tympanum [62]. As stated in [19], understanding the mechanism of the middle ear pressure regulation is important for both physiologists and practising clinicians; especially for the practising clinicians in their decisions on how to treat their patients. The surface area of the mucosa in the middle ear, especially the one covering the mastoid air cell system, is therefore a valuable parameter for physiological studies of gas exchanged between the air cells and the capillaries present in the mucosa lining the air cells [18, 19, 62].

Quantitative measurement of the entire mastoid air cell system aeration is, however, only reported in few studies [17, 64, 78]. A plausible explanation of this sparse literature emanates from the fact that there is no technique available to allow a direct measurement of the mucosa surface area for the entire mastoid air cell system in vivo. An alternative is to scan the mastoid bone through X-ray computed tomography (CT), and consider the walls of the mastoid air cell system as a surrogate to the very thin mucosa invisible on clinical CT scans. The volume of gas occupied within the mastoid air cells, also important to estimate when investigating how well a mastoid is pneumatized, is though easier to estimate using X-ray CT and has been reported in several studies [84, 17, 37, 48, 54, 64, 77, 88].

Recent advances in micro-CT scanning technology allows scanning such complex structures using very high resolution, in turn producing more accurate statistics. The aim of this work is to investigate the use of micro-CT scanning of human temporal bone specimens, to estimate the surface area to volume ratio using classical image processing methods, to compare with results from previous studies where conventional clinical CT scans were used, and finally to assess whether the obtained estimates help to further understand the anatomy and physiology of the mastoid air cell system.

Micro-channels were discovered while visualizing the temporal bone specimens using different visualization settings. A structural analysis of the micro-channels within the bone was therefore assessed. The results from the structural analysis further suggested to enhance the original data by reducing the noise level to a minimum, while slightly enhancing their representation.

Inspired from the idea of the micro-channels forming a supplementary blood supply to the mucosa lining the mastoid air cell system, the last part of the study aimed at investigating how the diameter of the air cells, understood as the cell size, influence the level of pneumatization of the temporal bone, at the level of the mastoid air cell system during inflammation. To achieve this aim, a compact representation of the mastoid air cell system was obtained through the use of a structure tensor analysis on a distance transform computed from the enhanced data. A shape analysis was possible through the use of a medial unit composed of a medial surface, a medial axis and medial balls. Clustering of the medial balls in terms of size and three-dimensional location could help investigate the size variation, but also where each size class is located in the mastoid. Although the medial balls are not a direct representation of the mastoid air cells, their use in a shape analysis opens a new door into the analysis of the mastoid in terms of pneumatization, which is beyond the more conventional surface area and volume estimations.

Overall, this thesis is an attempt to open new doors when analyzing the mastoid air cell system using image analysis while using micro-CT scans.

1.2 Thesis outline

The thesis is divided into nine parts. The thesis is formed in a way such that non-medical readers can obtain a brief and basic introduction to the anatomy and physiology of the human temporal bone in Chapter 2, and obtain more information about the mastoid process alone in Chapter 3. Chapter 4 gives a brief description of the image modalities often used when investigating the mastoid, from conventional clinical CT scanner up to histological sections, together with the image modality that was used to pro-

duce the high resolution scans, i.e. X-ray micro-CT. This chapter can be of interest for both clinical and technical readers who are not so familiar with the different modalities. Chapter 5 states the aims for this overall work based on observations made from the micro-CT scans supported by histological sections at some occasions. For non-technical readers, chapter 6 introduces some necessary concepts in image processing. The second half of chapter 6 describes a method used forms the basis for Papers III, V, VI, and VII. Chapter 7 is devoted to a 3D shape analysis and stands on its own due to the use of tools that belong more to discrete geometry and pattern recognition than image processing. Chapter 8 summarises the contribution of each paper. Chapter 9 provides a discussion about the presented work and ideas for future work, followed by a section illustrating the presence of mucosal strands in the mastoid air cell system which has not been found in previous literature. Two unpublished anatomical findings more related to the bone than soft tissues, are also reported.

***N.B.:** It should be noted that besides **Figs. 2.4, 2.5, 2.6, 2.7, 2.8, 2.9, 3.1, 4.5, 4.6, 4.7, 5.6, 9.3**, all remaining illustrations presented in this thesis were produced by the sole author Olivier Cros. Permission to use these illustrated should be asked beforehand. Moreover, some illustrations are from the sole author but were published in some of the papers. Fig. 5.1 is published in Paper V and therefore permission to use this illustration has to be granted by IEEE. Fig. 6.29 is published in Paper VI and therefore permission to use this illustration has to be granted by IEEE. Figs. 3.5, 7.10, 7.15, 7.19, and 7.22 are submitted in Paper VII and therefore permission to use this illustration has to be granted by Elsevier.*

1.3 Publications

- I **Olivier Cros**, Hans Knutsson, Mats Andersson, Elin Pawels, Magnus Borga, Michael Gaihede. *Determination of the mastoid surface area and volume based on micro-CT scanning of human temporal bones. Geometrical parameters depend on scanning resolutions.* **Accepted** and published in the medical journal Hearing Research Special Issue MEMRO, Volume 340, pages 127-134, 2016.
- II **Olivier Cros**, Magnus Borga, Elin Pawels, Joris JJ Dirckx, Michael Gaihede. *Micro-channels in the mastoid anatomy. Indications of a separate blood supply of the air cell system mucosa by micro-CT scanning.* **Accepted** and published in the medical journal Hearing Research Special Issue MEMRO, Volume 301, pages 60-65, 2013.
- III **Olivier Cros**, Michael Gaihede, Mats Andersson, Hans Knutsson. *Structural analysis of micro-channels in human temporal bone.* **Accepted** for IEEE International Symposium on Biomedical Imaging (ISBI), New York, United States of America, pages 9-12, April 2015.
- IV **Olivier Cros**, Anders Eklund, Michael Gaihede, Hans Knutsson *Enhancement of micro-channels within the human mastoid bone based on local structure tensor analysis.* **Accepted** for IEEE sixth International Conference on Image Processing Theory, Tools and Applications (IPTA'16), IEEE, Oulu, Finland, December 2016.
- V **Olivier Cros**, Michael Gaihede, Anders Eklund, Hans Knutsson *Surface and curve skeleton from a structure tensor analysis applied on mastoid air cells in human temporal bones.* **Accepted** for IEEE International Symposium on Biomedical Imaging (ISBI), Melbourne, Australia, April 2017.
- VI **Olivier Cros**, Anders Eklund, Michael Gaihede, Hans Knutsson *Local size descriptor of mastoid air cells in human temporal bone based on a structure tensor analysis from an Euclidean distance transform.* **Submitted** to the journal Computerized Medical Imaging and Graphics, March 2017.

1.4 Related Publications

- I Gunnar L G., **Olivier Cros**, Hans Knutsson, Magnus Borga. *Non-ring filters for robust detection of linear structures*. Note: little contribution. IEEE 20th International Conference on Pattern Recognition (ICPR), pp. 233-236, 2009.
- II **Olivier Cros**, Hans Knutsson, Mats Andersson, Elin Pawels, Magnus Borga, Michael Gaihede. *Mastoid structural properties determined by analysis of high-resolution CT scanning*. Note: medical abstract. Hearing Research, 263:1-2, pages 242-243, 2010.

1.5 Abbreviations

This list provides abbreviations used in this thesis, along with their meanings.

CT	Computed Tomography
MAC	Mastoid Air Cell
MACS	Mastoid Air Cell System
ME	Middle ear
TYMP	Tympanum
TM	Tympanic membrane
OM	Otitis media
EC	Ear canal
EDT	Euclidean distance transform
ET	Eustachian tube
SNR	Signal to Noise Ratio

1.6 Acronyms

This list provides the different acronyms related to anatomical locations.

- **Lateral** towards the outside,
- **Medial** towards the midline,
- **Anterior** towards the front,
- **Posterior** towards the back,
- **Superior** towards the top,
- **Inferior** towards the bottom,

- **Endo-** inner part of a structure,
- **Meso-** middle part of a structure,
- **Ecto-** outer part of a structure,
- **Exo-** external to a structure,
- **Epi-** outside a structure,
- **Hypo-** under the structure,
- **Hyper-** over the structure,
- **Inter-** between structures,
- **Intra-** within a structure,
- **Peri-** surrounding a structure,
- **Retro-** behind a structure.

General Anatomy & Physiology

2.1 Introduction

Before introducing the temporal bone as a whole entity, taxonomy of anatomical terms and spatial positions is briefly introduced, with the six most relevant positions illustrated in Fig. 2.1, i.e. inferior, superior, posterior, anterior, medial and lateral. A list of acronyms is provided in Chapter 1.

Figure 2.1: *Anatomical directions.*

These terms are used alternatively, globally or locally. As a simple example, the temporal bone is located laterally in relation to the human skull, while the cochlea is located medially in relation to the temporal bone. Combination of locations are also possible as, for instance, the superior retrosgmoid air

cells are located at the superior part and at the back of the mastoid process. A structure located at the periphery of another structure is for example the lateral perisinus cells, referring to air cells at the periphery of a sinus located in the back of the mastoid towards the lateral side of the bone. As a side note, mastoid air cells located at the centre of the mastoid can also be named medial or central air cells, while air cells located towards the outside of the mastoid bone will be called lateral or peripheral air cells. This notion should be kept in mind while reading the thesis. To further help the reader, taxonomy of the different terms used when describing a structure orientation in medical images is resumed in Section 1.6 and represented graphically in Fig. 2.1.

2.2 Anatomy of the temporal bone

The temporal bone houses the organ of hearing. The temporal bone consists of five parts: the *temporal squamosa*, the *petrous portion*, the *mastoid process*, the *tympanic bone*, and the *styloid process*, see Fig. 2.2.

Figure 2.2: *The temporal bone viewed alone from a lateral side (left) and a medial side (right), respectively viewed from the outside and the inside.*

The temporal squamosa (from Latin squama, "scale"-shape structure) is the biggest part of the temporal bone, lying superior to all parts. The petrous portion of the temporal bone, located inside the skull, is recognizable with its pyramidal shapes housing the middle ear and the inner ear. The word petrous originates from the Latin word petrosus, and relates to the very hard portion of the bone housing the internal auditory organs. The mastoid process, from the new Latin "mastoides" resembling a nipple or breast, is located posteriorly and inferiorly to the squamous part.

The tympanic bone is a small portion situated inferior to the temporal squamosa, anterior to the mastoid process, and superior to the styloid process. The styloid process runs inferior to the tympanic bone, and is shaped

like a thorn pointing downwards. It is used as an anchor point for several muscles. The described structures are illustrated in Fig. 2.2. A coronal section of a right temporal bone reveals the complexity inside the temporal bone, see Fig. 2.3.

Figure 2.3: *Cut through a temporal bone to reveal all its internal structures.*
Source: *Atlas of Skull Base Surgery & Neurotology*, Thieme, 2009.
Image copyrighted by RK Jackler. Permission granted for non-profit educational use.

From a physiological point of view, the temporal bone can be decomposed into three main parts: the **external ear**, the **middle ear**, and the **inner ear**. These parts are described in the following sections.

2.2.1 Outer ear

The outer ear is the external portion of the ear, which consists of the pinna, the external auditory meatus (also known as the ear canal), and the tympanic membrane, see Fig. 2.4.

The **pinna**, also known as the auricle, is the visible portion that is generally referred to as "the ear". Its function is to localize sound sources and direct sound into the ear. The folds of the pinna allow some specific frequencies to be amplified, while other frequencies can be damped.

The **external auditory meatus (EAM)**, also named the external auditory canal or simply the ear canal, extends from the pinna to the tympanic membrane and has a length of about 26 millimetres (mm) and a diameter of about 7 mm. The size and shape of the ear canal vary among individuals and between the left and right ears.

Figure 2.4: *The outer ear.* Legend *EAM: external auditory meatus.* Source: *Atlas of Skull Base Surgery & Neurotology*, Thieme, 2009. Image copyrighted by RK Jackler. Permission granted for non-profit educational use.

The **tympanic membrane**, also known as the eardrum, is a cone-shaped structure separating the outer ear from the middle ear and protects the middle and inner ear from foreign objects. The tympanic membrane resonates in response to sound pressure waves. The displacement during vibration is extremely small, about one-billionth of a centimetre.

2.2.2 Middle ear

The middle ear cavity is a combination of two cavities, the *tympanum* and the *mastoid air cell system*. Since the mastoid will be described further, only the tympanum is briefly introduced. The tympanum is the narrow air-filled space of the middle ear, where the ossicles are located, see Fig. 2.5. The tympanum resembles an oblique rectangular room with a floor, a ceiling, and four walls. Sound waves traveling through the ear canal will hit the tympanic membrane. This wave information travels across the air-filled tympanum via a series of delicate bones called the ossicles, see Fig. 2.6. The ossicles are composed of the malleus (also known as hammer), incus (also known as anvil) and stapes (also known as stirrup), see Fig. 2.6. They form an ossicular chain.

While the handle of the malleus, also known as the manubrium, articulates with the tympanic membrane, the footplate of the stapes articulates with the oval window, a membrane-covered opening which leads to the vestibule

Figure 2.5: *The tympanum, being part of the middle ear. Source: Atlas of Skull Base Surgery & Neurotology, Thieme, 2009. Image copyrighted by RK Jackler. Permission granted for non-profit educational use.*

of the inner ear, located at the opposite side of the tympanum. Tendons attaching the head of the malleus and the incus are attached to the tegmen tympani, a bony wall separating the cranial cavity from the superior part of the tympanum, the epitympanic recess or attic, see the vertical tendons in Fig. 2.6. Behind the head of the malleus and towards its upper part, the posterior wall of the tympanum is mostly occupied by the aditus ad antrum, a path towards the mastoid air cell system via the antrum.

2.2.3 Inner ear

The inner ear, shown in Fig. 2.7, can be divided into two labyrinths: a bony labyrinth and a membranous labyrinth.

The osseous labyrinth consists of the cochlea, the vestibule, and the semi-circular canals, a series of bony cavities within the petrous temporal bone. These bony cavities are lined with periosteum and contain perilymph; extracellular fluid at the periphery of the structure. The oval window, articulated by the footplate of the stapes, is an opening in the lateral wall of the vestibule of the osseous labyrinth.

The membranous labyrinth is composed of communicating membranous sacs and ducts housed within the osseous labyrinth. It is cushioned by the surrounding perilymph and contains the endolymph within its confines; an extracellular fluid inside the structure. The membranous labyrinth also

Figure 2.6: *The ossicular chain with the malleus in direct contact with the tympanic membrane, and the stapes viewed laterally attached to the oval window separating the vestibule from the tympanum. The incus transmits the movements induced by the malleus to the stapes. Also, notice the tendons attaching the ossicles to the tympanum. Source: Gray plate 919 made public available.*

houses a cochlear, a vestibular, and semicircular canals components.

Cochlea

The cochlea is the auditory portion of the inner ear. The spiral-shaped cavity within the cochlea has three fluid-filled sections along its main axis: the scala vestibuli, the scala tympani, and the cochlear duct. The organ of Corti, located in the cochlear duct on the basilar membrane, transforms mechanical waves into electric signals sent to the neurons in the brain.

Vestibule

The vestibule is the central part of the osseous labyrinth. The vestibular apparatus is composed of the utricle and the saccule. They respectively sense linear acceleration in the horizontal and vertical planes. Within these organs are hair cells. The cilia of these cells are intricately associated with a membranous substance containing calcium carbonate granules, or "otoliths", also commonly known as stones in the cochlea. Movement of the head induces a shearing of the hair cells by the mobile otoliths. This directional change is sensed by the brain via the superior division of the

Figure 2.7: *The inner ear.* Source: *Atlas of Skull Base Surgery & Neurotology*, Thieme, 2009. Image copyrighted by RK Jackler. Permission granted for non-profit educational use.

vestibular nerve at the utricle, and via the inferior division of the vestibular nerve to the saccule. Together, the otolithic organ organs of both ears are of prime importance for directional sensation.

Semi-circular canals

The third main apparatus in the inner ear is the set of three semi-circular canals. Each stands at right angles in relation to each others. The superior, posterior, and lateral semicircular canals are located posterior and superior to the vestibule.

2.2.4 Eustachian tube

The Eustachian tube originates in the posterior part of the nose, runs slightly uphill, and enters the tympanum inferiorly, see Fig. 2.8 (left). The cartilage provides a supporting structure for two thirds of the Eustachian tube, while the part closest to the tympanum is made of bone. The tissue lining the Eustachian tube is similar to that inside the nasal cavity, and may respond in the same way (swelling) when presented with similar stimuli.

Function of the Eustachian tube

The primary function of the Eustachian tube is to ventilate the middle ear space, ensuring that its pressure remains at near normal environmental

Figure 2.8: *Course of the Eustachian tube from the middle ear (left) and its entrance from the nasal cavity (right), see the red vertical arrow.*
 Source: left. Hill, M.A. (2015) Embryology Gray0915.jpg, right: [1].

air pressure; known as middle ear pressure. The secondary function of the Eustachian tube is to drain any accumulated secretions, infections, or debris from the tympanum.

Several small muscles located in the back of the throat and palate control the opening and closing of the tube. Swallowing and yawning cause contractions of the muscles located in the back of the throat, and facilitate the regulation of the Eustachian tube function. If it was not for the Eustachian tube, the middle ear cavity would be an isolated air pocket inside the head, which would be vulnerable to every change in air pressure, under- and over-pressure, leading to a pathological middle ear.

Normally, the Eustachian tube is closed, which helps to prevent the inadvertent contamination of the middle ear space by the normal secretions found in the back of the nose. A dysfunctional Eustachian tube can lead to chronic ear infections. A much more common problem is a failure of the Eustachian tube to effectively regulate air pressure. Partial or complete blockage of the Eustachian tube can cause sensations of popping, clicking, and ear fullness, and occasionally moderate to severe ear pain. Such intense pain is most frequently experienced when sudden air pressure changes arise while traveling by airplane, particularly during take-off or landing.

As the Eustachian tube function worsens, air pressure in the middle ear falls, the ear feels full and sounds are perceived as muffled. Eventually, a vacuum is created which can then cause fluid to be drawn into the middle ear space (termed serous otitis media). If the fluid becomes infected, a common ear infection (suppurative otitis media) will eventually be developed [77].

Mucosa lining the tympanum and the Eustachian tube

The mucosa is the innermost layer of hollow organs. All bone surfaces of the ME cleft are covered with a mucosa. The structure of the mucosa is subdivided in an epithelial layer (mucosal epithelium) located on the top of a basal membrane (also called basal lamina), and a connective tissue layer called lamina propria, adherent to the the outer part of the underlying bone; the periosteum layer [83]. The lamina propria is formed of loose connective tissue, characterized by a less organized appearance and by being composed of relatively few cells (mezenchymal cells, basal cells, fibroblasts, lymphocytes, macrophages, etc) in a large volume of extracellular matrix. The major matrix components are collagen and elastic fibres. Generally, the loose connective tissue contains many blood vessels, nerve endings, lymphatic vessels and, interestingly, misses a basal membrane making it easier to be crossed by macromolecules and cells [83].

Starting from the nasopharynx (located between the mouth and the nose; see Fig. 2.8 (right image to the left of the black arrow)), the epithelium is respiratory, with one to three layers of columnar or cubical cells, often ciliated or with microvili, secretory and non-secretory, and intercalated with goblet cells [66]. The role of these ciliated cuboidal cells is to carry waste towards the Eustachian tube.

The antero-inferior part of the tympanum (on the frontal part and on the lower part), is structurally similar to the mucosa from nasopharynx and the eustachian tube. In the postero-superior part (the back part of the tympanum on the top part) of the tympanic cavity and the mastoid, the epithelium is mono-layered, flat, alternating with very rare cuboidal cells [32]. The postero-superior and antrum mucosa also seem to be more abundant with superficial blood vessels, compared to the tympanic cavity [5].

The common structures have now been introduced, and the next chapter is solely dedicated to the mastoid bone and the mastoid air cell system, along with some information about its physiological role in pressure regulation.

The Mastoid

3.1 Introduction

In this chapter, more facts about the mastoid are provided. Two sections deal with the developmental aspects from a bone point of view, before and after birth. A third section then provides a personal interpretation on how the air cells invade the mastoid over time. A fourth section then describes a fully functional adult mastoid with the different air cell clusters described as regions. Because not all mastoids have a similar air cell system, different levels of pneumatization are briefly introduced. The following two sections are more related to the physiology of the mastoid air cell system, as a medium for pressure regulation where one section shortly describes how pressure is regulated via the mucosa through gas exchange, and finally what happens in cases of negative pressure and the possible pathology following an middle ear inflammation. It should also be noted that due to the complexity of the mastoid air cell structure, the physiology of the mastoid air cell system is still not fully understood. Therefore the aim of these last three sections is to provide a very basic understanding of the mastoid physiology, especially useful in Chapter 7.

3.2 Development process

Slightly before the prenatal development is over, three principal components represent the temporal bone: the squama, the petromastoid, and the tympanic ring. Because only the mastoid bone is of interest in this work, the remaining parts are further described in [27].

The petromastoid part is developed from four different centers. Their appearance occurs around the fifth or sixth month in the ear capsule still in a cartilaginous state. One ossification forms part of the cochlea, vestibule,

superior semicircular canal, and medial wall of the tympanic cavity, also known as the proötic ossification or anterior ossification. A second ossification shapes the floor of the tympanic cavity and vestibule, surrounds the carotid canal, forms the remaining part of the cochlea, and spreads towards the more interior and inferior part of the internal auditory meatus. This ossification is known as opisthotic. A third ossification, named pterotic, casts the roof of the tympanic cavity and antrum. The fourth ossification, called epiotic, appears near the posterior semicircular canal, and is also responsible for the formation of the mastoid process.

The three main components, originally representing the temporal bone, then fuse together in a chronological order. Shortly before birth, the tympanic ring unites with the squamous part of the bone. During the first year at about the same time, the petromastoid and squamous parts attach to each others as well as to the tympanic part of the styloid process.

3.3 After birth development

After birth, a series of evolution processes also occur, and are briefly resumed in the following. The tympanic ring extends outward and backward to form the tympanic part. The mandibular fossa deepens and directs more inferiorly. The zygomatic process projects like a shelf at a right angle to the squama. The mastoid portion, at first flat, grows towards the anterior and towards the inferior to form the mastoid process with the air cell development. The growth of the mastoid process towards the inferior and the anterior pushes the tympanic part forward, resulting in a rotation of the original floor of the ear canal to become the anterior wall, see Fig. 3.1.

3.4 Origin of the mastoid air cells in the newborns

In the newborn, the mastoid bone is practically one single air space, later on named the antrum, surrounded by diploic bone, i.e. mostly containing marrow. As the mastoid process develops, the marrow space hollows out and the air cell system is being formed, while becoming lined with a highly vascular cuboidal epithelium layer emanating from the endoderm. This cuboidal epithelium will later on form the mucosa of the mastoid air cell system. The endoderm is one of the three primary layers in the very early embryo. The air cells grow out primarily from the attic of the tympanum to the aditus ad antrum, creating a complex structure of cavities separated by bony walls communicating with each other, see Fig. 3.2(A and B) for two different stages of pneumatization.

Figure 3.1: *Difference in shape from a newborn skull when compared to an adult skull highlighted at the level of the mastoid process. Source: Sobotta plates 40 and 104 publicly available.*

Similarities with the alveoli of the lungs, where extensive gas exchange takes place, have been found due to a close contact between the mucosa and the blood vessels, see Fig. 3.2(C). This has lead to the mastoid air cell system being known as a mini-lung.

Figure 3.2: *Two different stages (A & B) where the epithelial lining covers the newly formed air cells during the pneumatization process. C. gives a simplified illustration of a set of alveoli formed by an epithelial lining and surrounded by interstitial space with blood vessels at the proximity for gas exchange. Abbreviation in C.: (B) blood vessel, (Alv.) alveolus.*

A broader representation of the growth of the mastoid air cell system is illustrated in Fig. 3.3. As can be observed, the mastoid air cell system originates from the tympanum, and spreads over time.

A very interesting study has been carried out to simulate a morphogenetic

Figure 3.3: *Transverse plane of the hearing system. A. First stage: single air cell later called antrum, B. Beginning of the cellularization, C. Enlargement of the mastoid air cell system, D. Completion of the pneumatization with small air cells at the periphery of the main MACS. Abbreviations: (A) antrum, (CB) cerebellum, (DM) diploic mastoid, (EAC) external auditory canal, (M) mastoid, (P) pinna, (SCC) semi-circular canals, (SS) sigmoid sinus, (T) tympanum, (TMJ) temporomandibular joint.*

model of the cranial pneumatization, based on an invasive tissue hypothesis. Zollikofer et al. [94] simulated the invasion of airspaces in cranial bone along with the mucosa lining the airspaces. Not only did they simulate the growth of the mastoid air cell system by a mathematical model, but they also indirectly showed how the airspaces can be constrained by surrounding structures. This would explain the possible variation of the shape of the mastoid air cells, in relation to their location within the bone.

3.5 The adult mastoid

After approximately 17 years, the mastoid has reached an adult size and stops growing. The mastoid air cell system, according to A. Allam [3], with a further extension by M. Tos work [85], can be categorized into various sub-regions:

- **Antral** - anterior to the antrum at the proximity and towards the ear canal,
- **Periantral** - at the periphery of the antrum,
- **Tegmental** - located at the level of the tegmen tympani at the superior level of the mastoid bone above the ear canal,
- **Lateral and medial zygomatic** - at the level of the zygomatic bone laterally and medially in relation to the external auditory canal,
- **Sinodural** - at the postero-superior level of the mastoid in the vicinity of the cranial cavity,
- **Central mastoid** - all air cells immediately inferior to the cells in the mastoid antrum,
- **Superior retrosigmoid** - at the supero-posterior part of the mastoid bone near the sigmoid sinus,
- **Medial and lateral Perisinusal** - all air cells at the periphery of the sigmoid sinus towards the medial and lateral parts of the mastoid process,
- **Inferior retrosigmoid** - at the infero-posterior part of the mastoid bone near the sigmoid sinus,
- **Perifacial** - at the mid level of the posterior wall of the external ear canal,
- **Medial and lateral apical** - at the tip level of the mastoid process respectively towards the anterior and posterior walls of the external ear canal and towards the pinna.

Altogether, these air cells form the so-called mastoid air cell system (MACS), as represented in Fig. 3.4. The mastoid portion is represented by eleven groups: (1) the *antral* cells, (2) the *periantral* cells, (3) the *tegmental* cells, (4) the zygomatic cells, (5) the *sinodural* cells, (6) the *central mastoid*, (7) superior retrosigmoid cells, (8) medial perisinusal, (9) inferior retrosigmoid cells, (10) perifacial cells, (11) apical cells.

The representation of the air cell system, as illustrated in Fig. 3.4, is an updated version of the original drawing from M. Tos ([85]), based on observations from micro-CT scans of temporal bone specimens. It should be noted that this mastoid air cell system representation varies considerably from one bone to another, though some common features are found. To further illustrate this shape variation, Fig. 3.5 shows an interesting case where the apical part presents very large air cells while the remaining part

Figure 3.4: *Exposed mastoid reflecting the mastoid air cells present within the bone. This illustration is inspired from [85] but modified based on observation from scanning of several temporal bone. The numbers represent the different types of cells.*

of the bone is similar to other mastoid bones.

Figure 3.5: *Volume rendering of a mastoid process used during this overall study with the presence of large apical cells [14].*

3.6 Level of pneumatization

The mastoid air cell system is often categorized based on 3 different types of pneumatization:

- *sclerotic* where the pneumatization is absent,
- *diploic* where the pneumatization is poor due to the presence of bone marrow inside the air cells,
- *pneumatic* where the pneumatization is complete.

A recent attempt to measure the level of pneumatization in a more systematic manner using clinical CT scans has been proposed by Han et. al. [30], as illustrated in Fig. 3.6. To show the variation in pneumatization among subjects, four different cases, each representing a certain level of pneumatization, are presented.

Figure 3.6: Rough illustrations of different levels of pneumatization for four different patients inspired from [30] with the addition of the designated levels L1, L2, L3, and L4 together with the dotted line between the malleoincudal complex and the superior line between L1 and L2, as well as the three coloured dots providing the landmarks at the level of the sigmoid sinus. Note that the multiple air cells are rendered as a single space for the sake of simplicity in the drawings.

The four levels of pneumatization are designated as L1, L2, L3, and L4 with one of the levels given in red to highlight which degree of pneumatization

it belongs to. The procedure uses the so-called malleoincudal complex, i.e. the malleus and the incus (see Fig. 2.6), and the groove receiving the sigmoid sinus. Three parallel lines, normal to the line formed by the malleoincudal complex are drawn on the slice image where the malleoincudal complex appears as an ice-cream cone. The upper line corresponds to the most anterior part of the sigmoid sinus, see the green dot in Fig. 3.6. The second line passes through the most lateral aspect of the sigmoid sinus groove, see the blue dot in Fig. 3.6. The third line passes through the most anterior point of the sigmoid sinus, represented by the purple dot in Fig. 3.6.

3.7 Pressure regulation and gas exchange

In organisms, gas exchange is carried on by diffusion governed by Fick's law. The principal factors behind Fick's law are the surface area of the membrane where diffusion occurs, the thickness of the membrane, the concentration gradient, and the speed at which the molecules diffuse. According to [77], the tympanum is essentially a single large air-cell. However, whereas the mucosa covering the tympanum respects many of the conditions for an efficient gas exchange, the surface area to volume ratio is limited by the tympanum principally being a single large air cell, see Fig. 3.7.

Figure 3.7: *Enhanced representations of the tympanum at different depths, to expose the regularity of its walls viewed from top to bottom. The data was maximally cropped around the tympanum, using clipping planes, to mostly concentrate on the tympanum itself, allowing a better visualization.*

Opposite to the tympanum, the mastoid air cell system presents a very complex structure with a very intricate surface, especially at the level of

the antrum from where most air cells emanate, see Fig. 3.8. The walls of each individual air cell often appear smooth. However, when observing the locations where the air cells split from each other, these conducts display spicules and columnar structures which greatly increase the surface area in relation to the overall volume of the mastoid air cell system.

Figure 3.8: *Representations of the mastoid air cell system at different depths, to expose its shape complexity viewed from top to bottom.*

The gases present in both the tympanum and in the mastoid are identical to those found in the blood and in the atmosphere, namely: O_2 , CO_2 , N_2 , Ar (argon), and H_2O in the vapour form. In gas exchange, the different gases are often presented in terms of partial pressure, as $P(O_2)$ standing for partial pressure of oxygen. In the middle ear, the partial pressure of oxygen and carbon dioxide are slightly lower and higher than in the venous blood respectively, meaning that the middle ear is consuming oxygen and producing carbon dioxide at a moderate level. Both oxygen and nitrogen are absorbed in the blood via the mucosa from the different airspaces of the middle ear, while the carbon dioxide and water vapour are diffused from the blood to the airspaces of the middle ear via the mucosa. Therefore, the gas exchange depends on the functional properties of the cells of the mucosa, the specific diffusion rate of each gas, and how the vascular system behaves locally.

The gas exchange regulation allows the tympano-ossicular complex to vi-

brate in an optimal manner, the tympano-ossicular complex being composed of the tympanic membrane together with the ossicular chain composed of the malleus, incus, and staples. To ensure an optimal environment, the air pressure inside the middle ear, i.e. tympanum and mastoid air cell system, has to be close to the atmospheric pressure, which is 760 mmHg. In ambient air, atmospheric pressure corresponds to the sum of the partial pressures of the four gases in the air, that is oxygen (158 mmHg), carbon dioxide (0.3 mmHg), nitrogen (596 mmHg), and water vapour (5.7 mmHg). However, because the middle ear is a closed cavity directly connected to the nasopharynx via the Eustachian tube, the gas entering the tympanum consists more of exhaled gas with a reduced amount of oxygen and more carbon dioxide than found in the ambient air. Because gas diffusion occurs between the arterial and venous blood system and the middle ear, the gas composition also varies. Oxygen in the arterial blood is about 93 mmHg, while carbon dioxide is around 44 mmHg in the venous blood. These differences lead to a gradient from the middle ear to the capillaries of 57 mmHg for O_2 , and a gradient from the capillaries to the middle ear of 39 mmHg for CO_2 . The sum of the partial pressures of oxygen and carbon dioxide in the middle ear is therefore 90 mmHg which is lower than the equivalent in ambient air (150 mmHg). Because of its very slow diffusion towards the capillaries, nitrogen is the gas exerting the higher partial pressure (623 mmHg).

To ensure a gas pressure balance, the two most relevant systems are explained in the following. The first system corresponds to the opening and closing of the Eustachian tube, as explained in the previous chapter. The second system is the vascular system. Variations of the blood flow in the middle ear, impacting the permeability of the vessels, allow adaptations to normal gas pressure fluctuations. Because the surface area of the mucosa lining the mastoid air cells is known to be larger than the surface area of the tympanic mucosa, the mastoid air cell system represents the main vascular gas exchange effector. Ars et al. [6] consider the gas contained in the mastoid air cells to act as a passive pressure buffer.

In fact, the anatomical properties of the pneumatization are important to take into consideration: the greater the volume, the more compliant the system is, the larger the surface area the more gas can be exchanged through the mucosa. As already introduced, the size of the mastoid air cell system varies considerably from person to person. According to Ars et al. [6], the smaller the mastoid air cell system volume, the faster the deviation from normal pressures small air cell systems are therefore more prone to pathology. Under normal conditions, as long as the small mastoid can level out pressure differences arising in the middle ear, the mastoid is considered normal or healthy. However, when the mucosa lining the middle ear becomes inflamed, the small mastoid has a higher probability to fail in

compensating for a too high negative pressure.

3.8 Inflammation of the mucosa lining the air cells

Back in 1928, J.P. Stewart [74] described the effects of inflammation on the mastoid air cell system, in the case of mastoiditis, meaning inflammation of the mastoid air cells in the mastoid as well as in the neighbouring parts. When the Eustachian tube or a group of air cells is blocked, an oedema occurs. The single layer of flattened epithelium becomes swollen by the accumulation of cells in the lamina propria layer. The outer layer of the epithelium locally erupts and allows cells to pass through it. The capillaries in the mucosa are also saturated, which leads to angiogenesis, i.e. building of new capillaries. Leucocytes of polymorphonuclear type are filtrated out from the blood to the surrounding tissues. When the epithelial lining becomes/is erupted, inflamed connective tissue from the lamina propria enters the lumen, i.e. the airspace, through the defects of the epithelial lining. This process is known as exudation, i.e. the escape of fluid, cells, and cellular debris from blood vessels and their deposition in or on the tissues typical in inflammation. Through the gaps formed on the erupted epithelium, the newly formed capillaries together with fibroblasts form tissue bridges.

When the inflamed mucosa from two opposite sides of an air come in contact, connections form between two gaps and form tissue bridges. Epithelial cells form along these bridges, leading to mucosal strands. Mucosal strands are also found from a study performed by P. Cayé-Thomasen et al. [11] and described as adhesions. Such mucosal strands may eventually form sheets obstructing an air cell completely from the rest of the air cell system, or form some sort of ropes. The remaining bridges are known as pseudocysts or pearls. Eventually, when the fibrosis is complete, the mucosal strands shrink in size. The mucosal strands either become flattened or degenerate, and their content fills out the lumen of the air cell over time with fibrous tissue. This presence of stagnating serous fluid in the lumen of the air cells may eventually turn into bone forming tissue, leading to bone and thus enforcing the mastoid to become diploic up to sclerotic. This series of events is commonly referred, depending on the degree of the infection to otitis media. Different types of otitis media are described in the next section.

Depending on the local size and the shape of the air cells, inflammation of the mucosa can have severe impact on the pneumatization of the mastoid. In Fig. 3.9, three possible cases are provided, where the pneumatization is changed during and after the mucosal inflammation; as illustrated by the three rows pointed with black arrows between the rows to show the evolution over time. The three cases represent air cells having different shapes. In the first case, as shown in Fig. 3.9(A), an air cell is represented

as a finger-like elongated structure. A second case where two air cells communicate via a narrow conduit is illustrated in Fig. 3.9(B). Another pair of air cells with a more weakly defined splitting is presented in the third case, see Fig. 3.9(C).

Figure 3.9: *Illustration of the effect of mucosal thickening on three different air cell configurations (A, B, C) when the mucosa becomes inflamed and after recovery.*

In Fig. 3.9(A) on the second row, the swollen mucosa from both sides of the thin-elongated air cell come in contact almost all the way along its principal axis. In Fig. 3.9(B), the mucosa from both sides is only coming in contact at the level of the narrow conduit. The dashed lines in the air cell below the occlusion site refers to the presence of a liquid resulting from the exudation. In Fig. 3.9(C), the mucosa from opposite walls, though swollen, does not come in contact with each others, still allowing gas exchange to occur locally. On the third row labelled as *Recovery*, the mucosa from both sides have in the first case, as shown in Fig. 3.9(A), merged into one through the fibrosis and is no longer active in gas exchange. Over time,

the soft tissue may turn into bone, thus completely filling the air cell which in turn will turn diploic or sclerotic. In the second case, see Fig. 3.9(B), the mucosa has recovered its original geometry in most places excepted at the level of the narrow conduit where a strand is created, similar to the ones found in [11]. It should be noted that the occluded air cell may also become diploic or sclerotic. In the third case, illustrated in Figs. 3.9(C), the mucosa has fully recovered its original shape and the air cell is fully functioning.

Figure 3.10: *Different combinations of mucosa adhesion present in the cavity of the air cell. In (A), the adhesion forms a full membrane. In (B, C, D), singular threads are formed instead of a membrane as in (A). In (C) and (D), multiple strands join together.*

Typical types of mucosal strands or folds are illustrated in Fig. 3.10, joining two opposite sides of mastoid air cells. In Fig. 3.10(A), a full epithelial membrane is dividing the air cell into two compartments. Only half the membrane is represented in this illustration. In Fig. 3.10(B), a simple string - also cut in half - is shown. In Fig. 3.10(C and D), a small bundle of epithelial tissue is visible at the crossings of the strings. In Chapter 9, illustrations of these mucosal strands will be provided for real cases. Though they have been reported previously in [11, 74], they have not been illustrated in the bone in 3D.

Before proceeding, it should be noted that this process does not occur in all the air cells at the same time, and all cells are not necessarily hampered with the same degree. To further extend this study, estimation of the range of air cell sizes and their locations in the temporal bone may help understanding the extent of the occlusion when the mastoid is subject to inflammation, which can in turn lead to a further increase in the negative pressure and therefore worsening the pathology, such as in otitis media.

3.9 Otitis media

Otitis media is a group of inflammatory diseases of the middle ear. Two main types of otitis media can be found: acute otitis media (AOM) and otitis media with effusion (OME). AOM is an infection of abrupt onset that usually presents with ear pain, resulting in pulling at the ear, possible fever, and poor sleep for young children. OME is not typically associated with symptoms. OME is defined as the presence of non-infectious fluid in the middle ear that has remained for more than three months. OME is also known as serous otitis media (SOM) or secretory otitis media (SOM), and commonly referred to as glue ear due to an accumulation of serous or purulent fluid that occurs within the middle-ear spaces.

Another type of otitis media is the chronic suppurative otitis media (CSOM), corresponding to a middle ear inflammation of greater than two weeks leading to episodes of discharge from the ear. Complications from an acute otitis media often results in a CSOM. The common cause of all forms of otitis media is a dysfunction of the Eustachian tube, originating from an inflammation of the mucous membranes in the nasopharynx, itself possibly caused by a viral infection or allergies at the level of the upper respiratory system.

In the next chapter, the typical imaging modalities used nowadays (2017) are presented. As will be further explained, though the histological sections remain the optimal imaging technique in terms of information content, the reconstruction from two-dimensional images into a three-dimensional volume is difficult, and requires a physical slicing of the temporal bone. At a very high resolution, X-ray micro-CT scanning allows the obtention of a three-dimensional volume directly without any slicing of the bone specimen but at the expense of losing some information.

Imaging Modalities

4.1 Introduction

Both clinical CT and micro-CT scanners use X-ray technology. Before stating the differences between these two types of scanners, some common features are discussed.

X-ray Computed Tomography (CT) imaging consists of exposing an object with X-rays from multiple orientations, while measuring the intensity decay when going through different materials. This decrease in intensity is described in terms of the X-ray energy, the length of the X-rays path, and the coefficients directly related to the material attenuation. From the signals generated during the scanning, reconstruction using specialized algorithms is necessary to produce the final image data.

The principal components of an X-ray tomography system are:

- an X-ray source,
- a series of detectors that measure X-ray intensity attenuation, usually located on the opposite side of the X-ray source in relation to the scanned object,
- a rotating device either housing the X-ray source or on which the scanned sample is spinning.

Most X-ray CT machines, especially in the clinical world, use an X-ray tube. Other X-ray sources, such as a synchrotron or a gamma-ray emitter can also be used. Relevant characteristics concerning the X-ray tube are listed below:

- the target material,
- the peak X-ray energy,
- the current (expressed in Ampere),

- the focal spot size directly impacting the spatial resolution of the final scan volume.

There exists three main configuration types. As illustrated in Fig. 4.1(A), the X-rays are collimated in a linear fashion and collected by a linear detector array, resulting in a so-called planar beam type of scanning. In Fig. 4.1(B), parallel-beam scanning is performed using a synchrotron beam line as an X-ray source. In cone-beam scanning, Fig. 4.1(C), the linear array is replaced by a planar detector, and the beam is no longer collimated. Compared to Fig. 4.1(A), cone-beam X-ray and parallel-beam scanning do not have a collimator restricting the X-ray path.

Figure 4.1: Three types of scanner: (A) Planar Fan Beam, (B) Parallel Beam, (C) Cone Beam.

In the case of a planar fan-beam configuration, scattering of the X-rays, producing spurious additional X-rays outside the path going from the X-ray source and the detector, can be reduced using collimators. Such linear arrays are more efficient than planar ones, such as in Fig. 4.1(B and C), at the expense of producing a single slice per scan. The aperture of the linear array sets the thickness of the resulting slice. To get a 3D volume, the scanned object needs to be moved normal to the path formed between the X-ray source and the linear detector.

For the parallel-beam type of scanning, a synchrotron beam line is used as the X-ray source. A good feature of parallel-beam scanning is the lack of distortion in the resulting data. However, the width of the X-ray beam limits the size of the object to be scanned. Synchrotron radiation generally has a very high intensity leading to a quick acquisition of the data and objects with a size of up to 6 cm in diameter can be imaged.

In the case of a cone-beam CT scanner, a planar detector replaces the need for a collimator. The data is acquired during a single rotation and then reconstructed into images using a cone-beam algorithm. Some of the

downsides of a cone-beam CT scanner, are the blurring and distortion of the data at the extremities of the object being scanned. When using high energy X-ray for better resolution, X-ray scattering artefacts also hamper the final data.

There are variants from these CT configurations, such as multiple-slice acquisition in which a planar detector is used, but where the generated data are reconstructed using a fan-beam reconstruction algorithm, along with spiral scanning where the sample is displaced during the acquisition.

4.2 Clinical X-ray CT scanner

A conventional clinical CT scanner typically uses a planar fan-beam type of configuration. A typical clinical X-ray CT scanner is illustrated in Fig. 4.2.

Figure 4.2: *Drawing of a conventional clinical X-ray CT scanner.*

Clinical CT has been used daily for medical applications since the seventies. The bed on which the patient is lying is sliding through a doughnut-shaped gantry. The gantry of the scanner houses the X-ray source, located on one side of the ring, and the set of detectors on the opposite side, see Fig. 4.2. The frame holding both the X-ray source and the detectors rotates around the patient. The succession of radiographs taken along the displacement of the bed forms a 3D image. The main parameters the radiology technologist needs to fine-tune, which contribute to minimizing the radiation dose, are:

- the tube current (mA),
- the peak kilovoltage (kVp),
- the pitch (degrees),
- the gantry cycle time (sec.)

The number of electrons accelerated across the x-ray source tube per unit of time defines the tube current expressed in milliamperes (mA). This parameter is not only an important factor for the resulting scan image quality but more importantly for the amount of radiation dose imposed to the patient [40]. Lowering the tube current leads to less radiation dosage, but affects the resulting scan images with an increased noise level [65].

The peak kilovoltage (kVp) corresponds to the energy of the emitted X-rays. Each tissue type has its own range of density which will affect the X-ray beam's attenuation. Larger and/or denser objects will require a higher energy peak voltage to make sure that a sufficient number of photons, from the X-rays, exit the body part being scanned, and then become collected by the detectors. To reduce the radiation dose imposed to the patient, the peak kilovoltage can also be reduced at the expense of an increased image noise reducing the contrast-to-noise ratio, often leading to an increase of the tube current to compensate for less noise [65].

Pitch, expressed in degrees ($^{\circ}$), is a factor mostly related to the moving table, on which the patient lies, in a helical CT scanner. Pitch is calculated as the proportion between the table feed, expressed in centimetres per full rotation of the gantry, and the total width of collimated x-ray beam along the z direction [59]. When increasing the pitch while keeping the tube current per unit of time constant as the table moves, the radiation dose is decreased.

Decreasing the radiation dose can also be achieved by decreasing gantry rotation time, expressed in seconds (sec.); the faster the gantry rotation, the lower the dose. Doubling the speed of rotation per full rotation of the gantry reduces the dose essentially by half. As for the previous parameters, decreasing the dose level while tuning the gantry rotation time leads to an increase in image noise.

The scanning time is significantly longer compared to routine scans, by setting both a small collimator size and a low pitch value to increase the scan resolution. The effective radiation dose for a head CT for temporal bone imaging is usually around 2 millisievert (mSv).

CT scanning of the temporal bone is a common routine still nowadays. It enables diagnostics of bone fracture, detecting tumors in the bone, and inflammations such as cholesteatoma. It is also used during pre-operative planning and for post-operative assessment. Patient follow up can also involve regular CT scanning. The literature on CT scanning of the temporal bone is broad. To name a few, two studies concerning the general aspect of CT scans of the temporal bone are provided in [71, 81]. In terms of research, CT scanning is a great tool for gaining knowledge about the anatomy of the ear in both normal and pathological cases, thus enabling more in depth

understanding of the physiology governing the temporal bone. More related to the mastoid, several studies have used CT as the main modality [37, 46, 64, 79, 80, 88]. A clinical CT scan of a patient's head, as well as a clinical CT scan of a temporal bone specimen alone, are illustrated in Fig. 4.3.

Figure 4.3: *A clinical CT scan of a patient's head and of a temporal bone specimen. The yellow square reveals the mastoid of the patient of the right ear (always seen on opposite direction in CT images). Clinical CT scan of a temporal bone specimen giving more details about the mastoid air cell system through adjustment of the scanner settings. Resolution: around 0.256mm in all directions.*

4.3 Cone-Beam CT scanner

A more recent X-ray scanning device, dating from 1988 originally targeted towards dentistry, is the cone beam CT scanner (CBCT). In some hospitals, CBCT became a standard diagnostic tool when imaging the temporal bone. CBCT is being more and more used in diagnosis and treatment planning in fields such as dentistry for oral surgery. In ENT, CBCT is also involved in preoperative planning or postoperative assessment in for instance sinuses and nasal fossae or inflammatory pathology, where effusion or mucosal thickening take place or in the visualization of how well pneumatized a temporal bone is, [34, 61]. In the presence of chronic otitis with a dense pars petrosa and opaque middle ear, cholesteatoma and related inflammations can be diagnosed using CBCT. Similar to [2], temporal bone fracture can also be studied via CBCT scans.

The CBCT scanner rotates around the patient's head, obtaining up to about 600 slices. Similar to a conventional clinical CT scanner, the scanning software collects the data and reconstructs it, producing a three-dimensional representation of the scanned body part. As its name stands for, the fan beam corresponds to Fig. 4.1(C). The resolution of a CBCT

scanner ranges from $125\text{ }\mu\text{m}$ down to $80\text{ }\mu\text{m}$ in all directions. A downside of CBCT is the increase of noise as voxel size and slice thickness decrease. A simplistic representation of the CBCT scanner is provided in Fig. 4.4.

Figure 4.4: A CBCT scanner where the patient sits on the chair and the X-ray and detector spin around his/her head.

However, though available for the study, even higher resolution scans were needed during this study and therefore CBCT scans were not used. Another type of images, extensively used in medicine, are the histological sections. In the next section relevant information about histological sections is given.

4.4 Histological sections

Histology is the microanatomy study of cells and tissues. It is commonly performed through examination under a microscope, either light-based or electron-based. After the specimen has been sectioned using a microtome, stained, and mounted on a microscope slide, pictures are taken from the field of view. Histology is considered as a daily routine in biology and medicine. This is especially the case when investigating the presence of disease, such as a tumor, where a small sample is taken from the patient and analysed for histopathology by pathologists. Similar procedure can be applied for basic science with a subpart of an organ or bone specimen. The Temporal Bone Histology Laboratory in Zurich, Switzerland, performs such histological sectioning of human temporal bones. Prof. Dr. Med. Alexander Huber was kind enough to provide such high detailed histological

sections of the mastoid air cell system from the temporal bone. Fig. 4.5 is a histological section from the Temporal Bone Histology Laboratory, where the mastoid air cells are visible on the right hand side mostly at the lower right.

Figure 4.5: A histological section of a temporal bone specimen. Legend: (FN) facial nerve, (MACS) mastoid air cell system, (SCC) semi-circular canals, (SS) sigmoid sinus, (TS) trabecular spaces, (TYMP) tympanum. Source: Temporal Bone Histology Lab, Otorhinolaryngology, University Zurich (permission granted by Prof. Dr. Med. Alexander Huber).

Figure 4.6: Magnified region from the blue rectangle shown as a reference on the lower right corner from Fig. 4.5. Source: Temporal Bone Histology Lab, Otorhinolaryngology, University Zurich (permission granted by Prof. Dr. Med. Alexander Huber).

When only looking at the magnified region marked with the blue rectangle, more accurate information about how the air cells are built can be displayed, see Fig. 4.6. The advantage of using histological sections is that they provide very high details of the structures due to the high magnifications. For instance, in Fig. 4.6, the calcification layers of the bony septae forming the walls of the more medial mastoid air cells is visible. Capillaries as well as red blood cells stuck into vessels provide information about the local vascularization. Mucosal strands, which will be discussed in the last chapter of this thesis, appear to separate one air cell into two cells as the one shown in the upper left corner of Fig. 4.6.

To obtain such high detailed images, the temporal bone is usually fixed in 10% formalin, decalcified using ethylenediaminetetracetic acid (EDTA), embedded in celloidin (paraffin) and serially sectioned with a certain thickness, i.e. 20 microns. Decalcification is a technique for removing mineral from the bone together with other calcified tissues. EDTA works by capturing the calcium ions at the surface of the apatite crystal, allowing a slow reduction of the sample size. Though the process is very slow, it makes it possible to produce very high quality images. The section is then stained with hematoxylin and eosin, also known as H&E, to reveal structures of interest. H&E is one of the most common stains in histology, and is often considered as a golden standard when used in medical diagnosis for instance in the search for a suspected cancer. The staining involves the application a complex formed from aluminium ions and hematein called hemalum, turning the cell nuclei and calcified materials blue. The eosin counterstaining, eosin being an acid aniline dye, colors structures in various shades of red, pink and orange. Among the structures highlighted while using H&E staining, muscles will appear in dark red, erythrocytes in cherry red, collagen such as in decalcified bone matrix in pale pink.

In agreement with [89], small degrees of physical deformation may be introduced during the sectioning; including stretching and folding. Unless taken care of before taking the pictures, resolving these distortions can be very challenging, if not impossible. However, using histological sections has its own drawbacks. Depending on how the pictures were taken, blurring can occur due to slightly wrong focus. Note that this limitation is reduced greatly nowadays. Another issue using histological sectioning is the lack of the third dimension, i.e. the accumulation of consecutive images. When available, the sections are often separated over a certain distance and therefore interpolation techniques need to be used to assemble the sections so as to form a final 3D model. The tiny bony walls forming the air cells may also break during the sectioning, leading to a final incomplete 3D reconstruction.

As a side note, a 3D printed representation of the mastoid air cells based

on contours drawn from histological sections over an entire temporal bone as early as in 1938, see [4] for more details on how it was produced.

The challenge to obtain histological sections in a consistent manner led to considering another type of scanning, to represent the mastoid air cell system in a resolution as high as possible but with the advantage of easier 3D reconstruction, namely X-ray micro-CT scanning. In the next section the micro-CT used during the study is presented.

4.5 Micro-CT scanner

Micro-CT provides much higher-resolution and quality scans compared to a conventional clinical CT scanner, but obviously provides less details compared to histological sections. From a general setup, instead of a rotating of the X-ray source and the detectors around the patients or object to be scanned, the object is placed on a rotating platform while being scanned, but the principle is similar to CBCT and the fan beam corresponds to Fig. 4.1(C).

Figure 4.7: *Picture of the micro-CT scanner used in this study, from the department of Physics and Astronomy at Ghent University in Belgium [57].*

In this study, a custom-built micro-CT scanner available at the department of Physics and Astronomy, Ghent (Belgium), was used for scanning of 8 bone specimens applying the scanning protocol as in [56]. Fig. 4.7 shows the micro-CT scanner with the X-ray source located to the left, the turntable placed in the middle on which the sample is carefully positioned, and the

flat panel detector. Each bone specimen was positioned in a plastic cup on a computer controlled turntable, so as to avoid displacements during rotation while being scanned.

Depending on the sample size, the peak kilovoltage of the X-ray source was set to 120 kVp with 3 mm of aluminium of beam filtration, see Paper II for details. A full cone beam rotation scan was performed in 2 s of exposure per projection, to ensure a high signal to noise ratio. A series of about 2100 cross-sections was reconstructed using a software package called Octopus, developed locally at the department of Physics and Astronomy in Ghent [87].

The resolution of the final images depends on several components such as:

- the X-ray detector resolution,
- the focal spot size,
- the geometric magnification,
- the filtering algorithm used to reconstruct the images,
- the stability of the turntable on which the specimen is placed.

While dose radiation is an important factor to minimize to preserve the health of the patient being scanned, scanning a temporal bone specimen is not affected by such a factor. However, a clinical CT scanner has some pre-settings that cannot be overloaded by the radiological technologist. Therefore, typical scan settings were used even for the bone specimens. A slice from a micro-CT scan of a temporal bone specimen is illustrated in Fig. 4.8.

The signal-to-noise ratio (SNR), determining the quality of the CT scan, depends on the total exposure. SNR is proportional to the tube current and the total exposure time. The SNR also increases with kilovoltage (kV) but not proportionally. In medical scanners high power is used, expressed in kilowatts, but a full scan is performed in a few seconds. In micro-CT scanning, low power is used, expressed in watts, but the specimen is exposed for a very long time instead, typically a thousand seconds.

Estimated price

The price is also an important parameter. A medical scanner is around 1,000,000 € and possibly more. On the other hand the price of a typical micro-CT scanner can be anywhere between 50,000 € and 1,000,000 €, but 250,000 € is a common average price.

The price of a single micro-CT scan is charged per hour of actual measurement time at the department of Physics and Astronomy in Ghent, and is

Figure 4.8: *A micro-CT scan of a temporal bone specimen. Resolution: $36\text{ }\mu\text{m}$.*

fixed to 100 €. Clinical CT scans would have a similar price.

After presenting the equipment used to produced the data together with other modalities used when imaging the temporal bone, the next chapter introduces the motivation of this overall study based on observations made from micro-CT scans, eventually supported by histological sections.

Observations from the data and aims of the thesis

5.1 Introduction

In the literature, the focus has been on the mastoid air cells as an overall system, [37, 46, 64, 79, 80, 88] and Paper I in this study. However, very little is known about the air cells when studied as a single entity. Interestingly, during this overall study, many questions were asked about the shape and size of a single air cell. Knowing the shape and size of an air cell alone is not so relevant, but to gain information about how much they vary in shape and size could provide information on the effect of a surrounding structure locally. For instance, large air cells are thought to have very few constraints until some boundary conditions are large enough to stop their evolution in size. Also, some large air cells may be the result of smaller air cells combining into one. It should be kept in mind that the mastoid bone is originally very flat, and possibly had similar characteristics as the cranial bone.

This chapter is an attempt to address these questions and proposes a preliminary study of how a mastoid air cell can be better characterised. Four main questions, that were most frequently asked, are stated in the remaining of this chapter with a brief answer to each of them. But before, there is a need to define what is the midline of the mastoid air cell system, since it will be frequently used in this chapter.

5.2 Midline of the mastoid air cell system

The antrum, the transition between the tympanum and the mastoid air cell system, forms a major part of the midline. However, due to the variation of the mastoid air cell system from one bone to another, the midline

might be defined as a combination of the antrum along with the central mastoid air cells directly protruding from the antrum, also known as central cells. Therefore, the notion of midline remains a relative geometrical notion rather than just the antrum itself. How far the midline extends vertically towards the apex of the mastoid bone is also relative to the overall pneumatization of the bone under study. Now that the midline is defined, the questions are stated and a brief answer is proposed.

5.3 What is a mastoid air cell in terms of shape?

A cell, from the latin word cell meaning *small room* is used in multiple ways in the British Dictionary as a small simple room commonly found in a jail, convent, or in a monastery. A cell is also defined as a small compartment, found for instance in complex structures such as isolating materials or honeycomb. In the field of biology, the word cell is also used to either define the basic structural and functional unit of living organisms or any *small cavity*. This last definition is the one used throughout this overall study.

Figure 5.1: Surface rendering of a complete MACS together with the airspaces forming the tympanum. Source: [16].

As already introduced in Chapter 3, A. Allam followed by M. Tos classified the air cells within the mastoid bone either based on their location or on anatomical structures in the vicinity, [3, 85]. However, a mastoid air cell could also be classified based on its location in relation to the midline of the MACS: a *terminal* air cell or an *intermediate* air cell. A *terminal* air cell could be defined as an air cell from which there is no more projection from. On the opposite, an *intermediate* air cell has the extra feature of extending to another air cell(s); eventually to a terminal air cell. Terminal air cells

are easily observable at the outskirts of a rendered complete mastoid air cell system, see Fig. 5.1. To ease the reader's understanding, a surface rendering of the air cell system shown in Fig. 5.1 can be interpreted as a resin cast used to reveal their shape without the bone in which the air cells are carved into. It should be kept in mind that some terminal air cells are also present halfway between the outskirts and the midline. Fig. 5.2 illustrates the distinction between the terminal and intermediate air cells.

Figure 5.2: *Distinction between terminal and intermediate air cell.*

Communicating air cells may also be on a single branch, as represented in Fig. 5.3 in a very simplistic manner.

Figure 5.3: *Drawing of a simplistic isolated branch with a communicating air cell at its centre, and a terminal cell to the left separated by a conduit forming a necking.*

From Figs. 5.2 and 5.3, another term is also introduced: *conduits*. From observations on micro-CT scans of temporal bone specimens, a *conduit* is represented as a tubular structure that connects two air cells with each other, whether of communicating type or of terminal type. Their length may vary quite a lot; from very short where the two air cells are almost in contact with each others, to very long with a fairly constant diameter along its main axis in most cases. In the presence of very short conduits, the word *necking* could also be used due to the large variation of diameter

along its main axis. The word pore is also frequently used in other fields, such as stereology or in biology, but to keep the description simple, only the name conduit is used. Another, and more challenging, type of cells are the finger-like type of cells. Though normally seen vertical, Fig. 5.4 gives an idea of how they look like.

Figure 5.4: *Finger-like elongated type of air cells normally observed as vertical in the mastoid, but represented horizontal for more practical purposes.*

To illustrate how these structures appear in real cases, Fig.5.5 provides two different views of an exposed mastoid process from a micro-CT scan using volume rendering.

Figure 5.5: *Complexity of the mastoid air cell shapes revealed from a lateral view (left) and posterior view (right) of the same temporal bone using volume rendering and clipped so as to exposed its internal part. Legend: (red arrows) finger-like elongated air cells, (yellow arrows) short conduits, (blue arrows) necking forming pores between two cells, (green stars) terminal cells in the more medial part of the MACS.*

5.4 What is the size of an air cell?

Park et al. [64], followed by many other other research teams, provided a framework on how to assess the surface area and volume of the overall mastoid air cell system, while using image processing on a stack of CT images, [37, 46, 64, 79, 80, 88]. The estimates of surface area and volume for a single mastoid air cell, though, remain unknown. When observing a surface rendering of the complete mastoid air cell system, as shown in Fig. 5.1, the size of a singular air cell appears to vary considerably. One downside of such full representation is that, it does not allow visualization of the more medial, i.e. more internal, cells and information on how they communicate with the midline. As explained in [13], clipping planes can be used to partially hide some of the rendered data in order to see the more internal parts of the air cell system.

Unfortunately, the complexity of the air cell system is so high that even using clipping planes for visualization, it does not allow a clear understanding on how the air cell system is structured. It is although possible to clip most of the air cell system and only preserve a small fraction to be rendered and study how air cells are locally connected to each others and eventually to the midline. However, the connectivity with the remaining structure is not visible and may lead to interpretation errors. Also, it is possible to perform such a study by carefully studying a limited block of data at a time, and build a final three dimensional map, but besides the fact that it would take a tremendous amount of time to go through the overall volume, a lot of errors may accumulate along the way.

From the new nominations given in the previous question, the original question of what is the possible size of an air cell is changed to what is the size of these structures, or more exactly how large these air cells are locally. When taking into consideration the fact that the air cells might be related to surrounding structures, it is also considered relevant to investigate whether the variation in size of these structures is related to the regions over a full temporal bone as defined by [3, 85].

5.5 Air-mucosa versus mucosa-bone surface area

A striking but not surprising observation is the difference between the surface area of the interface air-to-mucosa interface that is much smaller than the corresponding surface area of the mucosa-to-bone interface. A typical representation from a histological section is illustrated in Fig. 5.6. The mucosa appears thinner in some places while thicker in more complex lo-

cations. This local non-uniform thickness is possibly due to blood vessels, which participate in the gas exchange and therefore regulate the pressure locally.

Figure 5.6: *Nonuniform covering of the mucosa along the mastoid air cell walls.*
 Source: Temporal Bone Histology Lab, Otorhinolaryngology, University Zurich (permission granted by Prof. Dr. Med. Alexander Huber).

A second but less obvious reason is the local thickening of the mucosa resulting from the remnants of a previous inflammation as detailed in Section 3.8. Extensive cold is quite common and therefore it is very common to observe these thickened mucosa from place to place. Therefore, there will be a distinction between the surface area formed by the walls of the mastoid air cells, representing the interface bone-mucosa, and the more functional surface area defined by the interface mucosa-air.

Ideally, estimating the surface area from the air-to-mucosa interface would be preferred. However, on micro-CT scans, the mucosa is too thin and barely visible even for the highest scan resolution obtained in this overall study, i.e. $32\text{ }\mu\text{m}$. Estimation of the mucosa thickness was estimated by Matanda et al. [58] distances were $32\text{ }\mu\text{m} \pm 1.26$ ($\text{mean} \pm \sigma$) in the normal mucosa. As a side note, it is believed that even for locally enlarged mucosa, the distance between air to blood, an important factor for gas exchange, is kept constant and therefore this differentiation is not considered critical in this study. Also, the thickness of inflamed mucosa was reported to $22.32\text{ }\mu\text{m} \pm 0.59$ ($\text{mean} \pm \sigma$) but does not reflect the swollen epithelial lining due to the fact that it does not not impact the gas exchange from air to blood. However, from a surface area point of view, the surface area of the air-to-mucosa interface is certainly smaller than the surface area estimated from mucosa-to-bone interface. As previously stated in Sec. 3.8, the mucosa can swell up to 80 times its original thickness.

5.6 Discovery of micro-channels

Another observation is the discovery of numerous micro-channels found at the cortex of the mastoid bone between the exterior and the lateral air cells, but also present between the air cells. These micro-channels have not been described in previous literature. These micro-channels are likely to represent a separate blood supply for the mastoid mucosa, and have important implications for further understanding the middle ear physiology including its pressure regulation and gas exchange.

Figure 5.7: *Volume rendering of a temporal bone specimen with a special transfer function, in order to reveal the micro-channels running in the bone between the mastoid air cells and outer layer of the bone. Clipping planes were used to expose the internal parts of the bone. Both the use of special transfer function and clipping planes are detailed in Sec. 6.4.*

To resume this section, the advantage of having at disposal micro-CT scans of complete mastoid air cell systems from several temporal bone specimens, and to compare the results from previous studies where conventional CT scanning was used is considered a high priority since more and more ENT departments perform such pneumatization studies involving surface area and volume estimation. Performing a similar study using micro-CT scans, which is an uncommon procedure if not unique, may provide useful information for the rest of the community. The discovery of micro-channels,

though not related to the size of the air cells, needs further investigation and is also considered a high priority in terms of new knowledge to bring to the field. Investigation of the shape of the mastoid air cells as singular entities is also an important topic, but belongs more to basic science in terms of its use in the field. Also, studying the shape of singular air cells in order to obtain a size distribution over a complete mastoid bone requires more tools from different fields. The next section provides a timeline on the research has been carried out during the overall study.

5.7 Aims of the thesis

From the previous sections, the aims of this overall PhD study is twofolds: (1) Geometrical aspects of the mastoid air cell system related to pneumatization of the mastoid, (2) Interpretation and geometrical description of the discovered micro-channels penetrating the mastoid process. For each topic, a set of aims are defined and presented in the following.

Geometrical aspects of the mastoid air cell system

Aim 1. Determination of the surface area and the volume of the mastoid air cell system from X-ray micro-CT scans is necessary for estimating the pneumatization of the mastoid. The goal is to investigate whether scanning ex-vivo temporal bone specimens using X-ray micro-CT allows the detection of a broader range of mastoid air cells in terms of size, especially towards air cells of smaller size, thus affecting the estimated surface area and volume measures reported in previous literature.

Aim 2. Reducing the complexity of the mastoid air cell system representation. Because the mastoid air cell system is very complex structure-wise, its binary representation remains complex in shape. It is believed that by reducing the complexity of the binary representation, while respecting the mastoid air cell system topology, analysis beyond the estimation of the surface area and volume such as a 3D shape analysis is possible.

Aim 3. Investigate in a simplistic manner how inflammation of the mucosa lining mastoid air cells could impact the level of pneumatization of the mastoid through a 3D shape analysis. From the previous chapters, under inflammation, the mucosa lining the mastoid air cell might expand to a level where some air cells will be clotted. Depending on how large the mastoid air cell system is locally, part of the air cell system may become temporally or permanently dysfunctional. Measurements on the local size

of the mastoid air cell system is therefore relevant to have.

Aim 4. Investigate whether the mastoid air cells do form groups or regions as described by A. Allam and M. Tos. As described in the previous chapters, from observation of exposed mastoids, the mastoid air cells have been grouped based on their location or according to neighbouring structure such as the sigmoid sinus. However, there has not been any systematic way of confirming or rejecting this classification. A tool is therefore necessary to investigate this question.

Interpretation and geometrical description of the discovered micro-channels

Aim 5. Interpretation of micro-channels in the cortex of the mastoid process. Though already presented as an observation in the previous section due to the high relevance, a more in depth description both in terms of their shape and their possible physiological role is needed.

Aim 6. Investigate the geometrical properties of the micro-channels. Discovery of a new anatomical structure such as the micro-channels penetrating the bone of the mastoid process always leads to questions relating to their geometry. Investigating their geometrical properties is therefore important and further helps gaining knowledge about their potential role and possible connection with surrounding structures.

Aim 7. Enhance the representation of micro-channels to get a better definition of their course and connectivity. From the previous chapter, though the X-ray micro-CT provides high resolutions scans, the resolution imposed by scanning the complete mastoid air cell system, together with artefacts commonly found using this technology, corrupt micro-channels of smaller sizes. Incomplete representation of the micro-channels can lead to misinterpretation of their connectivity and their implicit role from a physiological point of view. Enhancement of their course is therefore needed for future analysis.

In the following two chapters, the most relevant image processing tools needed to fulfil these aims are presented. In Chapter 6, image processing tools commonly used in the field of ENT are introduced. Then follows the description of more advanced image processing tools used at several occasions in this study. An extra section is dedicated to volume rendering, clipping planes, and special masking. In Chapter 7, tools that do not belong to image processing but relate to 3D shape analysis are introduced.

Image Processing

6.1 Introduction

This chapter provides a theoretical background of image processing, from basic tools to more advanced ones. Common image processing tools often encountered in clinical studies are resumed first. The reader can find a broader coverage in the literature such as in [25, 38]. More advanced methods are then presented, and form the main framework for this overall study. A section about volume rendering is also included in this chapter, since it relates to visualization of results from image processing along with original data.

6.1.1 Segmentation by thresholding

This section only deals with the type of segmentation used in this field of hearing research. Many other segmentation methods exist, such as the local adaptive threshold, but their description is beyond the scope of this thesis and are seldomly used in the clinical field [91, 93]. The most common technique used in this field to isolate a structure from the rest of the data is called binary thresholding, which is applied in the following manner. Every pixel whose value is above a certain threshold, or within a certain range of values, is interpreted as foreground value, usually set to 1, while all remaining pixels are considered as background, respectively set to 0. Fig. 6.1 shows an example of an image segmented at different thresholds. In this example, the airspaces pixels below the threshold are segmented out and represented as white in the binary mask. As can be observed, the higher the threshold the more structures are segmented along, including the bone marrow in the trabecular spaces when the threshold is above 60.

Figure 6.1: *Binary segmentation of the airspace in the mastoid bone using different threshold values (40, 60, 80, 100, 120) compared to the original data (top left).*

Also, it can be noticed that some of the segmented structures are attached to each others, while they are not supposed to. This arises when the threshold is close to the bone range of grayscale values, and where the segmentation classifies some pixels belonging to bone as foreground (air) and not as background (bone). This type of error is illustrated in the simple representation shown in Fig. 6.2.

At the boundary of a structure, the CT intensity values change from the level of one tissue to that of another tissue. However, this change is gradual rather than abrupt, giving a blurred boundary interface. The interface, in perfect condition, is measured exactly halfway between the two CT number levels, known as half maximum height, or HMM, representing the mean value of the two CT levels from both sides of the transition. Binary segmentation does not use any a-priori knowledge of the underlying structures to be segmented. Therefore, all pixels or voxels above a certain value from different structures will be segmented at the same time, forming different components either connected or disconnected.

Figure 6.2: *Left: single structure, middle: two structures far apart leading to the segmentation of the two structures, right: the structures are too close, which leads to a single object in the segmentation result. Note that the pixels with a value below the threshold are set to white.*

Because segmentation using the binary thresholding technique does not contain information of the structure to keep, a manual selection of that structure needs to be done, see Fig. 6.3 to the right. The only requirement is to make sure that the structure is disconnected from all segmented structures. The selection was performed using a region growing tool, by first placing a seed in the structure of interest, and letting the algorithm segment all voxels belonging to the structure [36].

Figure 6.3: *Selection of a region of interest (right) from the binary segmentation (middle) of the original data (left).*

The result of the binary segmentation is often postprocessed using different morphological operations, which are described in the following section.

6.1.2 Morphology on binary images

Results from a binary segmentation can lead to misdetected or over-segmented regions, while the remaining parts appear acceptable. One solution to circumvent these issues is through the use of morphological operations. Four morphological operators are commonly used: dilation, erosion, opening,

and closing. Dilation consists primarily of enlarging the boundary of all the segmented regions in the data, see Fig. 6.4.

Figure 6.4: Results from dilation operation (right) compared to the original binary image (left). As can be observed, each object grows in size.

Erosion, on the contrary, consists of shrinking the boundary of all the segmented regions in the data, see Fig. 6.5.

Figure 6.5: Results from the erosion operation (right) compared to the original binary image (left). Contrary to Fig. 6.4, as can be observed, each object shrinks in size.

The last two morphological operators, namely opening and closing, are in fact combinations of dilation and erosion in different orders.

Opening, corresponding to a succession of erosion followed by dilation, can be seen as a morphological noise filtering. The small structures are first removed from the binary data, and then all shrunk boundaries from the erosion part are restored back to their original shape by the dilation step. An example of applying a morphological opening on binary data is shown in Fig. 6.6.

Figure 6.6: *The results from an opening operation (right) compared to the original binary image (left). Note that many small structures have disappeared.*

Closing is, on the other hand, a succession of dilation and erosion. While opening removes noise in the binary data, closing can be interpreted as a filling small cavities present inside a binary structure. An example of applying a closing operation on a binary image is given in Fig. 6.7.

Figure 6.7: *The result of a closing operation (right) compared to the original binary image (left), see the gray squares where merging occurred.*

Despite a simple implementation and its usefulness for very simple shapes, such as a cube or a sphere in 3D, results from morphological operations are very often unsatisfactory for natural images (i.e. medical images).

In some cases, even after a single pass of dilation or closing, some unwanted holes reside in the segmented object. A hole can be defined as a region of pixels only having a background value (0), surrounded by a border of pixels with foreground value (1). The algorithm, in that case, iteratively changes the connected background pixels to foreground pixels and stops when the

boundaries defining the hole of the segmented object are filled, see Fig. 6.8.

Figure 6.8: *A (top left): binary segmentation of the bone, B (top right): results from the hole filling where some holes were not processed due to incomplete borders (sides of the image), C (bottom left): Modified version of A with the left, top, and right sides of the image covered with a foreground value line. D (bottom right): successful filling of the holes in the bone.*

The cavities present in the bone in Fig. 6.8(A), when fully surrounded by foreground value, are fully covered with foreground value, see Fig. 6.8(B). However, it can be noticed that some cavities were not filled due to the boundary not fully surrounding the cavity. To illustrate this problem, vertical and horizontal lines were added to the side of the image where bone is present, see Fig. 6.8(A). This step ensures that all cavities would be fully be surrounded by a boundary, see Fig. 6.8(C). Only the lower right part of the image was not covered with a foreground line. When applying the hole filling algorithm on Fig. 6.8(C), all cavities inside the bone are then

successively fully filled, see Fig. 6.8(D).

Such a technique can further help to extract a structure, such as the bone, with all its content. However, when looking at the binary results, it is almost impossible to know what lies under the voxels added during the morphological step(s) being used. A possible way to verify what information lies under these added pixels/voxels is to reassign the original data over the binary data, as explained in the following subsection.

6.1.3 Masking original data over a binary segmentation

A possible way to recover the original information from a binary segmentation, is to simply multiply the binary data by the original data, see Fig. 6.9. All pixels or voxels that have a zero label will remain 0, while all the other pixels or voxels that are of label 1 from the binary image data will recover the original grayscale value.

Figure 6.9: *Reapplying the original grayscale data on the binary data (right), by multiplying the binary segmentation (middle) with the original data (left).*

In an image processing pipeline for visualization, such a method could be relevant. However, when used as an intermediate step, masking from a binary mask may introduce sharp edges on the side of the segmented structures and can further lead to additional errors, and can worsen the results further when drawing statistics from such operations.

It may sometimes become a challenge to fix problems either related to the original image itself, or resulting from one image processing operation. Eventually, it may become impossible to know where the problem comes from. For instance, changing the threshold in the binary segmentation may result in using an iteration of morphological operators, in order to allow the hole filling to work appropriately. Another example arises when seeking the true boundary of a bone. The CT level corresponding to the half maximum height should be measured locally. This is the reason why it is difficult to

decide a global threshold to the overall image.

Global thresholding techniques, such as the one introduced by Otsu ([69]) which is commonly used in clinical research, are very fast and generally produce good results for the sole purpose of visualization. However, when the segmentation is aimed at measuring some metrics accurately, global thresholding will not be adequate. Local thresholding techniques estimate a different threshold for each pixel, using the grayscale information from the neighbouring pixels. However, local thresholding is rarely used in clinical research, possibly due the need for replicability of their study in relation to others or simply by not knowing its existence.

Also, an early stage segmentation discards a lot of information from the original data that could be used in further processing. A good practice is to perform a segmentation at the very last stage, where all information from the data needed for the study have been taken into account.

As the mastoid air cell system is now isolated from all other structures in the data, metrics such as surface area and volume can be estimated, as explained in the following.

6.1.4 Measuring surface area and volume

Estimating the volume from 3D data is fairly straightforward. It is usually done through counting the number of voxels belonging to the structure, and multiplying this number by the voxel size. Voxels containing a mixture of several tissue types within the same voxel, also known as the partial volume effect, also need to be included since they contain a portion of air in them. A solution is to reapply the original data on the binary segmentation through the masking operation, as explained previously. The result from the masked data is then inverted, so that it can be used as a percentage representation of the content of air in the mastoid air cells.

When the voxels only consist of air (white), it will represent 100% of the voxel volume. Additionally, when a voxel contains a mixture of tissues, the volume of air in that voxel will provide a percentage proportional to the amount of air in it. As an example, if a voxel contains 50 % of both air and soft tissue, the volume of air will be 0.5 of the voxel volume corresponding to air. This method produces a slightly more accurate volume estimate, compared to directly using the binary segmentation. It should be noted that the presence of noise, as in in Fig. 6.9, can also have an influence on the volume measurement for such a technique, but is still believed to be a more accurate measure than estimating the volume from the binary segmentation.

Figure 6.10: *Surface mesh of a complete mastoid air cell system generated from a regular clinical CT scan (512×512 pixels with a slice thickness of 0.625mm).*

The surface area, however, is more challenging to estimate. One technique, commonly used in many medical image processing applications, is to fit a triangular mesh from a binary segmentation [53], from which it is possible to measure the surface area. For each triangle composing the mesh, the area is computed and the sum of each triangle area is then accumulated, to give the total surface area over the whole mesh. Fig. 6.10 provides an illustration of such a triangular mesh, from which the surface area can be estimated. Using such a triangular mesh, fitted to a binary segmented volume, depends heavily on the resolution of the original data. If the resolution is too low, a triangular mesh will be composed of relatively large triangles, and will therefore give a rough estimate of the surface area. When the original data is of much higher resolution, the triangles composing the mesh will be much smaller, thereby giving a much more accurate estimate of the surface area.

All the mentioned methods are based on the grayscale values of the original data. As introduced earlier, noise will corrupt the grayscale intensities and will inevitably lead to using different "tricks" to circumvent some of the encountered issues. Moreover, it may be difficult to define a specific order when using different image processing tools. Nonetheless, it is often necessary to assess the results from one step, go back to a previous step, re-tune the parameters for that step, and move on to the following steps, leading to a hit-or-miss type of global method. While some methods can be used globally for the whole image, some require local adaptivity.

Also, it can be noted that this line of thinking leads to segmenting the structure of interest in order to extract information from it. Another strategy is to define the conditions for the structures of interest to be investigated, and use such a-priori information in order to extract these structures from the original input data.

6.2 More advanced image processing

In this section, a different line of thinking is presented. It relates to extracting all relevant a-priori information about the structure of interest, and then use this information to guide the algorithm to detect these structures of interest. Before defining the conditions to find such a-priori information, several topics need to be introduced to understand this line of thinking. For humans, as for many other species, vision is a very important part of our sensory system. Detecting a special object of interest in the surrounding involves different sources of information such as for instance colours, patterns, and shapes.

In fact, an object is mostly characterized by its size and its edges, whether being sharp or smooth. Lines are also important in the vision system, since they allow us to delineate between different structures or objects. Using such information is therefore essential when looking for a specific structure in a 2D image or 3D volume.

Figure 6.11: *Examples of lines and edges.*

Even if an object of interest globally has a very complex structure, it can locally be represented with very simple structures such as lines and edges, all with different orientations, see Fig. 6.11. A real case is shown in Fig. 6.12.

Such transitions, whether a line or an edge, can be seen as simple signals in 1D. One approach is to generate a one-dimensional template for a line and for an edge, rotate these templates in several directions (2D or 3D), and detect them in the original 2D image or 3D volume (using convolution).

Figure 6.12: *Lines and edges in different regions of an image.*

Relevant tools from digital signal processing can be used when building such filters [45]. As will be explained in the following of the thesis, information about the frequency domain along with the spatial domain is needed. In fact any signal, such as a complete image, can be decomposed in terms of sinusoids with different frequencies. For instance, if a profile is taken along a single image from a data scan, the grayscale intensity profile along that line can be extracted, see Fig. 6.13. It should be noted that a line in 2D corresponds to a sheet or a surface in 3D.

Figure 6.13: *Grayscale intensity profile (right) extracted from the white line overlaid on the original image (left).*

According to Fourier theory, the signal describing the profile can be decomposed into sinusoids each having a unique frequency, see Fig. 6.14. When

combining all these sinusoids with respective frequencies, it is possible to reconstruct the original profile shown in Fig. 6.14. In this example, the reconstruction from three different set of frequencies are shown. It can be observed that the reconstructed profile in the lower graph is very similar to the intensity profile shown to the right in Fig. 6.13.

Figure 6.14: *Reconstruction of the intensity profile given in Fig. 6.13 from three different sets of sinusoids.*

The operation of representing this intensity profile into sinusoids of different frequency was performed through a so-called Fourier transform. The Fourier transform plays a central role in signal and image processing, and is used when defining the template for line and edge detection in this study.

In the following section, all elements related to the design of the line and edge filters are briefly introduced focusing more on the philosophy rather than on the mathematical aspect, though some mathematical expressions will be given. A more complete introduction is given in [26].

6.2.1 Filter design

As introduced earlier, the edges and lines defining the structures can be decomposed in terms of sines and cosines. To pick up sine and cosine information from a natural signal, a solution is to create a filter which is only sensitive to a certain wavelength or range of frequencies. For illustrative purposes, only a single frequency is shown in the following figures. Before

digging into the filter design, the Fourier transform of both the cosine and sine needs to be resumed first. For the cosine, defined as $f(t) = \cos(\omega t)$ with $\omega = 2\pi f$ for a 1D signal in the spatial domain, the Fourier transform becomes

$$\begin{aligned}
 F(u) &= \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt \\
 &= \int_{-\infty}^{\infty} \cos(\omega t) e^{-i2\pi ut} dt \\
 &= \int_{-\infty}^{\infty} \cos(\omega t) [\cos(-2\pi ut) + i \sin(-2\pi ut)] dt \\
 &= \int_{-\infty}^{\infty} \cos(\omega t) \cos(-2\pi ut) dt + i \int_{-\infty}^{\infty} \cos(\omega t) \sin(-2\pi ut) dt \\
 &= \underbrace{\int_{-\infty}^{\infty} \cos(\omega t) \cos(-2\pi ut) dt}_{\neq 0 \text{ when } u = \pm \omega} - i \underbrace{\int_{-\infty}^{\infty} \cos(\omega t) \sin(2\pi ut) dt}_{= 0 \text{ for all } u.} \\
 &= \frac{1}{2} [\delta(u + \omega) + \delta(u - \omega)].
 \end{aligned} \tag{6.1}$$

The Fourier transform of a cosine is illustrated in Fig. 6.15, where two impulse responses are located at $-\omega$ and $+\omega$.

Figure 6.15: *Fourier transform of a cosine.*

Similarly for the sine wave, defined as $f(t) = \sin(\omega t)$ with $\omega = 2\pi f$ in 1D in the spatial domain, the Fourier transform becomes:

$$F(u) = \frac{1}{2} i [\delta(u + \omega) - \delta(u - \omega)]. \tag{6.2}$$

As can be observed, the Fourier transform of the sine is imaginary, see Fig. 6.16.

Figure 6.16: *Fourier transform of a sine.*

Fig. 6.17 shows side by side, using three axes, the Fourier transform of $\cos(\omega t)$ and $\sin(\omega t)$. The Fourier transform of the cosine lies along the plane formed by the real axis and the frequency axis. The Fourier transform of the sine lies on the plane formed by the imaginary axis and the frequency axis, normal to the real plane.

Figure 6.17: *Fourier transform of cosine located along the real plane (left), and Fourier transform of the sine is located along the imaginary plane, normal to the real plane (right).*

To avoid involving complex notation, a possible solution is to perform the Fourier transform of $i \sin(\omega t)$ instead of $\sin(\omega t)$, leading to a rotation of the Fourier transform of the sine along the real plane instead of the imaginary plane, see Fig. 6.18

Figure 6.18: *Fourier transform of cosine and the complex sine ($i\sin$) are both lying along the real plane.*

Obviously, since the Fourier transform of a cosine and a sine each produce only a set of impulses for both ω and $-\omega$ whether positive or negative, using such an impulse response to design a filter would make no sense, since their inverse Fourier transform would have an infinite support in the spatial domain. A filter by definition needs to be local, in order to extract local information from an image. To achieve locality, by forming a pulse train around a centre frequency and adding the contribution of each impulse then produces a filter that becomes more and more local. As an example, observe a cosine and a sine for 5 different frequencies, see Fig. 6.19.

Figure 6.19: *5 cosines (left) and sines (right) with the same respective frequencies, represented both in the spatial domain and in the frequency domain. Only one period is represented for clarity purpose, but each impulse in the frequency domain represents a infinite long cosine or sine of the specified frequency.*

When adding cosines with different frequencies together, a major signal

component appears at the centre with some ripples on both sides, see Fig. 6.20 (left column). The same applies for the sine components added together, Fig. 6.20 (right column).

Figure 6.20: Summation over 5 frequencies in the spatial domain for both the cosine (left) and sine (right), and representation of the impulses that were used during the summation.

The quest of localization for the filter is then to suppress the ripples on both sides of the signal, while keeping the major component. To give more emphasis on the main lobes at the centre of the signals in Fig. 6.20, while attempting to discard the ripples, a weighting window or function can be applied to the impulses used during the summation. In Fig. 6.21, two different types of weighting are applied. A first weighting function is made linear corresponding to a pyramid, with for instance a weight of 0.1 for the outmost frequency in relation to ω (respectively $-\omega$), a weight of 0.5 is given to the innermost impulse, while the middle impulse, ω or $-\omega$, is given a weight of 1. In the second part a non-linear weighting function (Gaussian) is used instead of a linear one, which for this example has the weights: 0.2, 0.7, 1, 0.7, and 0.2 respectively.

The ripples are minimized by the nonlinear weighting function, making the filter in the spatial domain very localized. Overlapping the resulting signals from the non-linear weighting - i.e. cosine and sine - in the spatial domain, we can observe a good match in localization, see Fig. 6.22.

Figure 6.21: Linear and non-linear weighting functions applied to the impulses before summation, making the signals in the spatial domain more localized.

Figure 6.22: Overlapping the results from the non-linear weighting of both the cosine (even) and sine (odd).

From Fig. 6.22, the left signal is even in relation to its centre and is defined in the real part as illustrated in Fig. 6.17 to the left. The second signal to the left is odd in relation to its centre, and is defined in the imaginary part as illustrated in Fig. 6.17 to the right. If the even part is denoted $f_e(x)$ and the odd part denoted $f_o(x)$, we can see that $f_o(x)$ is the Hilbert transform of $f_e(x)$. This property leads to the notion of quadrature, where these two components form a complex pair in the spatial domain. The real part is defined as a line filter and the imaginary part is defined as an edge filter. The cosine can be seen as a line detector, and the weighted sine can be seen as an edge detector. Together, these two filters form a quadrature filter. When convolving the quadrature filter over an image, when a line structure is met, the response of the filter will be dominant in the real part, while when the filter encounters an edge, it will produce an imaginary response. It is possible to map the type of structure encountered in the image, using the so-called local phase and local amplitude or magnitude [43]. The local

phase is defined as

$$\theta(x) = \arctan \left\{ \frac{f_o(x) * f(x)}{f_e(x) * f(x)} \right\} = \frac{\text{Im}(q)}{\text{Re}(q)} \quad (6.3)$$

where $*$ denotes the convolution operation. The local magnitude is expressed as

$$A(x) = \sqrt{[f_e(x) * f(x)]^2 + [f_o(x) * f(x)]^2}. \quad (6.4)$$

The local phase can be mapped on a complex plane where θ represents the angle and A represents the magnitude of the vector, illustrated in Fig. 6.23.

Figure 6.23: Coding of the local phase.

The advantage of using the local phase from quadrature filters is the fact that it is contrast independent. While gradient-based type of filters can most of the time only detect edges formed by rapid transitions between two distinct regions of different intensities, the gradient-based filters will very often fail in detecting smooth edges or edges having low contrast [47]. In the case of a line, hardly detectable using common gradient-based filters, a quadrature filter will provide the same local phase for a strong line as for a weak line. A strong magnitude provides a high confidence in the detected structure type. In locations where the magnitude is close to zero, meaning low confidence, the local phase mainly represents noise. In Fig. 6.23, the argument of the local phase is kept constant to a maximum on the unit circle.

The quadrature filter used in this work is presented in the next section.

6.2.2 The quadrature filter

To pick up energy in all possible orientations, the quadrature filter is built as polar (i.e. spherically) separable in the Fourier domain with an arbitrary but positive radial bandwidth function $R(\rho)$ and a direction function $D_k(\hat{\mathbf{u}})$, according to

$$F_k(\hat{\mathbf{u}}) = R(\rho)D_k(\hat{\mathbf{u}}). \quad (6.5)$$

Mathematically, the quadrature filter is in the frequency domain, $F_k(u)$, as further detailed in [43] is defined as

$$F_k(u) = \begin{cases} R(\rho)(\hat{\mathbf{u}}^T \hat{\mathbf{n}}_k)^2 & \text{if } (\hat{\mathbf{u}} \cdot \hat{\mathbf{n}}_k) \geq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (6.6)$$

where u is the coordinate vector in the Fourier domain, $\hat{u}$ is a unit vector directed along u , $\mathbf{n}_k$ is the direction of filter k , and $\rho = |u|$. The radial basis function and the direction function are detailed in the following subsections.

The radial basis function

The radial function $R(\rho)$ is typically designed as a bandpass function based on a center frequency and a bandwidth. A suitable radial function for estimating local image structure is based on the log normal function, a normal distribution in the logarithmic scale, i.e.

$$R(\rho) = e^{-\frac{4}{B^2 \ln^2} \ln^2(\rho/\rho_0)} \quad (6.7)$$

where B is the relative bandwidth of the filter and ρ_0 is the center frequency of the filter, illustrated in Fig. 6.24.

Figure 6.24: The lognormal function defined by the bandwidth B , the low and high cutoff frequencies ρ_L and ρ_H defining the center frequency ρ_0 .

The radial basis function determines in what scale the filter has its sensitivity, meaning what the filter will detect. Practically, the center frequency and the bandwidth, are respectively given by:

$$\rho_0 = \sqrt{\rho_{\text{low}} \cdot \rho_{\text{high}}}, \quad (6.8)$$

and

$$\beta = \frac{\log_n(\frac{\rho_{\text{high}}}{\rho_{\text{low}}})}{\log_n 2} (dB), \quad (6.9)$$

with ρ_{high} and ρ_{low} corresponding to the high and low frequency bounds defining the bandwidth of the filter. The centre frequency corresponds to a wavelength (measured in pixels) of

$$\lambda_0 = \frac{2\pi}{\rho_0}. \quad (6.10)$$

The spatial size of the filter needs to be at least one wavelength, and in practice 1.5 or even 2 wavelengths. Table tab:Chap6Table1 show the spatial size for three different center frequencies.

centre frequency μ_0	wavelength λ_0	linewidth($\lambda_0/2$)
π	2 pixels	1 pixels
$\pi/2$	4 pixels	2 pixels
$\pi/4$	8 pixels	4 pixels

Table 6.1: Correspondance between the centre frequency, the wavelength, and the width of the line to be detected.

The direction function

The direction function, $D_k(\hat{\mathbf{u}})$, varies as $\cos^2(\phi)$, where ϕ is the angle between $\mathbf{u}$ and the filter direction $\hat{\mathbf{n}}_k$, given by

$$D_k(\mathbf{u}) = \begin{cases} (\hat{\mathbf{u}} \cdot \hat{\mathbf{n}}_k)^2 & \text{if } (\hat{\mathbf{u}} \cdot \hat{\mathbf{n}}_k) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (6.11)$$

In 3D, it is common to use six quadrature filters along different directions, see [26] for further explanations about this specific choice. The six directions are defined as

$$\begin{aligned} n_1 &= c \begin{pmatrix} a & 0 & b \end{pmatrix}^T, \\ n_2 &= c \begin{pmatrix} -a & 0 & b \end{pmatrix}^T, \\ n_3 &= c \begin{pmatrix} b & a & 0 \end{pmatrix}^T, \\ n_4 &= c \begin{pmatrix} b & -a & b \end{pmatrix}^T, \\ n_5 &= c \begin{pmatrix} 0 & b & a \end{pmatrix}^T, \\ n_6 &= c \begin{pmatrix} 0 & b & -a \end{pmatrix}^T, \end{aligned} \quad (6.12)$$

with $a = 2$, $b = 1 + \sqrt{5}$, and $c = \frac{1}{\sqrt{10+2\sqrt{5}}}$.

The combination of the radial function and the angular function is summarized in Fig. 6.25. The radial function is by itself smooth, but to get

a smooth transition between the different quadratures, a $\cos^2$ function is used as a directional function.

Figure 6.25: Sketch of the radial and angular functions for different directions k , with their respective smooth variations.

Without the $\cos^2$ weighting, the sharp transitions would result in undesired effects on the final output data. In 3D, a spherical shell with a thickness corresponding to the bandwidth of the radial function would be formed. A pseudo-representation of the radial function in 2D is given in Fig. 6.26, illustrating the lognormal function used as a radial function viewed from above.

Figure 6.26: Pseudo-representation of the lognormal radial function in 2D.

The magnitude of the quadrature filter output is computed as

$$q_k = \left\| \frac{1}{2\pi} \int S(\mathbf{u}) F_k(\mathbf{u}) d\mathbf{u} \right\|, \quad (6.13)$$

where $S(\mathbf{u})$ is the Fourier transform of the original data set.

Because both spatial and frequency domains are considered at the same time, an operation called filter optimisation is needed [44, 45]. An example of the real and imaginary parts of an optimised quadrature filter is given in Fig. 6.27, and the corresponding optimized quadrature filter in the frequency domain is illustrated in Fig. 6.28.

Figure 6.27: *Real and imaginary part of an optimized quadrature filter: left) real part, b) imaginary part.*

Figure 6.28: *An optimized quadrature filter in the frequency domain.*

For the local neighborhood of each visited voxel, the filter response provides a contribution from six directions. . The contribution of the filter response from each direction of detected local structure needs to be stored, for each visited voxel in the volume. Such a contribution cannot be represented by a point or a vector alone. A structure to englobe the total information for each voxel is therefore needed.

The need of storing the contribution from each direction for each voxel has been encountered in many fields, where physical properties need to be locally described. A mathematical field has emerged where tensors can represent all the information in each voxel. The notion of tensor should not to be confounded with the muscle type, such as the tensor tympani used to dampen sounds generated while chewing. Derivation of the tensor analysis is briefly stated in the following section.

6.2.3 Tensor analysis

While a vector is defined as a list of values, and a matrix is defined as a table of values, the natural evolution of such taxonomy would be a list of tables or a table of lists. Therefore, tensors can be seen as a generalization of matrix of N dimension. When considering more than one direction, such as in the case of the quadrature filtering with the contribution from six different directions, a tensor is an adequate solution.

The next step in this study is to compute the structure tensor, which in each voxel contains information about the local structure. The naming of structure tensor will be used during the rest of the thesis. The estimate of the structure tensor can now be constructed as

$$\mathbf{T}_{est} = \sum_{k=1}^6 |q_k| M_k, \quad (6.14)$$

with $|q_k|$ being the magnitude of the quadrature filter output, and M_k is defined as

$$\alpha \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T - \beta \mathbf{I} \quad (6.15)$$

where $\hat{\mathbf{n}}_k$ is the direction of quadrature filter k , $\mathbf{I}$ is the identity matrix, α is $5/4$, and β is $1/4$ in the 3D case, see [26] for more details on the choice of α and β . Because the tensor matrix is symmetric, the tensor components can also be written as follows (for the 3D case):

$$\mathbf{T} = \begin{pmatrix} t_{11} & t_{12} & t_{13} \\ t_{12} & t_{22} & t_{23} \\ t_{13} & t_{23} & t_{33} \end{pmatrix} \quad (6.16)$$

The advantage of using a tensor resides in the possibility of extracting the eigenvalues (λ_i) and eigenvectors ($\hat{\mathbf{e}}_i$) from the local structures. The eigenvalues and eigenvectors inform us about the energy distribution of the signal in each voxel. The main reason of extracting the eigenvalues and eigenvectors is the possibility to visualize such information through the use of ellipsoids. Ellipsoids are defined with three axes, the eigenvectors, and the magnitude defined along each axis determines the shape of the ellipsoid through the eigenvalues, see Fig. 6.29.

Specific combinations of eigenvalues can be defined to separate different types of structures. Linear structures (also known as prolate or cigar-shaped structure), can be extracted when $\lambda_1 \gg \lambda_2 \simeq \lambda_3$. λ_1 is the dominant component, while λ_2 and λ_3 are similar but much smaller in magnitude compared to λ_1 , see Fig. 6.29(A). Linear structures are of type rank 2 tensor.

Figure 6.29: *Ellipsoids defining: (A) the linear case with $\lambda_1 \gg \lambda_2 \simeq \lambda_3$, (B) the planar case where $\lambda_1 \simeq \lambda_2 \gg \lambda_3$, and (C) the isotropic case where $\lambda_1 \simeq \lambda_2 \simeq \lambda_3$. Source: [15].*

Planar structures (also known as oblate or disc-like) can be extracted when $\lambda_1 \simeq \lambda_2 \gg \lambda_3$. In this case, λ_1 and λ_2 are similar and much larger than λ_3 . Therefore the two dominant components λ_1 and λ_2 form a plane with a slight protrusion along λ_3 , see Fig. 6.29(B). Planar structures are of type rank 1 tensor.

Isotropic structures (also known as spherical) can be extracted when $\lambda_1 \simeq \lambda_2 \simeq \lambda_3$. In that case, because all eigenvalues are equivalent, the ellipsoid becomes a sphere without any specific orientation and therefore called isotropic, see Fig. 6.29(C). Isotropic structures are of type rank 3 tensor.

6.2.4 Extraction of planar, tubular, and isotropic structures

The eigenvalues from the structure tensor can be used to estimate the probability of each visited neighborhood belonging to either a rank 1, a rank 2, or a rank 3 tensor, defined by p_1 , p_2 , and p_3 , where $\sum_{k=1}^3 p_k = 1$ which can be seen as probabilities [26]. The probabilities are calculated as

$$p_1 = \frac{\lambda_1 - \lambda_2}{\lambda_1}, \quad p_2 = \frac{\lambda_2 - \lambda_3}{\lambda_1}, \quad p_3 = \frac{\lambda_3}{\lambda_1}. \quad (6.17)$$

The micro-channels, as any other natural structure, are however composed of a mixture of these three cases and a more general ideal structure tensor can instead be represented by

$$\mathbf{T} = p_1 \mathbf{T}_1 + p_2 \mathbf{T}_2 + p_3 \mathbf{T}_3 \quad (6.18)$$

Examples of rank 1, rank 2, and rank 3 tensors are given in Fig. 6.30 for several elliptical cylinders, having different sizes and for different orientations. The rank 1 tensor is displayed as blue (planar), the rank 2 tensor is mapped with the red colour (linear or tubular), while the rank 3 tensor is

rendered in green (isotropic). To better locate the structure tensors locations in relation to the original data, the original data was rendered along with each presented case.

In Fig. 6.30(A), the original data is rendered semi-transparent. In Fig. 6.30(B), the rank 2 tensor is rendered alone and overlaid on the original data. In Fig. 6.30(C), the rank 1 tensor is rendered alone and overlaid on the original data. In Fig. 6.30(D), the rank 3 tensor is rendered alone and overlaid on the original data. In Fig. 6.30(E), the rank 1 and 2 tensors are rendered together. In Fig. 6.30(F), the rank 1 and 3 tensors are rendered together and overlaid on the original data. In Fig. 6.30(G), the rank 2 and rank 3 tensors are rendered together and overlaid over the original data. Fig. 6.30(H) shows the combination of all cases together, illustrating their complementarity. For the rank 2 tensor, the tubular structures are

Figure 6.30: *Extraction of rank 1, rank 2, and rank 3 tensors on elliptical cylinders of different sizes placed with different orientations. (A) original data, (B) rank 2 tensor alone, (C) rank 1 tensor alone, (D) rank 3 tensor alone, (E) rank 1 and 2 tensors rendered together, (F) rank 1 and 3 tensors rendered together, (G) rank 2 and 3 tensors rendered together, and (H) rank 1, rank 2 and rank 3 tensors rendered altogether.*

effectively representing the edges of the cylinders along with the very thin cylinder being represented by a rank 2 tensor on most of its path, except at its extremities where it shows rank 3 tensor information. Notice the appearance of the very thin cylinder displayed as semi-transparent in the rank 1 tensor case alone, while there should not be any information for this thin cylinder, see Fig. 6.30(C). In the sole presence of a rank 3 tensor, as shown in Fig. 6.30(D), or in the case of rank 1 tensor rendered together with rank 3 tensor, as represented in Fig. 6.30(F), the results overlaid on

the original data helps locating the isotropic information at the extremities of the very thin cylinder.

The respective p_1 , p_2 , and p_3 , along with their combinations p_1 & p_2 or p_1 & p_3 , or p_2 & p_3 are illustrated in Fig. 6.31. Although the separate cases - p_1 to the left, p_2 at the middle and p_3 to the right - alone on the top row clearly indicate their respective contribution in terms of energy, their relation with each others is difficult to visualize. Combining two cases at a time gives a better understanding of how the structures are related. The lower right image indicates where the tubular structures are located in relation to the neighbouring planar structures. The middle low image shows the isotropic part within the bone mostly prominent in the thick part of the bone, where the density changes very little in a small neighbourhood. The lower right image gives an indication of how the tubular structures mix with the isotropic structures.

Figure 6.31: *The three cases p_1 , p_2 , and p_3 and their combination p_1 & p_2 or p_1 & p_3 , or p_2 & p_3 . p_1 represents the planar part structures, p_2 represents the linear/tubular structures, p_3 represents the isotropic contribution.*

The complete combination, i.e. p_1 , p_2 , and p_3 altogether, is illustrated in Fig. 6.32. The original data, to the left, is represented simultaneously, to allow a better understanding of the results from the structure tensor in terms of p_1 , p_2 , and p_3 to the right.

Figure 6.32: *Combination of all cases p_1 , p_2 , and p_3 (right) compared to the original micro-CT data (left).*

6.3 Enhancement through adaptive filtering

When carefully observing the results in Paper III, the very thin micro-channels are not fully represented along their path. The noise further corrupts their definition. A natural step is, therefore, to reduce the noise level as much as possible, while at the same time enhance the weak micro-channels by favouring the linear/tubular structure information from the local structure tensor analysis. This process is known as anisotropic adaptive filtering and has previously been used in 2D, 3D, and even 4D [21, 29, 55, 76]. Similar to [76], the enhanced data f_a is an adaptive filtering defined as

$$f_a = f_{lp} + \sum_{k=1}^6 w_k f_{bp,k}, \quad (6.19)$$

where f_{lp} is a lowpass filtered version of the volume, w_k being a set of scalar values, and $f_{bp,k}$ represent the filter response from different bandpass filters with the same directions as the quadrature filters. The structure tensor

given in Eq. 6.18 can be estimated as

$$\mathbf{T}_{est} = \sum_{k=1}^6 \|q_k\| (\alpha \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T - \beta \mathbf{I}). \quad (6.20)$$

To express the adaptive filter in terms of the local structure tensor, the scalar values w_k can be estimated by projection of the local structure tensor onto a set of dual tensors M_k corresponding to $(\alpha \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T - \beta \mathbf{I})$ in the right hand side of Eq.6.20, such that the enhanced volume is defined as

$$f_a = f_{lp} + \sum_{k=1}^6 \langle \tilde{\mathbf{T}}, \mathbf{M}_k \rangle f_{bp,k}. \quad (6.21)$$

Enhancement can be performed using the local structure tensor, through tensor processing via the building of a control tensor $\mathbf{C}$. After smoothing the components of $\mathbf{T}_{est}$ using a small isotropic low-pass kernel, to prevent rapid changes in the tensor field, a soft threshold of the norm of the low-pass filtered tensor is computed

$$\tilde{\mathbf{T}} = \frac{\gamma(\|\mathbf{T}_{lp}\|)}{\|\mathbf{T}_{lp}\|} \mathbf{T}_{lp}, \quad (6.22)$$

where $\mathbf{T}_{lp} = h_{lp} * \mathbf{T}_{est}$ and γ is a smooth mapping corresponding to low-parametric sigmoid function defined for $\alpha \in [0, 1]$ as

$$\gamma(x, \sigma; \alpha, \beta) = \frac{x^\beta}{x^{\alpha+\beta} + \sigma^\beta}, \quad (6.23)$$

The γ -function controls the amount of anisotropic high-pass filtering, allowing a bandwidth selectivity. Well defined structures are therefore represented by a large norm while small fluctuations, such as noise, are represented by a small norm. The σ parameter can be understood as an estimate of the noise in the data. The α parameter is typically in the range of $[0, 0.5]$, and is used to accentuate weak structures that are right above the noise level. The parameter β sets the main slope of the γ -function in the noise-signal transition [26].

Fig. 6.33 gives an overview of the γ -function where only σ varies in a range of $[0, 1]$. When $\sigma = 0$, all frequencies are taken into account, noise included, while when $\sigma = 1$ only the low frequencies are preserved. The red curve illustrates the setting used when enhancing the micro-channels with $\sigma = 0.15$, $\alpha = 0.07$, and $\beta = 8$. Fig. 6.34 shows the effect of varying α in the γ -function. Low values of α will often result in a high-pass filter slowly tending to 1 as x tends to 1. Higher values of α will produce a peak right after the main slope, and will gradually converge back to 1 as x tends to 1. Fig. 6.35 gives an indication of how the slope of the γ -function behaves when only modifying the β parameter.

Figure 6.33: γ -function where only σ varies. A special case is shown with $\sigma = 0.15$, $\alpha = 0.07$, and $\beta = 8$ (red curve).

Figure 6.34: γ -function where only α varies. The red curve is a setting with $\sigma = 0.15$, $\alpha = 0.07$, and $\beta = 8$, allowing a large range of frequencies but with no amplification in the selected bandwidth.

Figure 6.35: γ -function where only β varies. The red curve is a special setting with $\sigma = 0.145$, $\alpha = 0.3$, and $\beta = 8$, accepting a fairly large range of frequencies but with a substantial amplification in the high frequencies.

The output of the γ -function corresponds to the largest eigenvalue, i.e. γ_1 , of the control tensor $\mathbf{C}$. Therefore, choosing the parameters governing the γ -function will influence the bandwidth selection of the high-pass filter in the enhancement scheme. All these parameters, σ , α , and β have an influence on each others, and their settings heavily depend on the application and how much noise there is in the data being processed. Noisy data will generally imply a low σ . If the goal is to only enhance very small weak structures in noisy data, α needs to be adjusted in consequent so as to only enhance weak structures, while preserving larger structures as they are. The fact that the structures are small, implies a high frequency representation as previously explained, and therefore α will have an impact on the noise level amplification, since it will also amplify the noise level.

Therefore, σ and α should be chosen carefully in relation to each other. The slope of the γ also needs to be considered. To control the shape of the adaptive filter, i.e. the orientational variation in the high-pass regions, the ratio of the eigenvalues of the local structure tensor can be used and remapped using the μ -function such as

$$\mu(x; \alpha, \beta) = \left[\frac{(x(1 - \alpha))^\beta}{(x(1 - \alpha))^\beta + (\alpha(1 - x))^\beta} \right], \quad (6.24)$$

where $x = \frac{\lambda_n}{\lambda_{n-1}}$. The parameter α is defined in the μ -function as the value of $\frac{\lambda_n}{\lambda_{n-1}}$ for which the μ -function is equal to $\frac{1}{2}$. A representation of α for the μ -function is given in Fig. 6.36, where α is varying from $[0, 1]$ and β is kept constant and equal to 8.

Figure 6.36: μ -function where only α varies, while β is kept constant. The red curve is an example of a special setting with $\alpha = 0.3$ and $\beta = 8$.

As for the β parameter, it defines the slope in the transition area in the same manner as for β in the γ -function, see Eq. 6.23. A representation of β for the μ -function is given in Fig. 6.37 where β is varying from $[1, 10]$ while α kept constant and equal to 0.3.

Figure 6.37: μ -function where only β varies, while α is kept constant. The red curve is an example of a special setting with $\alpha = 0.3$ and $\beta = 8$.

The resulting μ_n will correspond to the filter shape coefficients. The control tensor can now be constructed based on its largest eigenvalue γ_1 and the filter shape coefficients. Reusing the eigensystem from the original local structure tensor ensures a preservation of the orientation estimation. Therefore, the control tensor can be written as

$$\mathbf{C} = \gamma_1[\hat{\mathbf{e}}_1\hat{\mathbf{e}}_1^T + \mu_2[\hat{\mathbf{e}}_2\hat{\mathbf{e}}_2^T + \cdots + \mu_N[\hat{\mathbf{e}}_N\hat{\mathbf{e}}_N^T + \cdots]], \quad (6.25)$$

where $\hat{\mathbf{e}}_n^T$ are the eigenvectors of $\mathbf{T}_{lp}$. Eq. 6.25 can then be rewritten as a weighted sum of outer products of $\mathbf{T}_{lp}$ as

$$\mathbf{C} = \sum_{n=1}^3 \gamma_k \hat{\mathbf{e}}_n \hat{\mathbf{e}}_n^T, \quad (6.26)$$

leading to a modification of Eq. 6.21, such that

$$f_a = f_{lp} + \sum_{k=1}^6 \langle \mathbf{C}, \mathbf{M}_k \rangle f_{bp,k}. \quad (6.27)$$

A high-pass amplification factor was further added to Eq. 6.27 to favour enhancement of solely linear or tubular structures by using the eigenvalues from the structure tensor according to

$$f_a = f_{lp} + a_{hp} \sum_{k=1}^6 \langle \mathbf{C}, \mathbf{M}_k \rangle f_{bp,k}, \quad (6.28)$$

where a_{hp} is defined as

$$a_{hp} = \xi \cdot \frac{\nu_1 \lambda_1 + \nu_2 \lambda_2 + \nu_3 \lambda_3}{\nu_4 \lambda_1 + \nu_5 \lambda_2 + \nu_6 \lambda_3}, \quad (6.29)$$

with ξ set to 1.7 in Paper IV along with $\nu = \begin{bmatrix} 3 & 1 & -1 & 2 & -1 & 1 \end{bmatrix}$. The enhanced data can be resumed as the sum of a low-pass filtered version of the original data, applied only once, with a linear combination of anisotropic high-pass filters using the corresponding directions of the quadrature filter, and where the scalars w_k are mapped by the control tensor.

$$I_{enh} = I_{lp} + \sum_{k=1}^6 w_k I_{bp} \quad (6.30)$$

Examples of image enhancement using adaptive filtering are given in Figs. 6.38, 6.39, 6.40, and 6.41.

Figure 6.38: *Volume rendering of the enhanced data (right) versus the original data (left) of a partial bone specimen viewed from the medial side, with the antero-superior part of the tympanum at the top left and the antrum seen as the dark region central to the sample.*

In Fig. 6.38, a volume rendering of the enhanced data, seen partially, is placed side by side with the original data, to be able to see the achieved enhancement. Fig. 6.39 gives a zoomed representation of Fig. 6.38, principally at the lower right region. While it is difficult to appreciate the enhancement achieved on Fig. 6.38, due to the sole enhancement of fine structures corresponding to about 4 to 8 pixels, Fig. 6.39 allows a better distinction of weak structures from the original data when compared to its enhanced version.

Interestingly, the enhancement further accentuates the notion of depth and therefore 3D perception. In fact, the difference in depth perception relates, in this study, to a preservation of the transfer function based on the original data when observing the enhanced data. To optimize the comparison between the enhanced and the original data, the transfer function should be modified accordingly. Attempts to match the transfer function for both the original and the enhanced data were made, but relates more to a hit-and-miss operation, and therefore a similar transfer function was used for both datasets.

Figure 6.39: Zoomed version of Fig. 6.38 to better appreciate the details of the enhancement (right) compared to the original data (left).

Figure 6.40: Comparison of the enhanced (right) versus the original data (left), where the micro-channels are revealed while using a transfer function similar to Fig. 6.44 (D).

To fully appreciate the benefit of enhancing the micro-channels, for which the quadrature filters are optimized, Fig. 6.40 shows a comparison between the original data and the enhanced data. The micro-channels, visible to the right side of both the original and enhanced data, are more easily readable in terms of structure in the enhanced data than in the original data. In some locations, in the original data, the micro-channels appear disrupted by noise. In the enhanced data, the combination of lowpass filtering the noise in combination with amplifying weak structures clearly improve their connectivity.

Fig. 6.40 also reveals an important feature. Between the micro-channels, the noise appears more dense in some part than in others. If too dense, the noise will not be completely removed but will result in more blurry regions between the micro-channels. Also, notice in Fig. 6.40 some dark spots on the wall of some mastoid air cells. On the original data, such dark spots could in some extent be understood as local presence of noise, liquid, or remnant of some loose soft tissues. However, in the enhanced data, these

dark spots appear sharper and could indicate cavities with the walls of the mastoid air cells. After visual assessment, these cavities correspond to micro-channels opening into the mastoid air cells.

Similarly, in Figs. 6.41 and 6.42, connectivity of the micro-channels for the enhanced data (right) seems more accurate than in the original. In Fig. 6.41(right), the micro-channels, especially the ones in direct communication with the mastoid air cells, appear more complete along their path when compared to the original data, as visible in Fig. 6.41(left).

Figure 6.41: *Volume rendering of a clipped volume as in Fig. 6.40 with both the enhanced data (right) and the original data (left). Using a similar transfer function, the micro-channels appear more consistent.*

Figure 6.42: *Volume rendering of the same clipped volume as in Fig. 6.40 but a different orientation for both the enhanced data and the original data. Again, using a similar transfer function, the micro-channels appear more consistent. Noise reduction can be observed on the lower right side of each subfigure.*

As expected, in Fig. 6.42(right), the enhanced data provides more information than the original data (left) but may lead in a confusion while visually

assessing the results, thus requiring more attention when following the path of a micro-channel.

Micro-channels, present in the deeper part but hidden in the original data now become more visible in the enhanced data. Adjusting the control points of the transfer function may help in reducing this confusion, though both transfer functions should not differ too much for a fair comparison. To study the effect of the enhancement when assessing the p_1 , p_2 , and p_3 cases, Fig. 6.43 shows the p_1 , p_2 , and p_3 cases with the original data (left) and with the enhanced data (right) with a fairly high level of transparency, to allow a good colocation of p_1 , p_2 , and p_3 cases in relation to the underlying data.

When visually comparing both results, a higher contrast can be observed in terms of p_1 , p_2 , and p_3 for the enhanced data, when compared to the case with the original data. The reduction of the noise also greatly helps understanding their network, especially in relation to the depth.

Figure 6.43: *Combination of the enhanced data (right), estimation of p_1 , p_2 , and p_3 cases altogether, with the original data set (left) with a high transparency level. Colour legend: blue means p_1 or planar structures, red means p_2 or linear/tubular structures, and green means p_3 for isotropic structures*

To visualize the results from an analysis, such as a structure tensor analysis along with the original data, volume rendering is very helpful. The theory behind volume rendering is by far beyond the topic of this thesis, the interested reader is referred to [20]. Some features that were extensively used during this work are explained in the following section.

6.4 Volume rendering

Compared to visualization of the data in 2D through slicing back and forth, volume rendering allows rotating the loaded data in all possible directions in an interactive manner. Hiding unwanted structures, without removing the

information from the original data is also possible, allowing for instance to only visualize bone. Hiding data can be done in two manners, and possibly using a combination of the two methods when necessary. The first method is through the use of a transfer function, and the second method relates to the use of clipping planes, detailed in the following.

Transfer function

A transfer function is a technique to assign a colour (RGB: red green and blue) or a gray shade for each voxel in the 3D volume. To enable the hiding of voxels that are not of interest, a parameter is used to set the opacity (opposite to transparency) for each voxel, commonly called α . Together, the combination of the color and the alpha value is commonly known as a RGBA value per voxel. A one-dimensional transfer function maps the RGBA value for every grayscale intensity from the original data volume; a typical range is $[0, 255]$. Multi-dimensional transfer functions allow multiple RGBA values to be mapped to a single isovalue, but this topic is left to the interested reader [20].

For ease of manipulation, the transfer function is sometimes overlaid on a histogram of the visualized data. Interactive manipulation of the transfer function allows setting some control points on a curve, typically a piecewise linear function where the breaking points are defined by the control points. By setting a specific colour for these control points, it is possible to create linear colour gradients between two colours. The vertical placement of the breakpoint defines the opacity level, with a magnitude of 0 reflecting that specific grayscale value to be set to transparent, while a breakpoint positioning at a value of 1 enforces this grayscale value to be fully opaque.

The default transfer function, usually found in volume rendering, is a linear ramp ranging from 0 for the lowest grayscale intensity setting all the black pixels as transparent, and linearly increasing up to 1 for the highest grayscale intensity. For 2D medical images, this type of transfer function is not very useful, but it becomes practical in 3D. Hiding all black voxels from the data in 3D allows to visualize the objects of interest, while discarding the air and noise around the objects of interest. If not hidden, the visualized data will simply be a black cube. Different examples of volume rendering are given in Fig. 6.44. For each example, the transfer function is available below the rendering.

In Fig. 6.44(A), the default ramp is used all the way from the lowest grayscale value to the highest grayscale value. In Fig. 6.44(B), colours are assigned to some of the control points. Blue represents air, brown represents soft tissues, with two different types of brown to give a more realistic feeling, while bone is defined as different grayscale values to respect the traditional way of visualizing bone in volume rendering. One can observe the air within

Figure 6.44: *Volume rendering of a partial bone scan with different transfer functions.*

the mastoid air cells, while the brown in some of the cells clearly indicates that these cells pertain to trabecular spaces containing bone marrow. In Fig. 6.44(C), the blue colouring is removed and all air voxels are hidden, by setting the opacity to 0. In Fig. 6.44(D), the transfer function is modified so that only transition between bone and all remaining structures is visible. This makes it to provide information about the internal content of the bone. The right side of the transfer function in Figs. 6.44(A), 6.44(B), and 6.44(C) contains a linearly decreasing ramp down to 0. These respective ramps are intended to remove possible possible artefacts that contain high density grayscale values, when present. Another possibility to hide, but not physically remove, unwanted information is to use clipping planes.

Clipping planes

Clipping planes can be placed around the 3D volume along the six different locations: anterior, posterior, superior, inferior, left, and right. This option allows to interactively hide all voxels in front of these clipping planes, while keeping the remaining voxels behind the clipping planes visible, see Fig. 6.45.

Figure 6.45: *Clipping planes used to hide part of the data. A) The original data is fully rendered. B) The initial clipping planes are activated around the full data. C) The clipping planes are interactively displaced and rotated to only expose the part of interest from the original data. D) Once satisfied, the clipping planes are hidden and only the remaining region of interest is visualized. The white and yellow stars, defining mastoid air cells, are used as references in relation to subfigure A. Each subfigure is shown with a different zoom level.*

6.4.1 Multiple volume renderings

Rendering a sole grayscale volume is somewhat straightforward. However, rendering several volumes at the same time requires some strategical thinking on how to display these volumes, i.e. which order, which colour, which opacity, etc. In Paper III, the three volumes corresponding to p_1 , p_2 , and p_3 were combined into a single volume called p_m . To visualize p_m , the configuration illustrated in Fig. 6.46 was used. A so-called LUT (lookup table) combiner, allows a MeVisLab (a free software from Fraunhofer Institute for Medical Image Computing MEVIS, Bremen, Germany) user, to build a custom LUT where the colour channels, i.e. the red-channel, green-channel and the blue-channel, along with the α -channel, can be blended using their own respective primitive. The type of primitives used during this study was a sigmoid function, with a center and a width as illustrated in Fig. 6.46.

Notice, though, the green channel is a combination of two sigmoid functions, one with a base value of 0 and a peak value of 1, while the second has a base value of 1 and a peak value of 0. The red-channel was also built

Figure 6.46: Schematic representation of the volume rendering scheme used in Paper 3 to illustrate p_m .

with a sigmoid function with a base value of 1 and a peak value of 0. These primitives can be thought as a lowpass filter for the red-channel, a band-pass filter for the green-channel, and a high-pass filter for the blue-channel. Contrary to the RGB channels, the alpha-channel could only be built using a piecewise linear function, where control points could be adjusted accordingly.

The grayscale values from the resulting volume p_m , ranging from 0 to 1 could then be assigned a colour with red for tubular structures, green for isotropic structures, and blue for the planar structures and intermediate colours between the three cases. Note that p_m was multiplied with a thresholded version of the original micro-CT data, where only the bone was segmented so as to only visualize p_m within the bone and not for the overall data. A downside while using this visualization scheme is the manual assignment of the three colour channels. Selecting the wrong center and the wrong width of the sigmoid functions will lead to a misinterpretation of the structures from p_m . A positive aspect of this method, however, is its relatively fast rendering since only two volumes are loaded.

Another type of visualization consists of letting p_1 , p_2 , and p_3 , govern the colouring. Opting for this solution implies converting each of the cases p_1 , p_2 , and p_3 into a RGB volume, and blend their respective contribution in a LUT combiner. However, similarly to the first scheme, Fig. 6.46, the α -channel could only be built using a piecewise linear function with control points. The advantage of using this second scheme is the direct use of p_1 , p_2 , and p_3 without having to create p_m . A solution was to save the final RGB volume into a file and to later visualize it using a very basic

volume rendering scheme, see Fig. 6.47. Another advantage is the possibility to create a combination of the cases, such as p_1 & p_2 or p_1 & p_3 , or p_2 & p_3 . This second advantage allows to qualitatively assess the correlation between the three cases. The inconvenient part of this configuration is the loading of five volumes, corresponding to the three cases and the two grayscale volume, i.e. the volume representing the λ_1 and the original volume for the masking. The time used for the computer to assign each voxel the contribution from p_1 , p_2 , and p_3 was also a negative point. Due to a more accurate representation, the second scheme was preferred in relation to the first scheme, previously used in Paper 3.

Figure 6.47: A different schematic representation of the volume rendering that can be used to illustrate p_m , by simply loading the three volumes p_1 , p_2 , and p_3 , rescale them from 0 to 1 and convert them to a LUT mapping before being combined into a final LUT.

It is sometimes also desired to for instance show the complete result, while rendering a limited part from the original data the computations are based on. Two possible solutions are proposed. A first solution is illustrated in Fig. 6.48, where the result from the analysis is rendered twice while the original data is only rendered once.

A first set combines the original data together with a volume resulting from the analysis. This set is cached under the box named *caching of state during rendering*. The second set only involves the resulting volume. By

Figure 6.48: *Volume rendering scheme where the original data can be clipped while the secondary data remains fully rendered. To force the visibility of the secondary volume, it is present twice: under the caching box and the second one directly attached to the renderer.*

clipping the first set, the volume containing the results from the analysis is kept untouched from the second set. Yet another solution is to preserve the volume containing the results from the analysis untouched but clipping copies of the original data, as illustrated in Fig. 6.49. In addition to the regular clipping, the original data can also be masked with a binary volume containing only 1's, on which simple geometrical objects are applied by setting their inside as 0's. The resulting binary volume is then multiplied with the original volume and rendered on its own, while keeping the volume with the results from the analysis untouched.

Illustrations of what can be obtained using such clipping is provided in Fig. 6.50. In Fig.6.50(A and B), the clipping scheme presented in Fig. 6.48 is demonstrated with two different types of geometry, see the box cavity in Fig. 6.50(A) and the spherical cavity in Fig. 6.50(B). In Fig. 6.50(C), the original data is duplicated three times, one for displaying the result, another one for the front part, and one for the posterior part, all in all forming an inverting T shape. In Fig. 6.50(D), a mixture of a spherical cropping of the binary masking together with the set of clipping boxes from Fig. 6.49 were used.

In the next chapter, a 3D shape analysis on the local thickness of the mastoid air cells is further detailed. In this coming chapter, volume ren-

Figure 6.49: *Opposite to Fig. 6.47, the original volume is rendered twice while the secondary volume only once. A binary mask is used along the second rendering of the original data to crop simple shapes as spheres or boxes, allowing to display the secondary volume in those cavities. The three dots to the left indicates that several volumes can be added along.*

dering together with clipping and masking has been very useful, and form the ground for a qualitative analysis presented in Paper VII along with a descriptive quantitative analysis.

Figure 6.50: Results from applying specialized clipping planes and clipping volume to reveal the inside of the bone, to show how the blue spheres fit the mastoid air cells locally.

3D Shape Analysis

7.1 Introduction

This chapter provides a framework to answer the third question stated in the introduction of Chapter 4, namely what is the size of an air cell. Although the mastoid air cell system has been investigated as a very complex airspace, almost no reference has been found so far on determining the size of a singular air cell, and its possible variation in size. To the author's knowledge, only the work by Virapongse et al. [86] has been investigating this problem directly. From observations on several mastoid air cells using X-ray micro-CT scans, the size and shape of the mastoid air cells indeed vary dramatically.

Investigating the size of single air cells is not trivial. In [86], the authors only concentrated on so-called medium-sized cells corresponding to cells having a size, i.e. a diameter, ranging between 2 mm to 5 mm . Because of the limited resolution of their clinical x-ray CT scan, even after careful 3D reconstruction, i.e. $0.3\text{ mm} \times 0.3\text{ mm} \times 1.5\text{ mm}$, they could only perform a subjective analysis. All cells smaller than 2 mm in size would be too blurry to be correctly detected. In fact, they considered areas with small air cells as regions with incomplete pneumatization, and warned the reader that those regions should not be misunderstood as presence of a disease. The fact that the air cells were too small to be detectable on a clinical CT scanner is not a surprise, especially back in 1985. Presence of artefacts in x-ray CT scans, as previously introduced, leads to a false conclusion of the small air cells being hampered by mucosa thickening. Mucosa thickening can indeed happen, but it is not to be confused with physical limitations from the scanner, such as resolution or artefacts.

Despite some misleading parts, the work by Virapongse et al. [86] has pointed out an important link between the size of an air cell, together with the thickening of an inflamed mucosa, and the ability for the mastoid to

regulate negative pressure during pathology. This chapter aims at investigating the range of air cell sizes, and to observe how the cells are naturally distributed in the mastoid.

The structure of this chapter is now described. Investigating the local size of the mastoid air cell system can be seen as a shape analysis problem, therefore a brief introduction to shape analysis is provided. As will be explained further down, a series of maximally inscribed spheres are used as a shape descriptor of the local size. The maximally inscribed spheres are extracted along a medial axis, derived from a medial surface. The medial surface is itself extracted via the use of structure tensor analysis applied to a distance transform computed on the original data. Clustering of the extracted spheres, in the following referred to as medial balls, is then performed both in terms of size and location. While the size-based clustering enables sorting the medial balls in terms of their diameters, location-based clustering detects possible groupings of medial balls independent of their diameters. A further analysis is carried out by finding the proportion of each size-based cluster within each location-based cluster.

7.2 General definition of a shape

As defined by Kendall [41], a shape can be resumed by all geometrical information that remains when location, scale, and rotational effects are filtered out from an object. There has been a gradual shift going from shape description to shape representation.

Three-dimensional shape analysis deals with different types of challenges, ranging from the analysis of single complex shapes up to the analysis of an entire collection of shapes. Several main research directions, with specialized mathematical tools have emerged, see the review in [8]:

- **Feature detection:** *finding significant geometrical features in a robust manner.* Differential geometry such as curvature or detecting salient points on the shape surface are common methods. The smoothness of the walls forming the mastoid air cells can be represented using the curvature.
- **Segmentation:** *decomposing a 3D model into significant parts.* The mastoid air cells, or other parts of the temporal bone, can be decomposed into different objects.
- **Semantic labeling:** *assigning meaningful textual information to the constituting parts of an object.* The air cells can be labelled both in terms of location and in relation to neighbouring structures, as defined by A. Allam and M. Tos [3, 85].

- **Registration:** *finding correspondences between several objects.* Compare the outer shape of the mastoid air cell system from several temporal bone specimens.
- **Description:** *describing what an object looks like.* Machine-understandable signatures enable shape matching, shape retrieval, and shape classification using histogram or distance from one point to another. The distance from the antrum to the terminal air cells in the mastoid could be a description vector.
- **Matching:** *comparing objects through similarity measures.* From the shape descriptors, comparison can be done between several mastoid air cell systems (by comparing several features).
- **Retrieval:** *finding an object in a large database.* This is especially useful in large repositories of 3D models. When performing a very large shape analysis of the mastoid air cell system using many specimens, it can be useful to retrieve the object from one specific case in a large repository.
- **Classification and clustering:** *assigning an object to a specific class.* As will be explained later in this chapter, medial balls representing the local size of the mastoid air cells are clustered according to their size and assigned a specific size class.

Although the entire set of items in the list is very appealing to investigate, the shape analysis in this study only concerns two items: **description** and **classification and clustering**.

Zhang et al. [92] detail a classification of the different tools used in shape representation and description. Because only one tool is used in this study for the shape description part, only the flowchart representing most of the common tool is illustrated in Fig. 7.1. From a description and representation point of view, two main classes exist: *contour-based* and *region-based*. The difference between the two classes lies in the way the descriptors are extracted. In the contour-based method, the shape features are extracted from the contour of the object. In the case of the mastoid, the contour-based category would extract information from the contour of the mastoid air cells, i.e. the walls of the mastoid air cells. In the region-based category, the shape features are extracted from the overall volume forming the shape, i.e. all voxels forming the air cells including the contour voxels. When the shape needs to be represented in a discrete form, structural methods can be used. When the shape has a smooth representation (implying continuity), global methods can be used instead. A further sub-classification can be achieved by categorizing the method as either belonging to the space domain or to the transformation domain depending on how the shape descriptors are obtained, i.e. from the volume itself or from a transformation

such as the Euclidean distance transform [92].

Figure 7.1: Diagram resuming different subdivisions of shape analysis and the category the medial axis approach belongs to. The flowchart is inspired from a review written by Zhang et al. [92].

Since the micro-CT scans used in our study are discrete volumes made of voxels, the methods to be used are structural. Global methods can also be used through interpolation, forming a more continuous representation of the data. In order to use maximum information from the data, the region-based approach is preferred over the contour-based, which is valid for both representation and description. Among the structural region-based methods, one method was found to give the best results, namely the medial axis representation.

The medial axis, originally created by Blum et al. [9] during the second half of the twentieth century, allows a compact representation of the shape, while preserving sufficient information in order to reconstruct the shape of the original object as described by Kendall [41]. Obtained from the medial surface, the medial axis can be used to find maximally inscribed spheres, locally representing the thickness of the mastoid air cells. Fig. 7.2 shows the overall algorithm used in the current shape analysis, from the original micro-CT scan up to the obtention of the maximally inscribed spheres, i.e. medial balls, and how they are clustered based on their diameters, i.e. sizes, and their locations within the mastoid air cell system.

Examples of how medial balls along a medial axis can represent the size of the mastoid air cell system locally are provided in Fig. 7.3, for three

Figure 7.2: Flowchart explaining the extraction of the mastoid air cell size descriptor from the original micro-CT scan. A structural scheme based on medial balls extracted from a medial axis and a medial surface is used. The shape descriptor is used in a classification and clustering representation. Note the semi-transparent box illustrating the extraction of the medial balls from a skeleton using thinning instead, as a possible option which was used in this work for comparison purposes.

different cases. The colours represent the different range of sizes. The three cases represent air cells having different shapes. The first case, Fig. 7.3(A), represents a finger-like elongated air cell similar to Fig. 5.4 but in vertical orientation. A second case, where two air cells communicate via a narrow conduit, is illustrated in Fig. 7.3(B). Another pair of air cells, with a weakly defined transition or conduit, is presented in Fig. 7.3(C). These three cases correspond to the three cases illustrated in Fig. 3.9.

When the mucosa is not inflamed, the three cases fully participate in the gas exchange. During inflammation, however, the mucosa lining the mastoid air cell walls will expand as previously explained in Chapter 3. Swelling of the mucosa may eventually lead to a full closure of the finger-like air cell in case Fig. 7.3(A), close the lower air cell in case Fig. 7.3(B), possibly reducing the volume of gas of the airspace in case Fig. 7.3(C). For each case, i.e. A, B, and C, the medial axis is represented as illustrated in the top row. In the bottom row, maximally inscribed spheres are placed

along the medial axis, and colour-coded based on their sizes from dark blue for very small medial balls, up to purple for very large medial balls, with intermediate sizes respectively colour-coded as cyan, green, yellow, orange, red and magenta. In Fig. 7.3(A), the finger-link shape has a set of medial balls with a smooth variation in size. In Fig. 7.3(B), however, the size varies tremendously; from very small diameters up to very large diameters, with small spheres present between the two larger ones. Though the medial balls in Fig. 7.3(C) also vary in size, all the medial balls belong to a higher range of diameters.

Figure 7.3: Equivalent maximally inscribed spheres based on the Euclidean distance mapped onto the extracted medial axis.

From these examples, the diameters from all the medial balls provide a well-suited shape descriptor to represent the local thickness in a compact form. By clustering the medial balls in terms of size and location in 3D, it is possible to obtain a global map that may eventually be used in assessing the level of occlusion of the mastoid air cells, either locally or globally.

Before investigating the clustering part, two important tools needed in the extraction of the medial balls from the original data are first covered: the Euclidean distance transform and the skeletonization. Though the Euclidean distance transform is used in this study, extraction of a skeleton is merely done by comparing purposes between existing methods to the scheme proposed in this study. The Euclidean distance transform is first explained.

7.3 Euclidean distance transform

The original concept of the Euclidean distance transform originates from H. Blum in the 60's [9]. The transform is based on the grassfire principle, where the boundary of a grass landscape is set on fire. Under ideal conditions, the fire fronts advance only where there is grass, i.e. towards the centre of the landscape and shock scaffolds appear, see Fig. 7.4. These shock scaffolds appear where two fire fronts meet, and where the fire cannot go further.

Figure 7.4: *Original and final steps of grassfire theory (on the top row) according to Blum's theory which are given as analogy to a black rectangle surrounded by a white border (lower left) and its Euclidean distance transform (lower right) with the shock scaffolds displayed as dashed lines. The shock scaffolds are the skeletal elements of the black box. The black rectangle can represent a rough representation of an air cell on a micro-CT slice.*

For each pixel in the image, or respectively voxel in the volume, the distance transform assigns a number that corresponds to the distance between that pixel/voxel and the nearest nonzero pixel/voxel in the data. In mathematics, the Euclidean distance is defined as the straight-line distance between two points in the Euclidean space. The Euclidean distance transform is usually performed on binary data, though a grayscale version exists [35]. The grayscale representation is not used in this study.

For simple shapes like a parallelepiped, a sphere, or a cylinder, the Eu-

clidean distance transform is quite straightforward to mentally represent. Fig. 7.5 shows the Euclidean distance transform obtained on three simple shapes: Fig. 7.5(A) a sphere, Fig. 7.5(B) a parallelepiped, Fig. 7.5(C) cylinders of different heights, and for a more complex structure built as a combination of spheres with different diameters together with an ellipsoid, see Fig. 7.5(D). Each shape is partially cut, so as to show its internal representation. The colour representation used in this example maps low values close to the border as blue, while the largest values are represented as dark red, sometimes up to black. Notice how the how the central part of the three cylinders changes as they differ in height. However, when the investigated shape is more complex, its interpretation can become challenging.

Figure 7.5: *Euclidean distance transforms applied to different shapes: A) a sphere, B) a parallelepiped, C) cylinders of different heights, and D) complex set of shapes.*

To further illustrate the potential of the Euclidean distance transform and the challenges when working with complex structures such as the airspaces in the temporal bone, Fig. 7.6 illustrates the Euclidean distance transform of the airspaces in a temporal bone. The bone is partially represented but landmarks such as the lateral semi-circular canal, the utricle and a portion of the ear canal are represented. A partial representation of the mastoid air cells is also provided on the lower left side. As in Fig. 7.6, the colouring illustrates the distance from all air voxels to the bone voxels, the closer to

the bone voxel the more blue it becomes and the further away from the bone voxels the more red it becomes.

Figure 7.6: *Euclidean distance transform of the airspaces in the temporal bone. The lower right part represents part of the ear canal. The red spot to the left represents the central cell of the mastoid air cell system, and is known to be large anatomically. The bony part surrounded by the Euclidean distance transform represents part of the incus head and short process being part of the ossicular chain.*

The Euclidean distance transform, can be used to extract the local thickness of all airspaces. For instance, in Fig. 7.6, the red colouring present at the lower left of the image, representing the mastoid air cells, indicates that there is a set of large cavities. Indeed, these large cavities form the basis for the mastoid midline, with the antrum being prolonged in the central cells before splitting into smaller air cells. As will be shown later in this chapter, the Euclidean distance transform will be used to build a compact representation of the mastoid air cells used for the shape analysis. The notion of skeletonization, and how to extract the skeleton from the Euclidean distance transform, is explained in the following section.

7.4 Skeletonization

Four main classes of skeletonization methods exist: (1) thinning methods, (2) distance-field based methods, (3) general-field methods, and (4) geometric methods. There is a plethora of methods used in each of these skeletonization techniques. Their description is beyond the scope of this thesis, but interested readers can find several surveys, [67, 82]. Only the two first methods, i.e. the thinning method and the distance-field based

method, are considered in this thesis. The skeleton implemented in Paper VI is based on the Euclidean distance transform, and therefore belongs to the distance-field skeletonization based method. The thinning method used in Paper VI, for comparison purposes, is based on the thinning method. Before providing more details about the medial axis approach implemented in this study, the skeleton based on thinning-based method is first described.

The skeleton is directly extracted from a binary representation of the mastoid air cell system, obtained by thresholding explained in Sec. 6.1.1. The skeleton is obtained by an iterative morphological thinning procedure of the foreground of the obtained binary representation, where for each iteration a layer of boundary voxels is removed. Despite the need for multiple tests to respect certain conditions, this procedure is relatively straightforward.

Depending on the size of the binary structure to extract the skeleton from, this technique can take a large amount of time, in the order of an hour for a complete temporal bone from an X-ray micro-CT scan. A downside with such a skeletonization technique is the presence of loops which in 3D can also form blobs, i.e. an enclosed cavity. Such loops break one of the elementary rules of a true skeleton. Some methods propose to break these loops by cutting the shortest branch. Also, presence of noise in the original data can have a large effect on the extracted skeleton.

Any perturbation at the border of the binary representation will be propagated along the morphological thinning, leading to spur branches. Spur branches can be removed as a post-processing step, but when true fine structures are present at the borders, a decision on whether a branch is a spur branch to remove, or a branch to preserve, becomes challenging if not impossible [68]. A typical example of a skeleton obtained from the thinning method is illustrated in Fig. 7.7.

Beside the presence of noise, because this type of skeleton is discrete, it may eventually be locally represented by two voxels instead of only one voxel, breaking one of the fundamental rules for being characterized as a skeleton [72]. More practically, such a two-voxel representation may lead to further errors when attempting to extract maximally inscribed spheres locally, implying new conditions on which one of the two voxels to preserve. Similarly, the junctions of the skeleton are often represented by more than two voxels, often up to four or five voxels besides the junction voxel and its neighbouring voxels directly attached to it. For simpler shapes, these drawbacks can be corrected quite easily, but in the presence of very complex structures, such as the mastoid air cell system, correcting these issues would be very time consuming.

The second class of skeleton is based on the distance transform. As introduced in the previous section, shock scaffolds are formed where the fire

Figure 7.7: A 3D skeleton of the MACS obtained via binary thinning down to one voxel size branches. From the three windows, further zoomed to the right, multiple unuseful branches along with blobs are pointed out.

fronts meet. These intersections form a medial axis. As represented in Fig. 7.8, the shock scaffolds can have multiple shapes: a surface for the parallelepiped (Fig. 7.8(d)), a point for the sphere (Fig. 7.8(e)), and a line for the cylinder (Fig. 7.8(f)).

However, these medial axes are obtained by finding a global maximum for each of the cases. When looking at local maxima, the medial axes become more complex in shape. In Fig. 7.8(a), the surface is augmented by the presence of eight branches joining each corner of the surface to each of the corner. In Fig. 7.8(c), a conic shape is added to both extremities of the medial axis. Only the case with the sphere, in Fig. Fig. 7.8(e) does not change since only the place where the shock scaffold is present is at the centre of the sphere. To further help understanding the difference between Fig. 7.8(a,b,c) and Fig. 7.8(d,e,f), the EDT is also provided in colours. In Fig. 7.8(a), the Euclidean distance along each of the eight branches will vary from low value (blue) to the high values (red). The same applies in Fig. 7.8(c), where the base of two conic shapes will correspond to low values (blue) and where their respective tip will have the highest value (red). In the case of Fig. 7.8(d,e,f), a very compact representation of the original shape is obtained, but its reconstruction will be incomplete. However, in

Figure 7.8: *Natural medial axis formed by the shock scaffolds from three Euclidean distance on three simple cases: parallelepiped (a, d), a sphere (b, e), and a cylinder (c, f). In (a, b, c) the complete medial axis, obtained by detecting all local maxima, is represented with the black lines and further pointed out by the white arrows for each of the cases. In (d, e, f) only the global maximum was used to detect the medial axis. A cosine square of the EDT was used to represent the fire fronts at different levels, in order to visualize the shock scaffolds in each of the three cases. The fire fronts are represented with the gray surfaces seen as shells around the medial axis obtained from the global maxima. Clipping planes were used in order to visualize the inner parts of the three shapes.*

the case of Fig. 7.8(a,b,c), the reconstruction will be more accurate. It should be noted that the scaffolds are graphically represented using the cosine square of the Euclidean distance transform, i.e. $\cos^2(EDT)$, as gray surfaces mixed with the coloured Euclidean distance transform, as visible in Fig. 7.8. In the present work, reconstruction of the mastoid air cell system from the medial axis is not desired. However, because a good representation of the mastoid air cell system using medial balls is desired, detection of local maxima becomes necessary.

As will be seen later in this thesis, the medial axis is, in this study, extracted from the Euclidean distance transform using more advanced image processing, e.g. quadrature filters and structure tensor analysis, to allow a more robust detection of the local maxima. Also, when compared to the

morphological thinning, extracting a skeleton from a distance transform output allows a more continuous representation, and corresponds more to a curve skeleton than a discrete skeleton. In the case of a morphological thinning, an interpolation scheme is needed in order to produce a curve skeleton from its discrete representation.

In the next section extraction of the medial surface, used to extract the medial axis implemented in Paper VI, is further detailed.

7.5 Medial surface

From the enhanced data, the Euclidean distance transform is estimated, with the boundaries of the mastoid air cells used as starting points. The output of the Euclidean distance transform is then used to estimate the structure tensor. The resulting tensor is denoted $\mathbf{T}_{ms}$ with ms standing for medial surface. The quadrature filters used to estimate $\mathbf{T}_{ms}$ are the same as for estimating the structure tensor in the enhancement scheme. In addition, the local phase is estimated along with the norm of the structure tensor. To further help the reader's understanding, the eigenvalues obtained from $\mathbf{T}_{ms}$ are illustrated in decreasing order from left to right in Fig. 7.9. As can be observed, the largest eigenvalue λ_{ms1} already provides information about the medial surface. The local phase helps removing the responses from the boundary.

Figure 7.9: *Eigenvalues of the structure tensor estimated from the Euclidean distance transform, shown in decreasing order, for a mastoid air cell system: λ_{ms1} , λ_{ms2} , and λ_{ms3} .*

The local phase, briefly introduced in Sec. 6.2.1, is then used to separate the irregularities from the Euclidean distance transform from the data

boundaries. The local phase, rewritten in terms of the real and imaginary part of the quadrature filter response, is estimated as

$$\varphi_{loc} = \arctan\left(\frac{\sum_k \|\text{Im}\{q_k\}\|}{\sum_k \text{Re}\{q_k\}}\right), \quad (7.1)$$

with $\text{Re}\{q_k\}$ and $\text{Im}\{q_k\}$ being the real and imaginary parts of the quadrature filter responses in the k directions. To make use of the information provided by the local phase, a phase-based weighting factor w_φ is computed as

$$w_\varphi = \cos\left(\frac{\varphi_{loc}}{2}\right)^{\gamma_{loc}}, \quad (7.2)$$

with γ_{loc} equal to 4 in the present work. A simplistic illustration of the local phase, resulting from Eq. 7.2 is provided in Fig. 7.10. Signals with different widths, but with a similar magnitude, are illustrated in (a), while signals with different width and magnitude in (b). In all cases, the local phase detects the local maximum, as pointed out with the red circles. When using the Euclidean distance transform, the local phase will detect the local maxima resulting from the collision of the shock scaffolds.

Figure 7.10: Detection of the local maximum for signals with different widths (a) and also having different magnitudes (b).

The medial surface is then obtained by multiplying the norm of the structure tensor $\mathbf{T}_{ms}$ with Eq. 7.2, i.e.

$$M_{surf} = \|\mathbf{T}_{ms}\| w_\varphi. \quad (7.3)$$

An illustration of the medial surface before applying the local phase is provided in Fig. 7.11. The medial axis is represented in green, while the remaining part of the volume is represented in blue and considered as background. A zoom is added at the border of one air cell, to illustrate the need to remove unwanted ridges. To highlight the medial surface further, the transfer function used to represent the temporal bone corresponds to the transfer function introduced in Paper II.

Figure 7.11: Volume rendering of the medial surface (green) before removal of unwanted ridges via the local phase. The remaining part of $\|\mathbf{T}_{ms}\|$ is represented as blue. A zoom is also provided to illustrate the unwanted ridges from the boundaries. Source: [14].

Once the local phase is applied to $\|\mathbf{T}_{ms}\|$, the medial surface is extracted and illustrated together with the original data for two different angles in Fig. 7.12.

Figure 7.12: Volume rendering of the medial surface represented in yellow, along with the original data for two different orientations on the same bone specimen, (left) from a lateral view and (right) from a posterior view. Clipping planes were used to reveal the inner part of the bone.

The medial surface output volume is illustrated in Fig. 7.2, as the gray box with the number 1 next to it. Using the idea of extracting the medial

axis from the medial surface from [63], the medial axis is then obtained as explained in the following.

7.6 Medial axis

The medial axis can be directly extracted using

$$M_{axis} = p_{ms2} = \frac{\lambda_{ms2} - \lambda_{ms3}}{\lambda_{ms1}}, \quad (7.4)$$

where λ_{ms1} , λ_{ms2} , and λ_{ms3} are the eigenvalues of $\|\mathbf{T}_{ms}\|$. Further, it was decided to perform a structure tensor analysis on $\|\mathbf{T}_{ms}\|$, so as to gain more flexibility on the content of the medial axis. The estimated structure tensor is called $\mathbf{T}_{ma}$ standing for medial axis. The eigenvalues obtained from $\|\mathbf{T}_{ma}\|$ are then $\lambda_{ma1} \geq \lambda_{ma2} \geq \lambda_{ma3} \geq 0$. Using

$$p_{ma1} = \frac{\lambda_{ma1} - \lambda_{ma2}}{\lambda_{ma1}}, \quad p_{ma2} = \frac{\lambda_{ma2} - \lambda_{ma3}}{\lambda_{ma1}}, \quad p_{ma3} = \frac{\lambda_{ma3}}{\lambda_{ma1}}, \quad (7.5)$$

along with p_{ms1} and p_{ms3} , the medial axis can be extracted using

$$M_{axis} = (p_{ma1} - \alpha_{ma} p_{ms1}) + p_{ma2} + \beta_{ma} (p_{ms3} + p_{ma3}), \quad (7.6)$$

with $\alpha_{ma} = 0.75$, and $\beta_{ma} = 0.01$. The parameters α_{ma} and β_{ma} respectively allow the addition of a planar and an isotropic components to the medial axis. Although slightly violating the properties defining a skeleton or a medial axis as stated in [82], preserving a certain amount of planar and isotropic part from the medial surface, together with the tubular part of the medial axis, the medial axis has better continuity. Note that using a combination of the eigenvalues from both tensors, $\mathbf{T}_{ms}$ and $\mathbf{T}_{ma}$, can also be used to create a more sparse medial axis with less points along its course. Because the two tensors are computed over the entire volume, it is necessary to restrict the information along the medial surface. Therefore, the medial axis is finally calculated as

$$\tilde{M}_{axis} = M_{axis} M_{surf}. \quad (7.7)$$

An illustration of the obtained medial axis is provided in Fig. 7.13. To assess the quality of the skeleton, another skeletonization scheme (available over the Matlab file exchange (Mathworks, USA)) was used [50]. The package will in the following be referred as *Skeleton3D*. In Fig. 7.13(left), the medial axis extracted using the present scheme is presented, while Fig. 7.13(right) presents the skeleton resulting from *Skeleton3D*.

Figure 7.13: *Skeletonization using the structure tensor analysis from the EDT (left) compared with a skeleton based on the scheme proposed in [50] (right). Small air cells next to yellow stars were not detected in the Skeleton3D scheme (right), while they were detected using our scheme (left). The white arrow points to spur branches observed on the skeleton (right), while they are not observed in our scheme. The green star indicates the presence of an unwanted blob in the Skeleton3D scheme (right) at the superior part of the temporal bone, while it is not present in our approach (left).*

It should be noted that the skeleton provided in Fig. 7.13(right) at first sight looks less disrupted than the medial axis implemented in this study. However, many of the smaller air cells are not represented by the skeleton using the thinning scheme, see the yellow stars. Also, a blob appears at the top right of the air cell system to the left of the ossicular chain, see the green star in Fig. 7.13(right) as compared to Fig. 7.13(left). Another positive feature in the favour of the medial axis from the current study, is the absence of spur branches.

The medial axis volume is illustrated in Fig. 7.2, where p_{ms1} , p_{ms2} , p_{ms3} , p_{ma1} , p_{ma2} , and p_{ma3} are pointing to the gray box with the number 2 next to it. The medial balls are directly extracted from the medial axis.

7.7 Medial balls

From the medial axis, an implicit representation of the medial balls can be obtained. By reassigning the original Euclidean distance transform on the medial axis, each point of the medial axis then corresponds to the centre of

a maximally inscribed sphere, with a radius corresponding to the Euclidean distance transform value. These medial balls are obtained by combining the medial axis with the Euclidean distance obtained from the enhanced data, and are illustrated in gray in Fig. 7.2 with the number three next to it. Clustering can then be used to classify the medial balls according to their sizes or locations, clustering based on the size is first approached.

Figure 7.14: *Medials balls representing the maximally inscribed spheres within the mastoid air cells.*

The medial balls can then be used as a shape descriptor and in a clustering scheme, to find any specific patterns in the volume formed by the medial balls as represented in Fig. 7.14. Such a pattern could, in terms of the mastoid, be a local aggregate of air cells of a certain size range, or whether the air cells form natural groups as previously discussed in Chapter 3. To achieve a clustering analysis, two tools need to be employed. One tool is the K-means clustering, which is used to divide the medial balls into clusters. The second tool is the convex hull, which is related to the enclosed regions containing the clustered medial balls. These two tools are now described.

7.8 K-means clustering

Kmeans or K-means is a partitioned clustering algorithm that tries to find K non-overlapping clusters, each represented by their centroids. The K-means clustering process can briefly be described with the following steps. First, K centroids are initially manually selected, often based on a-priori knowledge about the data. In the case of the mastoid air cells, the number of regions

defined by A. Allam and M. Tos is a good starting point [3, 9]. Every point in the data is then assigned to the closest centroid. The collection of points assigned to a centroid defines a cluster. The centroid of each cluster is then updated based on the points assigned to that cluster. This process is repeated until the clusters are stable. In relation to the medial balls, the points represent their respective centres. More information about K-means clustering can be found in [90].

An important point to add, though, is how to select the number of clusters to divide the data into. While it is possible to select a predefined number of clusters using some a-priori knowledge about the data, a repeated iteration of the K-means algorithm with different cluster values does not always give consistent classification, visible when observe their convex hulls changing shape over iterations. If the results from manual selection of the number of clusters is not satisfactory, it is possible to use the Calinski-Harabasz criterion, also known as the variance ratio criterion, to get the optimal number of clusters [10]. The notion of a good clustering depends on what needs to be evaluated. Attaining high intra-cluster similarity and low inter-cluster similarity is the ultimate goal of clustering. It should be noted that a high value from this criterion does not necessarily lead to a natural clustering of, for instance, the air cells in the mastoid.

The notion of convex hull was introduced earlier. In the next section, a brief definition of the convex hull is provided.

7.9 Convex hull

To visualize the obtained clusters, a convex hull is determined for each cluster. A complete definition of a convex hull is beyond the scope of this thesis, but interested readers are referred to [7]. Its basic definition is resumed in the following.

A convex hull, in 2D, is defined as the smallest convex polygon P that contains all the points of a set S . A polygon P is said to be convex if: (1) P is non-intersecting; (2) for any two points p and q on the boundary of P , the segment pq lies entirely inside P . In 3D, given a set of n points P , the convex hull $CH(P)$ is the smallest polyhedron, so that all elements of P lie inside and at the boundary of the polyhedron. An example of a 3D convex hull is illustrated in Fig. 7.15.

K-means clustering is often used for N-dimensional data, requiring a definition of a convex hull in N-dimensions. Fortunately, as the data used in this work is three-dimensional, a 3D convex hull representation of the obtained classes can help to assess their variation in shape and size over the selected numbers of clusters, as explained in the next section.

Figure 7.15: A 3D convex hull (right) for a set of random 3D points (left).

7.10 Clustering the air cells

An example of clustering the temporal bone is illustrated in Fig. 7.16, where each cluster is represented by its convex hull. Each cluster is assigned a different colour for better distinction among the clusters. Each convex hull is rendered semi-transparent to reveal the centroid. It should be noted that the clustering in this example is the result of an earlier test where the number of clusters to be found was chosen to be 14. Within each cluster, the points belonging to the medial axis are used as the centres for the medial balls as illustrated in Fig. 7.15. When comparing to the regions defined by A. Allam and M. Tos [3, 85], the resulting clusters represented by their convex hulls, seem to correspond to the natural grouping of cells as observed in exposed temporal bones.

At least two types of clustering are possible: one clustering based on the size of the medial balls (defined by their diameters), and a second clustering based on the locations of the medial balls in the mastoid air cell system. The two proposed clustering approaches are now detailed.

7.10.1 Clustering in relation to size

The first clustering approach is aimed at defining a set of size ranges. Using the equivalent radii alone as a 1D vector, the K-means clustering can classify the vector into several classes based on a predetermined number of clusters. To avoid confusion with the location-based clusters, the eight size-based clusters are named from C_{S1} up to C_{S8} , with the letter S to emphasize that it relates to size-based clusters. The eight volumes, each corresponding to a separate volume, are represented by the gray boxes C_{S1} and C_{S8} with

Figure 7.16: *Example of a segmentation using K-means clustering. The temporal bone clusters are represented by semi-transparent convex hulls revealing the cluster centroids (red stars). Each cluster is represented with a different colour for better interpretation. Note that this example is based on 14 clusters with the exclusion of the tympanum.*

the number 4 on its left in Fig. 7.2. To obtain a detailed grouping of the medial balls, and after using the Calinski-Harabasz criterion, the number of clusters was chosen to be eight, where: blue represents the smallest medial balls, cyan represents a range of small medial balls, green corresponds to medium small medial balls, yellow is dedicated to medium medial ball sizes, orange defines the medium large medial balls, and where red magenta and purple emphasise the large up to the largest medial ball sizes.

In Fig. 7.17, the original data, Fig.7.17(a), and renderings of each the eight size clusters are provided; each rendered separately with a semi-transparent representation of the bone. While Fig.7.17(a) is partially clipped, the semi-transparent original data for each of the remaining subfigures is fully represented. As can be seen in Fig. 7.17, the K-means clustering was able to cluster the medial balls based on their sizes, i.e. diameters. By comparing the different size-based clusters, the smallest medial balls seem to be more concentrated towards the superior part of the bone, see Figs. 7.17(b), 7.17(c), 7.17(d). The clusters capturing the larger medial balls, i.e. C_{S6} to C_{S8} , show a clear separation between the central airspace in the proximity of the antrum, see Figs. 7.17(g), 7.17(h), 7.17(i). Note that this symmetry is often used as the midline of the mastoid, from the antral cells down to the apical cells at the lower part of the bone.

The remaining medial balls seem to be well spread over the whole mastoid, see Figs. 7.17(e) and 7.17(f), though a groove seems to already appear in Fig. 7.17(f), announcing a separation between the lower part from the superior part of the mastoid air cell system. When particularly observing

Figure 7.17: Volume rendering of the eight size-based clusters represented separately viewed from the lateral side of the mastoid: (a) original data, (b) C_{S1} represented in blue, (c) C_{S2} represented in cyan, (d) C_{S3} represented in green, (e) C_{S4} represented in yellow, (f) C_{S5} represented in orange, (g) C_{S6} represented in red, (h) C_{S7} represented in magenta, and (i) C_{S8} represented in purple. Source: [14].

the clusters C_{S5} and C_{S6} in Figs. 7.17(g) and 7.17(h), larger medial balls are also present at the mid-level of the mastoid. The eight clusters can be combined together to form a similar volume as the one illustrated in Fig.

7.14, with the exception that each medial ball is now coloured based on its size cluster as shown in Fig. 7.18. The combination of all clusters is represented in Fig. 7.18.

Figure 7.18: Volume rendering of different combinations of eight size-based clusters from the largest medial balls to the smallest one, to see how the different size clusters relate to each others. Legend: (a) original data, (b) C_{S8} , (c) $C_{S8}-C_{S7}$, (d) $C_{S8}-C_{S6}$, (e) $C_{S8}-C_{S5}$, (f) $C_{S8}-C_{S4}$, (g) $C_{S8}-C_{S3}$, (h) $C_{S8}-C_{S2}$, (i) $C_{S8}-C_{S1}$. Source: [14].

However, the relation between the cell sizes and their locations within the mastoid is not captured when clustering the medial balls in terms of size. In the next subsection, a clustering analysis of the medial balls based on their 3D location is presented.

7.10.2 Clustering in relation to 3D location

Compared to the size-based clustering, clustering the medial balls based on their 3D locations requires a more subtle selection of the number of clusters. One option when choosing their grouping in 3D relates to an a-priori knowledge on how they are naturally grouped in the temporal bone. As proposed by M. Tos [85], a set of eleven groups can represent the mastoid air cell system. Because the tympanum and the inner ear structure, i.e. the cochlea and the vestibular system, are also present in the original data, one more cluster was added leading to a total of 12 clusters to extract. As for the size-based clusters, the location-based clusters are named as C_{L1} up to C_{L12} , with L to emphasize that they relate to location-based clustering. C_{L1} up to C_{L12} are illustrated in Fig. 7.2 with the number 5 to the right side of it.

Fig. 7.19 and Fig. 7.20 provide the location-based clusters from a lateral side and from a posterior view. The relations between the clusters shown in Fig. 7.19 and Fig. 7.20 to the regions described by A. Allam and M. Tos [3, 85] (as illustrated in Fig. 3.4), are shown in Table 7.1. For each cluster, the corresponding regions are provided as names, but also according to the label numberings in Fig. 3.4.

	Description	Labels in Fig. 3.4
C_{L1}	tympanum	-
C_{L2}	antral, tegmental, zygomatic	1, 3, 4
C_{L3}	antral (medial), tegmental	1, 3
C_{L4}	antral (lateral), periantral	1, 2
C_{L5}	antral (medial), periantral	1, 2
C_{L6}	periantral, sinodural, superior retrosigmoid, medial persinusal	2, 5, 7, 8
C_{L7}	central (inferior), inferior retrosigmoid, perifacial (medial)	6, 9, 10
C_{L8}	central (inferior), inferior retrosigmoid, perifacial (lateral)	6, 9, 10
C_{L9}	perifacial (medial), apical (anterior)	10, 11
C_{L10}	perifacial (lateral), apical (anterior)	10, 11
C_{L11}	perifacial (lateroinferior), apical (anterior)	10, 11
C_{L12}	apical (posterior)	11

Table 7.1: *Anatomical correspondence of the 3D location-based clusters from Fig. 7.19 and Fig. 7.20 in relation to the regions defined by A. Allam and M. Tos ([3, 85]) as illustrated in Fig. 7.2. Source: [14].*

Size distribution histograms of each location-based cluster are provided in Fig. 7.21. A common observation is the fact that all distributions have a positive skewness, except for the last histogram, see Fig. 7.21(l). The skewness is to be expected since, as visible in Fig. 7.19(a, b, c) for the posterior view and Fig. 7.20(a, b, c) for the corresponding superior views, there are many more small medial balls compared to larger ones, thus leading to a shift of the mean towards the smaller cells. Also, as clearly visible in Fig. 7.21(g), the histograms often contain a mixture of distributions, which makes sense since several size-based clusters are combined into each of the location-based clusters.

The proportion of each size-based cluster in each location-based cluster is represented as a series of twelve pie charts, see Fig. 7.22. As previously explained, C_{L1} corresponds to the tympanum and stands on its own. However, as highlighted by the red rectangle, clusters C_{L2} up to C_{L4} have very similar patterns in terms of proportions. There is a dominating proportion of cells with a size range of type C_{S2} and C_{S3} . Similarly, C_{L5} up to C_{L7} present common features with more small medial balls than in the previous rectangle, but more medial balls of type C_{S3} and C_{S6} . C_{L8} has a well distributed amount of each size-based cluster with a lower amount of C_{S8} . No C_{S9} is present though. Cluster C_{L9} could be included in the blue rectangle but its lower amount of C_{S5} and C_{S7} implies C_{L9} to be on its own. The same applies for C_{L10} though it contains all types of size-based clusters. Clusters C_{L11} and C_{L12} are very similar in terms of proportions and were therefore paired together within the green rectangle. C_{L11} and C_{L12} correspond to the apical cells. A black arrow above C_{L1} is aimed at helping the reader to visualize the chronology of appearance from C_{S1} up to C_{L12} . The star added to the cluster title from C_{L3} up to C_{L9} indicates that no medial balls of size-type C_{S9} was found in those clusters.

A more clinical interpretation can be provided by these results. When the mucosa lining the mastoid air cell system is inflamed, it is expected that regions with a narrow space such as small air cells, narrow conduits, or small protrusions will be filled by mucosa, thus leading to a much lower surface area exposed to gas exchange as compared to the volume of gas within the air cells. Depending on the thickening of the mucosa, a first approach could in the future be to remove, for instance, all medial balls belonging to C_{S1} and C_{S2} from all location-based clusters C_{L1} up to C_{L12} , and observe the impact on the local pneumatization. For clusters C_{L2} up to C_{L4} , removing C_{S1} and C_{S2} would reduce their capacity by more than 50 %. For clusters C_{L5} up to C_{L7} , removing C_{S1} and C_{S2} would reduce their capacity by more than 45 %, while for C_{L8} C_{L11} and C_{L12} , the reduction would be slightly lower with about 40 %. In the remaining set of clusters, that is C_{L1} , C_{L9} , and C_{L10} , the reduction would be down by a factor of 4.

Using simulation, it may be possible to assess the effect of mucosal thickening on the gas exchange rate, and thereby on the negative pressure regulation.

In the next chapter, an overview of the published papers related to this thesis is provided.

Figure 7.19: Volume rendering of the twelve 3D location-based clusters together with a semi-transparent representation of the original data from a lateral view for better co-localisation. Source: [14].

Figure 7.20: Volume rendering of the twelve 3D location-based clusters together with a semi-transparent representation of the original data from a posterior view for better co-localisation.

Figure 7.21: Size distribution for each of the location-based cluster. Legend: x -axis represents the size, i.e. diameter, of the medial balls in mm; the y -axis represents the number of medial balls found per size bin.

Figure 7.22: *Pie charts representing the proportions of each size-based cluster within each of the location-based cluster. Source: [14].*

Summary of Papers

8.1 Introduction

This chapter provides brief summaries of the enclosed papers. Most of the materials presented in the papers have been covered in the first part of the thesis. From the previously stated observations, two main paths were taken. Due to the rare availability of X-ray micro-CT scans of complete temporal bones, especially including the mastoid bone, it was considered very important to investigate the surface area and volume occupied by the air in the mastoid air cell system, and compare the results obtained with previous studies (most of them relying on clinical CT scans).

Paper I relates to this part of the study. During the scanning of the temporal bones at the Department of Physics and Astronomy, Ghent (Belgium), while assessing the quality of the scans using volume rendering, together with Manuel Dierick and Elin Pawels, micro-channels running in the cortical part of the bone were discovered. As previously introduced, this discovery raised many questions. Therefore a divergence from the analysis of the mastoid air cells was unavoidable. Due to their complex structures, and the fact that their size is close to the scan resolution, further image processing was needed to better understand their geometrical inter-relation. Papers II, III, and IV relate to this part of the study.

Though further investigation of these micro-channels is relevant, the analysis of the mastoid air cells, not as a single system but as separate entities, was at this stage considered a natural extension of the primary work. In order to extract relevant metrics, such as the size distribution or the way the cells are grouped together in the mastoid, a preliminary shape analysis was designed to investigate the complete range of sizes the air cells have.

The analysis was applied on one complete temporal bone. In order to obtain a simpler representation of the air cell system, a new way of extracting a medial axis from an Euclidean distance transform was devised, and further

used to extract medial balls from which the local size was estimated. Using the medial balls, it was then possible to study the local variation in size of the mastoid air cells and assess whether the size variation correlates with the regions defined by A. Allam and M. Tos, [3, 85]. Paper V presents the novel way of extracting the skeleton from the Euclidean distance transform, while Paper VI provides improvements on obtaining the medial axis and details the clustering of the obtained medial balls in terms of size and location. These contributions, though not necessarily all relevant during clinical practice, are believed to facilitate the understanding of the mastoid in terms of anatomy, but also in terms of its physiological role. Each paper contribution is presented separately as follows.

8.2 Paper I

Determination of the mastoid surface area and volume based on micro-CT scanning of human temporal bones. Geometrical parameters depend on scanning resolutions. *Olivier Cros, Hans Knutsson, Mats Andersson, Elin Pawels, Magnus Borga, Michael Gaihede. Hearing Research Special Issue MEMRO 2016, Volume 340, pages 127-134, 2016.*

In this paper we describe the determination of the surface area to volume ratio, which may indicate that the mastoid air cell system is adapted to gas exchange, and assess the influence of the resolution on these parameters by downsampling the data four times. The surface area and volume were obtained through image processing tools commonly found in the field of hearing research, for comparison purposes though adjusted to the high resolution of the micro-CT scans. The current study found significantly larger surface areas and higher surface area to volume ratios than found in previous studies, performed at a limited resolution of around $625\mu\text{m}$ (commonly found in clinical CT scanners). The statistics resulting from this study are important for a more accurate modelling of the middle ear physiology.

8.3 Paper II

Micro-channels in the mastoid anatomy. Indications of a separate blood supply of the air cell system mucosa by micro-CT scanning *Olivier Cros, Magnus Borga, Elin Pawwels, Joris Dirckx and Michael Gaihede. Hearing Research. Volume 301. pages 60-65, 2013.*

This paper describes a discovery within the temporal bone named after its structural shape, i.e. the micro-channels. A qualitative analysis using

volume rendering first demonstrated the presence of these micro-channels, connecting the surface of the compact bone directly to the mastoid air cells, as well as forming a network of connections between the air cells. A quantitative analysis on 2D slices was then employed to determine the average diameter of these micro-channels, as well as their density at a localized area. These channels are believed to house a separate vascular supply for the mastoid mucosa. This study further helps improving our knowledge of the physiological properties of these micro-channels, which may have important implications for understanding the pressure regulation of the middle ear.

8.4 Paper III

Structural analysis of micro-channels in human temporal bone

Olivier Cros, Michael Gaihede, Mats Andersson, Hans Knutsson. Annual meeting of the IEEE International Symposium on Biomedical Imaging (ISBI), New York, United States of America, pages 9-12, April 2015.

In this paper a preliminary structural analysis of the micro-channels was implemented using structure tensor analysis. The eigenvalues obtained from the estimated local structure tensor were used to build probability maps representing planar, tubular, and isotropic tensor types. A visual interpretation of the three different tensor structure types was possible by assigning a respective RGB colour, and the full structure tensor was rendered along with the original data, to allow a smooth transition between the three tensor types. Such a structural analysis has provided new and relevant information about the micro-channels, but also their connections to the mastoid air cells.

8.5 Paper IV

Enhancement of micro-channels within the human mastoid bone based on local structure tensor analysis.

Olivier Cros, Anders Eklund, Michael Gaihede, Hans Knutsson. IEEE Sixth International Conference on Image Processing Theory, Tools and Applications (IPTA'16), Oulu, Finland, December 2016.

In this paper enhancement of the micro-channels was implemented using structure tensor analysis based on the framework in Paper III. In order to assess the quality of the results obtained with our scheme, a qualitative analysis was carried out followed by a comparison with another enhancement scheme. The enhanced micro-channels appeared more well-defined.

8.6 Paper V

Surface and curve skeleton from a structure tensor analysis applied on mastoid air cells in human temporal bones. *Olivier Cros, Michael Gaihede, Anders Eklund, Hans Knutsson. Annual meeting of the IEEE International Symposium on Biomedical Imaging (ISBI), Melbourne, Australia, April 2017.*

In this paper extraction of both a medial surface and a medial axis was achieved from an Euclidean distance transform itself derived from the norm of a structure tensor analysis from an enhanced version of the original data. The structure tensor analysis is very similar to the analysis for Paper IV, besides the addition of the local phase needed to isolate the medial surface from the Euclidean distance transform. Results were presented both on a test object, and on a limited region of one of the micro-CT scans available in this study.

8.7 Paper VI

Local size descriptor of mastoid air cells in human temporal bone based on a structure tensor analysis from an Euclidean distance transform. *Olivier Cros, Anders Eklund, Michael Gaihede, Hans Knutsson. Submitted to the journal Computerized Medical Imaging and Graphics, March 2017*

In this paper, a basic 3D shape analysis is implemented to investigate the local thickness of the mastoid air cell system via a clustering technique, where medial balls extracted from the medial axis (from Paper V) are classified based on their size and location. To allow more flexibility, the medial axis was revisited from Paper V, to include geometrical properties from the structure tensor analysis from both the medial surface and the medial axis itself. This preliminary 3D shape analysis can be used to foresee the impact of a mucosal thickening, due to inflammation, on the gas exchange.

Discussion & Conclusion

This chapter gives a discussion about the overall study. Some ideas for future work are also presented. First, the aims stated in Chapter 5 are critically assessed, both from a medical point of view and from a technical point of view interchangeably depending on the relevance of each part. Next, based on the work achieved, some clinical and technical methodological aspects are proposed as future work. Also, observations of three structures that have never been reported, especially not in 3D, are presented in a third section. Illustrations based on 3D volume renderings, together with personal drawings, suggest a potential reason of their presence but no claim about their actual physiological role is confirmed. These structures are thought to be mentioned before a general conclusion though.

9.1 Fulfillment of the aims or not?

To assess whether the aims of this PhD study have been fulfilled, each stated aim will be discussed and kept in its original main topic: (1) Geometrical aspects of mastoid air cell system, and (2) Interpretation and geometrical description of the discovered micro-channels.

9.1.1 Geometrical aspects of mastoid air cell system

Aim 1. Determination of the surface area and volume of the mastoid air cell system from X-ray micro-CT scans necessary for estimating the pneumatization of the mastoid (Paper I).

Measuring the surface area and volume from a discrete volume resulting from any scanning technique is always a difficult problem. The typical example is the measurement of the coastline of England, called the coastline paradox, stating that the coastline of a landmass does not have a well-defined length due to its fractal-like properties. The resulting measure-

ment depends on the method used to measure it. Similar to the coastline measurement paradox, the mastoid air cell system has features at nearly all scales, from centimetres to possibly nanometers and below. Thus there is no real obvious size of the smallest mastoid air cell that can be measured. Therefore, the perimeter is not well-defined and measurement of the surface area and volume of the complete mastoid air cell system will be impacted by this limitation. However, for the mastoid air cells to play a role in pressure regulation, the mucosa lining their wall may impose a minimal size on the smallest mastoid air cell. Also, previous measurements were solely done using X-ray clinical CT scans, for which the resolution is about sixteen times lower, compared to the X-ray micro-CT scans used in this overall study. It should be reminded that as the voxel resolution increases, the mixture of different tissue types becomes smaller, allowing a higher precision when separating tissue types. Also, the long exposure of the bone specimen to X-ray in a X-ray micro CT scanner, as compared to clinical CT scanner, allows images of much higher quality with less artefacts. Positioning the temporal bone specimen as close to the X-ray source as possible also permits for a limited volume of interest, to obtain images where tiny mastoid air cells could be detected, which is not possible using clinical CT scans.

A question was raised during a conference presentation, whether histological sections would be a better support to maximize the detection of the smallest mastoid air cells. Ultimately, the answer is yes but as to the author's knowledge, such a study has not been done for multiple reasons. For example, the obtained histological sections are so thin that fragile bony septae separating the mastoid air cells often fold and deform during the slicing. The slicing itself does not allow a perfect continuous concatenate stack of images. Three-dimensional image registration possibly using a GPU (graphical processing unit), is needed to reconstruct a complete mastoid air cell system with all the good features from a histological sectioning. Also, the time to prepare the final 3D volumes would be tremendous, as compared to two or three days to obtain the micro-CT scan from eight human temporal bones.

From a measurement point of view, because the main goal of this study was to show that the surface area and volume are underestimated using X-ray clinical CT scans, the measurement tools used in this study were similar to the tools from previous studies, [37, 64, 77]. There are more advanced image processing tools to perform such measurements, but it remains outside the scope of this particular aim [25].

From a clinical point of view, the aim is indeed fulfilled. From an oral presentation given at an international conference, some ENT professors engaged in this line of work stated that they could benefit from such work.

They realized the limitations of using X-ray clinical CT scanning when investigating the surface area, but some studies need to be done in vivo and therefore requires techniques such a cone-beam CT, or up to date clinical CT.

From a technical point of view, whether the aim is fulfilled or not, the answer is partly more vague. Many manual steps were needed before being able to obtain a measurement of the surface area and volume. Manual intervention is always prone to errors, and therefore more automatized techniques should be used as often as possible. Accumulation of the errors along the different steps may have a tremendous impact on the results. The time spent for each bone specimen was about four to five days. The complexity of the mastoid air cell system required a systematic visual check of the results. To fully fulfil the aim, more advanced image processing tools implying more automation would have been possible. It should be noted, however, that more advanced tools are sometimes too complex to understand for some medical doctors.

Aim 2. Reducing the complexity of the mastoid air cell system representation (Paper V).

The mastoid air cell system is known to be a very complex structure with its interconnected air cells, as illustrated in Fig. 3.5 using volume rendering, and in Fig. 5.1 as triangulated mesh representation. However, almost no literature exists on how complex the mastoid air cell system is. As seen in previous chapters, medial axis or skeleton extraction can be used to "shrink" the binary representation of the MACS down to a set of curves or line segments, while preserving the original shape of the MACS. In this thesis, the original idea of extracting geometrical properties from a structural tensor analysis of the Euclidean distance transform was used. Extracting a medial axis from a distance transform is known to be more accurate than a thinning method [82]. The presence of noise in the data, especially at the borders of the binary representation of the MACS, can create spur branches in the thinning method and unwanted ridges in the distance transform scheme. Using an enhanced representation of the original data, together with the local phase at the level of the structure tensor analysis, the medial axis obtained from the medial axis can be tuned to be less sensitive to noise. The local maxima issue from the Euclidean distance, directly relating to the medial surface defined by [33], is compensated using the local phase to focus on the symmetry of the signal and not on the magnitude. The medial axis can be extracted using only the rank 2 tensor (linear/tubular component) of the estimated structure tensor from the Euclidean distance transform. Only taking the rank 2 tensor may, however, introduce disconti-

nities. To compensate for these discontinuities, another structure tensor, estimated from the medial surface, can be used together with the planar and the isotropic components of the structure tensor estimated from the Euclidean distance, as presented in Eq. 7.6.

The final medial axis, after adjusting the amount of planar and isotropic components, results in a representation such as in Fig. 7.13(left). Compared to the skeleton scheme from [42, 49], the medial axis from this study does not seem to fully represent the topology at the first impression, as shown in Fig. 7.13 when comparing both schemes. However, the medial axis implemented in this study has several advantages over the scheme from [42, 49]. Many air cells in Fig. 7.13(right) are not represented by their skeleton, as pointed out by the yellow stars. Moreover, near the ossicular chain at the superior part of the bone in Fig. 7.13, the skeleton from [42, 49] forms blobs frequently observed along the complete mastoid air cell system skeletonization, as highlighted in Fig. 7.7. Also, Fig. 7.13(right) confirms the presence of spur branches at the level of the big air cell on the lower part of the bone, further highlighted by the white arrow. For the correspondence, the white arrow is also provided for Fig. 7.13(left), where the medial axis from the current study does not show any spur branches. The time to process the medial axis using this analysis was considered faster compared to the scheme from [42, 49].

From a technical point of view, the aim is partly fulfilled. The reason for not being fully fulfilled is due to the fact that the extracted medial axis is not fully continuous. To compensate for this lack of continuity, Eq. 7.6 should be revisited but is believed to be a minor issue since the medial axis does represent all air cells within the bone and is not prone to spur branches. In fact, it should be noted that the lack of continuity can also be interpreted as a compression of information in terms of the medial balls.

From a clinical point of view, extracting the medial axis using either method is already a fulfilment of the aim in itself. But if accuracy, in terms of air cell detection, is preferred over continuity, the medial axis from this study is preferred over the scheme from [42, 49]. Also, the medial axis from this study has its origin from the Euclidean distance transform, which is continuous as compared to a more discrete representation produced by [42, 49].

Aim 3. Investigate in a simplistic manner how inflammation of the mucosa lining mastoid air cells could impact the level of pneumatization of the mastoid through a 3D shape analysis. (Paper V).

From the extraction of the medial axis based on the Euclidean distance transform, medial balls representing the local thickness of the mastoid air cell system could be obtained. As already explained, the medial balls locally represent how broad the air cells are. Also, knowing that the walls of the mastoid air cells are lined with a mucosa susceptible of being inflamed, it is expected that the small air cells would be the first cells to be fully filled by mucosa. If the smaller air cells are only located at the periphery of larger ones, gas exchange will still occur in most locations in the mastoid. However, if small air cells or narrow conducts are present between larger air cells, the inflamed mucosa may eventually lead to isolation of air cell tracks from the rest of the network, thus reducing the gas exchange capabilities of the mastoid. Obtaining information of how large the mastoid air cell system is locally, and whether the extracted medial balls do form natural clusters within the mastoid air cell system, provides a simplistic mean of investigating how inflammation of the mucosa would impact the mastoid air cell system. These points should be taken with care though since they are only suppositions. Indeed, very little is known about the mastoid. But based on similar structures in the body, such as the mucosa lining the nostrils, the mucosa is believed to thicken so much that air cells will be completely filled with mucosa.

From a technical point of view, the aim is partly fulfilled due to several aspects. First, several medial balls may represent a single air cell. An ideal solution would be to fit a deformable shape for each mastoid air cell. Second, extraction of the medial balls is sensitive to the smoothness of the boundaries, defined by the walls of the mastoid air cells. As previously described, the medial axis is not continuous and therefore medial balls were not extracted in the missing parts. Therefore, a more continuous representation would be beneficial. The generation of the output volumes describing each size-based cluster and each location-based cluster was overly long, up two days. However, since it is a preliminary analysis based on an unexplored topic, it can be considered as a good contribution.

From a medical point of view, because it is a preliminary work, it is difficult to know whether such an analysis would be useful or not. Also, the fact that very little is known about the mastoid, especially in terms of physiology, leads to very cautious interpretations of any results, and can only lead to suppositions. However, the analysis was also aimed at pointing out that micro-CT scans of full mastoid air cell systems can be used to perform

analysis and simulation further, than just estimating the surface area and volume. From these aspects, in order to fulfill this aim, more work is needed.

Aim 4. Investigate whether the mastoid air cells do form groups or regions as described by A. Allam and M. Tos (Paper VI).

The very interesting observation from A. Allam, [3], further investigated by M. Tos [85], that mastoid air cells can be grouped based on both their locations and in relation to surrounding structures, is often used as a reference, but there has not been any systematic measure on the extent of the defined regions. Many studies used an inverse problem approach where geometrical aspects of a specific neighbouring anatomical structure are measured and inferences are made on the underlying pneumatization of the mastoid bone specimens investigated, as in [30]. Because very little is known on the mastoid, it remains challenging to accept or reject their hypothesis. The strength of the micro-CT scans is the possibility to investigate how the mastoid air cells are naturally grouped, via a combination of image and shape analysis. The cluster analysis reported in Paper VI, though representing clusters of medial balls, has the potential of investigating whether the groups proposed by A. Allam and M. Tos [3, 85] are consistent, and provides a mean to measure the extent of each group.

From a technical point of view, though it concerned medial balls and not the air cells, the principle remains the same and therefore this aim is considered as fulfilled. Because it is a preliminary work, many aspects need to be improved and more properties can be measured via the obtained clusters.

From a medical point of view, the main interest behind this study is believed to help obtain a mapping of the size range of medial balls for each group defined by A. Allam and M. Tos [3, 85]. This type of analysis may be fruitful when investigating the effect of mucosal inflammation on each group of cells for a specific bone, and statistics can be drawn for multiple specimens. Investigating whether the mastoid air cells do form groups, and obtaining information about the size variation for each of the groups in a systematic manner, can also point out the influence of neighbouring structures on the mastoid air cell system.

9.1.2 Interpretation and geometrical description of the discovered micro-channels

Aim 5. Description of the discovered micro-channels from an anatomical and physiological point of view (Paper II).

As pointed out in Paper II, the micro-channels are believed to house a separate vascular supply for the mastoid mucosa. The study carried out in Paper II also hinted on possible physiological properties, which could have important implications in understanding the role of the middle ear in pressure regulation.

From a technical point of view, the contribution is rather small. The use of a special transfer function enabled visualization of the micro-channels, and gave a good support for visual interpretation of their characteristics. Although not a technical challenge, the aim was fulfilled. It was, however, noticed that some of the micro-channels were not fully represented due to two main reasons: the diameter being close to the voxel resolution and inherent present of noise. This representation issue became the aim of another topic, and presented in the seventh aim.

From a medical point of view, a more recent study has already been carried out from the knowledge gained via the description of the micro-channels, and their possible role in the middle ear pressure [39]. The fact that the knowledge gained from this study was used in a more recent study leads to a positive answer and the aim is considered to be fulfilled.

Aim 6. Investigate the geometrical properties of the micro-channels (Paper III).

Compared to the previous aim, investigating the geometrical properties of the micro-channels relied more on the technical part than on the medical part. Indeed, a structure tensor analysis was implemented to extract the planar, linear/tubular, and isotropic components of the micro-channels. From the description presented in Paper I, the micro-channels do contain a planar part and an isotropic part; especially at their endings either directly in communication with mastoid air cells, or to the lateral part of the mastoid process. Through the analysis performed in Paper II, an interesting observation concerned the presence of triangular-shaped or star-shaped structures, which can be seen as hubs connecting multiple micro-channels together.

From a medical perspective, investigating the geometrical properties of the micro-channels allowed further insights on the question whether the micro-channels do house a separate vascular supply for the mastoid mucosa, as proposed in Aim 5. Also, whenever histological sections are available, there is a systematic search for micro-channels and possible hub structures within the bone matrix. The content of the observed micro-channels is also investigated, so as to gain more knowledge about them structure-wise, which in turn can help better understand their function. Also, the geometrical properties of the micro-channels, extracted from the structure tensor analysis,

further help understanding that their tubular shape has a strong connotation with blood vessels, but more work needs to be done to accept or reject this hypothesis. The only negative point is the lack of a quantitative analysis, normally seen as an important aspect of any study. It should be noted that it would be challenging to perform such a quantitative analysis. Whether the aim of this study was fulfilled, the answer is considered to be yes as it brought a new way of interpreting the geometry of these micro-channels directly from the micro-CT scans.

From a technical perspective, implementing the structure tensor analysis in 3D was not new [29, 55, 76]. The same applies for the determination of the planar, linear or tubular, and isotropic components from the structure tensor, originally done in [26]. However, the use of a structure tensor analysis combined with visualization using volume rendering for such complex data, such as the micro-channels, was new. Whether the aim is fulfilled from a technical view, the answer to this question is positive. The main reason being the uniqueness of such a study, where the geometrical properties highlighted by the structure tensor helps to better understand their characteristics.

Aim 7. Enhance the representation of micro-channels to get a better definition of their course and connectivity (Paper IV).

A point not discussed in the previous aim is the fact that some micro-channels did not have a well-defined path geometrically. By using the geometrical properties from the structure tensor analysis, the representation of the micro-channels could be improved quite significantly. Moreover, the noise hampering the overall data could also be reduced significantly, thus allowing more clear edges between bone and remaining structures, principally composed of air and soft tissues. The results from Paper II showed that enhancement of the micro-CT data was equally good when comparing to alternative enhancement schemes. In fact, for a fair comparison, the alternative scheme known as anisotropic diffusion partial differential equation (PDE) scheme from [24], was tuned based on the results obtained from the structure tensor analysis.

Enhancement of the original micro-CT scans was not specifically targeted to a medical use in the strict sense, rather as a pre-processing step before any further analysis. However, the enhanced micro-CT scans helped to better discern the course of the micro-channels in the bone and when rendered in 3D using volume rendering, it also provided more confidence on the true content of the data, instead of figuring out whether the observed data is simply noise or true weak structures. Whether the aim is fulfilled from a medical point of view, a psychological assessment of how the image enhancement has helped reading the micro-CT scans would be necessary.

In terms of interpreting the anatomy once the data has been enhanced, it has slightly improved the micro-channel representation, and therefore allows to mentally build a more clear map of their connectivity. The aim can therefore be considered fulfilled.

From a technical perspective, the results provided in Paper II, especially in Fig. 9, highlight the fact that the walls of the mastoid air cells do appear less grainy, and therefore more in line with their true definition. From the same paper, in Fig. 7, the combined visualization of p_1 , p_2 , and p_3 , resulting from a structure tensor analysis on the enhanced data (the last two rows), as compared to the original data (top row) clearly show a better definition of the micro-channels represented in red representative of linear/tubular structures. As a side note, the same enhancement scheme was used for the extraction of the medial surface and the medial axis in Paper V, and for the 3D shape analysis in Paper VI. From a technical point of view, this aim is fulfilled.

9.2 Future methodological aspects to consider

In the next two sections, future methodological directions are detailed, either belonging to clinical work or as implementation or improvement of useful technical tools.

9.2.1 Clinical aspects

This section highlights clinical parts that would be interesting to investigate in the future.

Combine histological section and micro-CT scans from the same bone specimen. Histological sections are known to be superior in terms of image content, and the field is very mature with many applications [28]. Though 3D reconstruction from histological sections is possible and getting better and better, X-ray micro-CT scanning will produce a better result by its less invasive nature. Histological sections have a much higher in-plane resolution. Mixing the information from both types of images through registration would be a compromise between the two modalities as explained in the next section, when related to image registration. Moreover, histological sections can provide information the X-ray micro-CT is not capable of, e.g. the mucosa lining the mastoid air cells as explained in the following. It should be noted that it is physically possible to obtain a micro-CT scan of the mucosa lining the mastoid air cells, but the scanning will need to be done on a very small sample taken from a bony chip, and several specimens

will need to be scanned with a similar protocol in order to get sufficient amount of data when inferring a region-independent mean thickness.

Analyze the thickness of the mucosa using more advanced image analysis. From Chapters 3 and 7, more information on the geometrical aspect of the mucosa lining the mastoid air cells is needed. Though the X-ray micro-CT scans cannot reveal the mucosa, since it is slightly smaller than the voxel resolution and further shrunk by the scanning time, histological sections can provide high level of details of the mucosa. Image analysis of the mucosa thickness from histological sections could be done in a more automated manner, rather than the tedious manual work of selecting locations from where to measure the local thickness.

Enhance weak structures such as the tympanic membrane. Enhancing the definition of the thin structures, such as the tympanic membrane, could be of great use when isolating the tympanum from the ear canal, or isolating the inner ear from the tympanum. Though this process can be seen as an extra step in the analysis, using locally enhanced structures could reduce manual intervention, where these membrane are manually over drawn by voxels with a high density value to mimic bone septae. This last method was extensively used in Paper III, to isolate the air inside the temporal bone from the outside air. A large amount of time was used to perform the task, which is prone to error requiring thorough checks.

Extract the medial axis of the micro-channels and study the local thickness. Despite substantial work on the micro-channels, many questions remain open. One possibility is to investigate statistics on the geometrical properties of these micro-channels, such as the mean lumen size, the mean length, the distribution, the curvature, or measuring whether they are spread all over the bone or if they are more concentrated in some regions than others. Investigating the network formed by the micro-channels could further help understanding their potential role in providing blood to the mucosa for gas exchange purposes, and whether their presence is due to a bone evolution resulting from a chronic pathology, or whether they were present from the early stages of the mastoid formation. Scanning children temporal bones using X-ray micro-CT could help in partly answering this question, as discussed in the following item.

Scanning multiple temporal bones from children. Very little information was accessible about the temporal bones scanned with the X-ray micro-CT scans. When observing the available temporal bone specimens using volume rendering, it was possible to tell that the mastoid was fully developed as described by M. Tos and Virapongse et al. [85, 86]. All specimens were considered well pneumatized and therefore healthy. As previously explained in Chapter 3, local thickening of the mucosa, result-

ing from an inflammation, can indicate a previous history of chronic otitis media. However, because this type of information is not available from the X-ray micro-CT scans, only the geometrical aspects of the mastoid air cells can inform us about plausible signs of pathological ears, which was not the case in the sample of bone specimens used during this overall study. Scanning temporal bones from infants of different age could further help answer whether the micro-channels do appear at an early stage, or whether these micro-channels is a natural property of the bone forming the mastoid process or whether they emanate from an osteoneogenesis resulting from a pathology such as a chronic otitis media.

It should be noted that otitis media is more affecting children temporal bones than adult ones, and therefore even if temporal bones from children are available, they might have already been hampered by such pathology. Scanning temporal bones from infants can provide relevant information about the mastoid air cell system development. Because X-ray micro-CT scanning required ex-vivo specimen, due to the lethal long X-ray exposure, only statistics obtained by scanning many specimens can provide information about the bone development. A sample of ten to twenty bone specimens for each age, running from year one to year seven, would be a very good sample population for a preliminary analysis. Together with a sample of adult specimens, such a bank of data could be used for multiple studies as explained in the next two topics: tympanum versus mastoid analysis, or investigating the effect of airspace occlusion due to a swollen mucosa leading to variations in surface area and volume in both airspaces.

Tympanum study. The analysis in Paper IV only concerned determination of the surface area and volume of the mastoid air cell system, but no comparison with the tympanum was provided. The reason for this missing comparison is multifold. A primary reason is the soft borders forming the anatomical interface between the exterior of the temporal bone and the tympanum, i.e. the tympanic membrane, the muscle responsible for opening and closing the Eustachian tube. Once isolated, the physical reduction of the temporal bone prior to scanning leads to exposition of some air cavities, at the level of the inner ear and of the tympanum. To avoid confusion, the mastoid was physically kept intact. Due to these inevitable and unwanted openings, soft membranes such as the round or oval windows, normally separating the tympanum from the inner ear, will not be considered as good separator due to their very thin thickness appearing disrupted during segmentation leading to an unwanted segmentation of inner structures such as the cochlea. Once all these unwanted openings tight, the tympanum and the mastoid air cell system could be segmented out as one structure with techniques reported in Paper IV. Splitting the MACS from the tympanum was then achieved by manually drawing a non-planar border, at the level of the tip of the short process of the incus within the antrum. The surface

area, volume, and surface area to volume are presented in Table 9.1.

Bone	Scan Res. (μm)	SA (cm^2)		V (cm^3)		SAV-ratio (cm^{-1})	
		Mast.	Tymp.	Mast.	Tymp.	Mast.	Tymp.
1	37.60	81	28.80	2.70	1.13	30.00	25.48
2	45.64	326	50.60	13.60	2.15	24.00	23.53
3	45.64	198.3	14.67	10.90	0.97	18.20	15.12
4	45.64	273.3	38.90	13.15	1.90	20.78	20.47

Table 9.1: Comparing the mastoid versus the tympanum in terms of surface area, volume, and surface area to volume ratio. Legend: (Mast.) mastoid, (Tymp.) tympanum, (Scan Res.) Scanning resolution - isotropic size, (SA) surface area, (V) volume, (SAV-ratio) surface area to volume ratio.

Compared to Paper IV, the results are only presented for the highest resolution, $37.60\ \mu\text{m}$ for the first bone and $45.64\ \mu\text{m}$ for the three remaining bones. The study only considered four bone specimens, but the resulting SAV-ratio of the mastoid already appears higher than the SAV-ratio from the tympanum. More bone specimens and more analyses are needed to confirm the results from this preliminary analysis, comparing the MACS versus the tympanum.

A possible way of segmenting the mastoid air cell system and the tympanum could in the future be eased by using the structure tensor analysis together with the medial surface through the watershed segmentation technique, as very briefly pointed in the following.

Segmentation of the mastoid air cells as separate units. The main motivation is to further improve the 3D shape analysis, by obtaining a more accurate representation of each air cell rather than fitting medial balls. A preliminary result using the watershed segmentation is provided in Fig. 9.1.

Each colour represents a separate label (air cell). By extracting each labeled region, further shape analysis could be carried out using contour-based methods, such as Fourier descriptors or shape signatures as pointed in Fig. 7.1. Segmentation of the mastoid air cells as separate units could be used in a simulation study, where some air cells are removed from the full segmented volume mimicking the occlusion of some air cells from chronic diseases such as otitis media. Therefore, controlling the number of regions to be included in the calculations would provide many possible simulations, and would allow to measure the possible impact of occlusion on the remaining mastoid air cells involved in gas exchange.

Figure 9.1: *Preliminary study on a watershed segmentation of the airspaces composing a temporal bone. Each colour represents an air cell.*

9.2.2 Technical aspects

This section highlights technical parts that would be interesting to investigate in the future.

Further work on the cluster analysis. As illustrated in Fig. 7.3, narrow conduits may block some underneath air cells, it may therefore be of interest to further investigate how each location-based cluster is structured, see Fig. 9.2. Indeed, by further studying the distance of each size-based clustered medial ball to the centroid within each location-based cluster, metrics could be extracted revealing how the medial balls of different size are located in relation to each others; e.g. going from very large medial balls towards very small medial balls as in Fig. 9.2(a), or simply in an unstructured manner as in Fig. 9.2(b) suggesting narrow conduits and possible occlusion for small cells.

Watershed segmentation of the cells. As introduced in the previous section, a watershed segmentation of the airspace would possibly be a great addition to the tools used for both a more advanced 3D shape analysis and for a simulation of occlusion of air cells that could happen due to a mucosal inflammation. Watershed segmentation is a mature technique and which is used in many applications. More information about Watershed segmentation can be found in [52, 60, 73].

Figure 9.2: Future analysis of the intra-location-based clusters for two different types of clusters: (a) structured, and (b) unstructured. The medial balls in (a) are following a certain order in appearance with larger balls close to the centroid and smaller medial balls at the outskirts. The medial balls in (b) are placed in a more random manner where for each size-class, the distance from the centroid varies a lot respectively.

Multi-scale integration. Only a singular scale was used during the implementation of the structure tensor analysis. The filters were only optimized for a specific center frequency, a multi-scale implementation could further help the determination of the Euclidean distance transform, where a binary input volume is needed. Instead of applying a classical thresholding on the original data, a multi-scale representation of the norm of the structure tensor analysis could be used for the thresholding, followed by the Euclidean distance transform. Multi-scale integration could also be used in combination with downsampling of the original data.

From the previous paragraph, a multi-scale integration of the norm of the structure tensor can be seen as a tensor certainty.

Tensor certainties Certainty measures can be used to govern further processing of the data, such as in the case of image enhancement or registration between a target and a reference volume. Using a quadrature filtering scheme for estimating the structure tensor, as previously described, several certainty measures can be extracted:

- structure certainty based on spatial information over a neighborhood,
- structure certainty based on scales,
- phase certainty over scales,
- phase certainty over orientation.

Table 9.2 resumes all possible combinations when either using the magnitude of the structure tensor or its phase over the space, scale, and orientation. As can be noticed, two options are not possible, i.e. phase certainty over space and structure certainty over orientation. The magnitude of the structure tensor does not carry any information about orientation, and therefore the structure certainty cannot be estimated over the different filter orientations.

	Structure	Phase
Space	✓	✗
Scale	✓	✓
Orientation	✗	✓

Table 9.2: Possible structure and phase certainties in relation to space, scale, and orientation.

Structure consistency over space provides a measure on how quickly the magnitude of the structure tensor changes in shape within a spatial neighborhood. Structure consistency over scale provides information on how much the magnitude of the structure tensor changes over different scales. Phase consistency over scale contains information about how quickly the phase changes over the scales for all the filters, for each single voxel. It is independent of the filter size. The certainty measures can be combined together, and weighting each of them respectively.

EDT obtained using a fast marching scheme where the voxels defining the norm of a structure tensor analysis are used as seeds.

One feature that has not been discussed in the estimation of the Euclidean distance transform, is the possibility to use a fast marching scheme to compute a signed distance transform, similar to the Euclidean distance transform based on seeds directly emanating from the norm of the structure tensor. In the case of the mastoid air cells, at a single scale, the norm of the structure tensor represents the edges formed by the walls of the mastoid

air cells with a certain thickness. When considering all the points from all the edges as seeds, a speed volume made of ones can be created, which a fast marching scheme can be applied on using the seeds. After a certain number of iterations, the resulting volume forms a distance-based volume similar to the Euclidean distance transform.

The tensor certainties from the previous discussion topic can then be used as a mask, to preserve the equivalent Euclidean distance transform corresponding to the mastoid air cells while excluding the remaining part. This line of work has been briefly tested and appeared to work as it should, but more work is needed to assess its quality. But, if the hypothesis is correct, there is no need to perform a binary thresholding for estimating the Euclidean distance transform in this line of work. The fast marching scheme used to perform this preliminary analysis is based on an implementation further detailed in [31].

Improve the medial axis equation. Eq. 7.6 is one possible way of extracting the medial axis, but not the only one. The main idea behind Eq. 7.6 was to increase the flexibility, by keeping a certain amount of planar and isotropic parts from the medial surface. As discussed, in Chapter 7, several criterions need to be respected in order to be defined as a valid skeleton: *homotopy*, *invariance*, *thinness*, *centeredness*, *smoothness*, *detail resolution*, *regularization* and *level of details*, *reconstructability*, and *computational scalability*.

These criteria are explained in [82]. For the current implementation, the last two features are not so important, since there is no interest in reconstructing the original shape from the medial balls, and the computational scalability related to computation power was not a critical point. However, while the *Skeleton3D* used for comparison was taking a long time, easily above an hour for a complete bone, the medial axis extraction for the complete bone took about five minutes. More interestingly, the *homotopy* of the medial axis, in the current implementation, seemed to be more representative of the shape formed by the mastoid air cell system. The *invariance* of the medial axis current scheme is clearly invariant to scaling, rotation and translation of the original data, since it is extracted from the Euclidean distance transform. *Thinness* is probably the more critical part of the current implementation, since planar and isotropic parts of the medial surface are kept along the medial axis to provide a more continuous representation. However, skeletons extracted from thinning often result in locations with multiple voxels representing a junction. Nevertheless, more work needs to be done to investigate how thin the current scheme can become. The *centeredness* is guaranteed by the Euclidean distance transform, and by ensuring that not too much noise hampers the borders of the air cell representation.

A qualitative observation of the current medial axis algorithm indicates a rather good *smoothness* but no metric was used to investigate on how smooth the medial axis is. As for the *detail resolution*, the medial axis algorithm implemented in this study seems to have a better topology than the skeleton obtained using the *Skeleton3D* module from MeVisLab, as further pointed out in Fig. 5.5 by the yellow stars. In terms of *regularization and level of details*, it should be noted that the current implementation is still at an experimental phase, and more work on Eq. 7.6 needs to be done to ensure a certain robustness in terms sensitivity and to compare the results with other techniques reviewed in [82].

Propose a conformal mapping. Another line of work is to consider a complimentary shape analysis that is more based on the surface representing the outer surface of the mastoid air cell system, as the one illustrated in Fig. 5.1. A potential idea is to apply a spherical conformal mapping transforming the original mesh illustrated in Fig. 5.1 into a unit sphere, where high and low concentration would locally form features. Once the spherical mapping is achieved, a rectangular mapping can be applied, making it easier to see any patterns on the mesh. Repeating the same principle over several temporal bone specimens would allow comparison of common patterns, and would also allow a statistical shape analysis. Similar work has previously been performed for the brain [51, 70].

Simulate the pneumatization through diffusion. A possible extension of this work would be to use the work from Zollikofer et al. [94]. In their study, bone pneumatization is simulated through a combination of a Laplace growth model and a Poisson growth model, where the Laplace growth model is used to mimic the invasion of mucous epithelial tissue while the Poisson growth model mimics the formation of air-filled spaces within a growing bone. Applying their method on the micro-CT scans may improve our understanding of the pneumatization process by placing a seed (reference point) at the level of the attic (in the tympanum) and letting the model grow up until it fills all the mastoid air cells.

Perform image registration between several bone specimens. To allow comparison between different temporal bone specimens, knowing that the mastoid air cell system can vary tremendously in shape, image registration needs to be done. As briefly introduced in [23], image registration can be defined as the process of mapping input images with a reference image. Indeed, when the number of bone specimens available is low (<10), a permutation approach can be used where image registration between two temporal bone specimens is performed many times. By investigating which permutation that provides the smoothest displacement fields, it is possible to use the resulting temporal bone as a template to which the remaining temporal bone specimens can be registered. When the number of tempo-

ral bones is larger, a possible solution is to select one temporal bone as a reference volume, then register all the remaining temporal bone specimens to that reference, calculate an average over all registered temporal bone specimens, which then becomes the template, and register each temporal bone specimen to that template (including the one used for the reference). Due to the large amount of data imposed by the micro-CT scanning, image registration can be done through the graphical processing unit (GPU) as further detailed in [22].

Before concluding, besides the discovery of the micro-channels, some additional unreported findings are presented. They deserve some attention as they may provide more information about the possible roles of the mastoid in the temporal bone. First strands were found deep inside the mastoid air cell system. Second, pores along the bony septae were commonly observed. Also, spicules protruding from the walls of the mastoid air cells has also raised the attention by their consistent findings. A section is dedicated to each of these unreported findings and possible roles are suggested though further research is needed.

9.3 Strands and membranes within the MACS

Figure 9.3: *Mucosal strands isolating the central air cell from the surrounding ones. Source: Temporal Bone Histology Lab, Otorhinolaryngology, University Zurich (permission granted by Prof. Dr. Med. Alexander Huber).*

Originally believed to originate from the fact that the bone specimens are dried, the mucosa normally lining the air cells was thought to be loose and form strands barely visible on the micro-CT scans. However, when carefully observing these strands, either on 2D slices and even better in

3D while using volume rendering, these epithelial strands have well defined pyramidal bases at their attachment locations. These structures, though very thin, have only been reported in the tympanum between the tympanic walls and the ossicles, known as the mucosal folds. Within the mastoid air cells, to the author's knowledge, this feature has never been illustrated in 3D in the literature so far.

This finding is recurrent and has been observed in almost all specimens available for scanning. Fig. 9.3 illustrates the presence of two strands closing the openings between bony septae in the MACS from a histological section, highlighted by the two yellow rectangles. In this histological section, the central air cell becomes isolated from the surrounding air cells. A positive aspect of these strands is that they further increase the surface area to volume ratio, especially if the inflammation is only temporary. Interestingly, it is possible to observe these strands using micro-CT scans. Fig. 9.4 illustrates the presence of strands found in the apical air cells of the mastoid pointed by arrows.

Figure 9.4: *Strands (pointed by arrows) together with membranes separating some air cells at the apex of the mastoid.*

The top three arrows show differences in structure: the top left yellow arrow shows a full membrane, while the top right yellow illustrates a standard strand. The bottom red arrow points towards a strand possibly containing a structure along its path. The two remaining yellow arrows point to more membraneous strand types as the top left arrow, proving that such strands can also exist further lateral in the mastoid air cell system.

In Fig. 9.5, multiple strands can be observed at the level of the aditus ad antrum towards the midline of the mastoid air cell system and pointed by

yellow arrows. Interestingly, at the intersection between multiple strands, a triangular patch of mucosa is visible. The yellow circle particularly points out a strand having a local increase in its thickness, possibly indicating the presence of a Pacinian corpuscle attached to it.

Figure 9.5: *Multiple strands from at the level of the antrum highlighted by the yellow arrows. Notice the pyramidal node at the intersection pointed by the yellow stars (scan resolution 38 μm).*

Another interesting observation is the presence of small holes carved into the septae separating the air cells, as further detailed in the next section.

9.4 Pores in the septae

Pores located within the mastoid air cell system in human temporal bones have not, to the author's knowledge, been reported before. Fig. 9.6 illustrates the presence of many pores in the bone septae, highlighted with the white stars. The origin and function of these ventilation pores is not understood. Plates in trabecular plates found in long bones, such as the femur, also contain openings similar to the pores found in this study.

In terms of physiology, when trying to find an analogy with the lungs, the pore of Kohn comes in mind. A Kohn pore is a tiny opening in the alveoli, allowing communicating between two adjacent alveoli. The main purpose of this shunt is to allow oxygen to reach alveoli which are isolated at the level of the alveolar ducts or higher up in the respiratory trunk or when the lungs are temporally collapsed. However, while the lungs can collapse, the mastoid air cell system is fixed in structure by the bone septae separating the air cells. The soft epithelial lining is not well represented in X-ray micro-CT, if at all, making it impossible at this level to understand their

Figure 9.6: *Ventilation pores in the mastoid air cell system. For better localization in the temporal bone, part of the ossicular chain is shown on the top of this coronal view showing the lower half of the mastoid bone, with the semi-circular canals at the top left. The white stars point out few of the ventilation pores.*

role. Some of the small pores might be fully covered by mucosa.

It is, however, believed that some pores may be too large to be fully covered by mucosa, and therefore still allow air to pass through them even if covered by mucosa. Moreover, these bone specimens are from corpses and were fixated. Therefore, the mucosa lining the walls of the mastoid air cells is no longer nourished by blood, and is no longer playing its role in pressure regulation. When in the presence of blood, the mucosa lining these pores may expand up to a point where the pore becomes temporally occluded, and open again when the mucosa shrinks in thickness. A mechanism of opening-closing of these pores could be possible. Interestingly, some pores can be found at the attic of the tympanum right at the level of the antrum near the posterior level of the malleus head, see Fig.9.7.

Figure 9.8: *Formation of the spicules at the early stages of the mastoid formation. Legend: (1) Classical organization of a bone with trabecular rods forming trabecular spaces, filled with bone marrow between two cortical bone layers. In (2), the rods are rupturing due to multiple sources: muscle pulling the bone downward, epithelial layer pushed by air coming from the forming tympanum. In (3), the tip of the ruptures rods slowly erode, giving the appearance of smooth spikes. In (4), the mucosa and epithelial lining cover all the walls and spicules of the mastoid air cell system.*

cosa lining the forming tympanum invades the mastoid bone. During this process, some of the trabecular rods and plates slowly split through bone resorption via the osteoblast-osteoclast interaction. It is possible that this process stabilises at eight years of age and hence the presence of spicules or incomplete plates within some of the mastoid air cells. In Fig. 9.8(1), a very simplistic representation of the mastoid bone with some trabeculae and bone marrow before the invasion of the epithelial mucosa is given. In Fig. 9.8(2), the epithelial mucosa has started to penetrate the cancellous part of the bone, with some bone resorption of some of the trabecular rods and plates illustrated. In Fig. 9.8(3), the epithelial lining takes more and more spaces, and the broken rods present a more spicule like shape. In Fig. 9.8(4), the epithelial lining is fully covering all the bone inner surface. At this stage, the mastoid bone is fully pneumatized. A volume rendering of the spicules using micro-CT, highlighted by yellow stars and covered with mucosa, is provided in Fig. 9.9. As explained earlier, the spicules are often the supports for mucosal strands, as pointed by the multiple arrows for the spicule located to the outmost right emphasized with the right yellow star.

Though cats have a bulla instead of a mastoid as in humans, Sula et al. [75] reported the occasional formation of irregular bony spicules in cats

Figure 9.9: *Spicules covered by mucosa and strands, respectively marked with yellow stars and arrows.*

with a history of chronically affected ears . Per Cayé-Thomasen et al. [11] reported the formation of bone in the presence of a pathological bone, such as in chronic otitis media. To investigate whether the bony spicules result from the original bone formation, or from a remodelling due to chronic pathology, more clinical work needs to be done and is beyond the scope of this work.

9.6 Conclusion

The content of this thesis was not aimed at being groundbreaking from a technological point of view, and does not intend to contradict previous methods used in the field of ENT, especially when measuring geometrical properties of the mastoid air cell system in human temporal bone via micro-CT scans. The overall main goal of this PhD was to open bridges between two fascinating worlds: the mastoid air cell system and advanced image processing tools, allowing to extract more information than the sole extraction of surface area and volume. Flexibility of the algorithm, while preserving a certain robustness of the methods, was considered as a higher priority when selecting the image processing tools. The flexibility allowed tuning the different algorithms based on the anatomical structures, while robustness was important to preserve consistency over the different studies. Through the thesis, visualization has been a very important tool. Indeed,

the complexity imposed by the mastoid air cell system implies a constant verification of any obtained result. The importance of visualization in this study has also shaped some of the scientific contributions to be more qualitative than quantitative. Another good reason for this choice is the fact that so little is known about the geometrical properties of the mastoid air cell system that there is a need to be acquainted with its intricate structures.

Bibliography

- [1] Illustration of the eustachian tube at the level of nasopharynx. <http://upload.wikimedia.org/wikipedia/commons/e/e3/Gray855.png>.
- [2] A. Ilea, A. Butnaru, S. Sfrângeu, M. Hedeşiu, C. Dudescu, P. Berce, H. Chezan, L. Hurubeanu, V. Trombiţaş, R. Câmpian, and S. Albu. Role of mastoid pneumatization in temporal bone fractures. *American Journal of Neuroradiology*, 35(7):1398–1404, 2014.
- [3] A. Allam. Pneumatization of the temporal bone. *Ann. Otol. Rhinol. Laryngol.*, 78:48–64, 1969.
- [4] B. Anson, J. Wilson, and J. Gaardsmoe. Air cells of petrous portion of temporal bone in a child four and a half years old: A study based on wax plate reconstructions. *Archives of Otolaryngology*, 27(5):588–605, 1938.
- [5] B. Ars and N. Ars-Piret. Middle ear pressure balance under normal conditions. specific role of the middle ear structure. *Acta Otorhinolaryngol. Belg.*, 48:339–342, 1994.
- [6] B. Ars, F. Wuyts, P. V. de Heyning, I. Miled, J. Bogers, and E. V. Marck. Histomorphometric study of the normal middle ear mucosa. *Acta Otolaryngol. (Stockh)*, 117:704–707, 2006.
- [7] M. Berg, O. Cheong, M. Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 3rd ed. edition, 2008. ISBN 9783540779735.
- [8] S. Biasotti, B. Falcidieno, D. Giorgi, and M. Spagnuolo. *Mathematical Tools for Shape Analysis and Description*. Synthesis Lectures on Computer Graphics and Animation. Morgan & Claypool Publishers, 2014. ISBN 9781627053631.
- [9] H. Blum. A Transformation for Extracting New Descriptors of Shape. *Models for the Perception of Speech and Visual Form*, Cambridge, MIT Press, pages 362–380, 1967.
- [10] T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974.
- [11] P. Cayé-Thomasen, A. Hermansson, and M. T. K. Prellner. Pathogenesis of middle ear adhesions. *The Laryngoscope*, 106(4):463–469, 1996.
- [12] P. Cayé-Thomasen, A. Hermansson, M. Tos, and K. Prellner. Bone modeling dynamics in acute otitis media. *The Laryngoscope*, 109(5):723–729, 1999.
- [13] O. Cros. Image analysis and visualization of the human mastoid air cell system. Licentiate thesis 1730, Linköping University, Department of Biomedical Engineering, 2015.
- [14] O. Cros, A. Eklund, M. Gaihede, and H. Knutsson. Local size descriptor of mastoid air cells in human temporal bone based on a structure tensor analysis

- from an euclidean distance transform. *Computerized Medical Imaging and Graphics*. Submitted.
- [15] O. Cros, A. Eklund, M. Gaihede, and H. Knutsson. Enhancement of micro channels within the human mastoid bone based on local structure tensor analysis. *IEEE The sixth International Conference on Image Processing Theory, Tools and Applications (IPTA16)*. 2016. In Press.
 - [16] O. Cros, A. Eklund, M. Gaihede, and H. Knutsson. Surface and curve skeleton from a structure tensor analysis applied on mastoid air cell in human temporal bones. *2017 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, 2017.
 - [17] Z. Csákányi, G. Katona, E. Josavai, F. Mohos, and I. Sziklai. Volume and surface of the mastoid cell system in otitis media with effusion in children: a case-control study by three-dimensional reconstruction of computed tomographic images. *Otology And Neurotology*, 32:64–70, 2011.
 - [18] W. Doyle. Increases in middle ear pressure resulting from counter-diffusion of oxygen and carbon dioxide into the middle ear of monkeys. *Acta Otolaryngol.*, 117(5):708–713, 1997.
 - [19] W. Doyle. Mucosal surface area determines the middle ear pressure response following establishment of sniff-induced underpressures. *Acta Otolaryngol.*, 119(6):695–702, 1999.
 - [20] R. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. *Computer Graphics*, 22(4):65–74, 1988.
 - [21] A. Eklund, M. Andersson, and H. Knutsson. True 4D image denoising on the GPU. *International Journal of Biomedical Imaging*, 2011:1–16, 2011.
 - [22] A. Eklund, P. Dufort, D. Forsberg, and S. LaConte. Medical image processing on the GPU - Past, present and future. *Medical Image Analysis*, 17(8):1073–1094, 2013.
 - [23] F. E. A. El-Gamal, M. Elmogy, and A. Atwan. Current trends in medical image registration and fusion. *Egyptian Informatics Journal*, 17(1):99 – 124, 2016.
 - [24] J. Fehrenbach and J.-M. Mirebeau. Sparse Non-negative Stencils for Anisotropic Diffusion. *Journal of Mathematical Imaging and Vision*, 49(1):123–147, 2014.
 - [25] R. C. Gonzalez, R. E. Woods, and S. L. Eddins. *Digital Image Processing using MATLAB*. Pearson Prentice Hall, 2004. ISBN 0130085197.
 - [26] G. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1995. ISBN 0792395301.
 - [27] H. Gray and H. Carter. *Gray’s Anatomy (Illustrated With 1247 Coloured Well Drawing Engravings)*. eKitap Projesi, 2016. ISBN 9786059654708.
 - [28] M. Gurcan, L. Boucheron, A. Can, A. Madabhushi, N. Rajpoot, and B. Yener. Histopathological image analysis: A review. *IEEE reviews in biomedical engineering*, 2:147–171, 2009.
 - [29] L. Haglund. *Adaptive Multidimensional Filtering*. Linköping Studies in Science and Technology. Dissertations No. 284, Linköping University, Sweden, 1992.

- [30] S.-J. Han, M. Song, J. Kim, W.-S. Lee, and H.-K. Lee. Classification of temporal bone pneumatization based on sigmoid sinus using computed tomography. *Clinical Radiology*, 62(11):1110–1118, 2007.
- [31] M. S. Hassouna and A. A. Farag. Multistencils fast marching methods: A highly accurate solution to the eikonal equation on cartesian domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(9):1563–1574, 2007.
- [32] E. Hentzer. Histological studies of the normal mucosa in the middle ear, mastoid cavities and eustachian tube. *Annals of Otolology, Rhinology and Laryngology*, 79:825–833, 1970.
- [33] W. H. Hesselink and J. B. T. M. Roerdink. Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(12):2204 – 2217, 2008.
- [34] C. Hodez, C. Griffaton-Taillandier, and I. Bensimon. Cone-beam imaging: Applications in ENT. *European Annals of Otorhinolaryngology, Head and Neck Diseases*, 128(2):65–78, 2011.
- [35] C. T. Huang and O. R. Mitchell. Rapid euclidean distance transform using grayscale morphology decomposition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 695–697, 1991.
- [36] N. Ikonomatakis, K. Plataniotis, M. Zervakis, and A. Venetsanopoulos. Region growing and region merging image segmentation. In *13th International Conference on Digital Signal Processing Proceedings*, volume 1, pages 299–302, 1997.
- [37] M. Isono, K. Murata, H. Azuma, M. Ishikawa, and A. Ito. Computerized assessment of the mastoid air cell system. *Auris Nasus Larynx*, 26(2):139–145, 1999.
- [38] B. Jähne. *Digital Image Processing. 5th revised and extended edition*. 2002. ISBN 3540677542.
- [39] P. Jensen and M. Gaihede. Congestion of mastoid mucosa and influence on middle ear pressure - Effect of retroauricular injection of adrenaline. *Hearing Research Special Issue MEMRO*, 340:121–126, 2016.
- [40] M. K. Kalra, M. M. Maher, T. L. Toth, B. Schmidt, B. L. Westerman, H. T. Morgan, and S. Saini. Techniques and applications of automatic tube current modulation for CT. *Radiology*, 233(3):649–657, 2004.
- [41] D. G. Kendall. The diffusion of shape. *Advances in Applied Probability*, 9(3):428–430, 1977.
- [42] M. Kerschnitzki, P. Kollmannsberger, M. Burghammer, G. Duda, R. Weinkamer, W. Wagermaier, and P. Fratzl. Architecture of the osteocyte network correlates with bone material quality. *Journal of Bone and Mineral Research*, 28(8):1837–1845, 2013.
- [43] H. Knutsson. *Filtering and Reconstruction in Image Processing*. PhD thesis, Linköping University, 1982.
- [44] H. Knutsson, M. Andersson, and J. Wiklund. Advanced filter design. In *Scandinavian conference on image analysis (SCIA)*, 1999.
- [45] H. Knutsson and M. T. Andersson. Optimization of sequential filters. In *In Proceedings of the SSAB Symposium on Image Analysis*, pages 87–90, 1995.

- [61] A. Miracle and S. Mukherji. Conebeam ct of the head and neck, part 2: Clinical applications. *American Journal of Neuroradiology*, 30(7):1285–1292, 2009.
- [62] M. Miura, H. Takahashi, I. Honjo, S. Hasebe, and M. Tanabe. Influence of the gas exchange function through the middle ear mucosa on the development of sniff-induced middle ear diseases. *The Laryngoscope*, 108:683–686, 1998.
- [63] I. Nyström, G. S. di Baja, and S. Svensson. Curve Skeletonization by Junction Detection in Surface Skeletons. In *C. Arcelli et al. (Eds): IWVF4 Capri, LNCS 2059, Springer-Verlag Berlin Heidelberg*, pages 229–238, 2001.
- [64] M. Park, S. Yoo, and D. H. Hoon. Measurement of surface area in the human mastoid air cell system. *J. Laryngol. Otol.*, 114:93–96, 2000.
- [65] S. P. Raman, M. Mahesh, R. V. Blasko, and E. K. Fishman. CT scan parameters and radiation dose: practical advice for radiologists. *Journal of the American College of Radiology : JACR*, 10(11):840–846, 2013.
- [66] J. Sadé and A. Ar. Middle ear and auditory tube: Middle ear clearance, gas exchange, and pressure regulation. *Otolaryngol. Head Neck Surg.*, 116:499–524, 1997.
- [67] P. K. Saha, G. Borgefors, and G. S. di Baja. A survey on skeletonization algorithms and their applications. *Pattern Recognition Letters*, 76:3 – 12, 2016.
- [68] L. Serino and G. S. d. Baja. Pruning the 3d curve skeleton. *2014 22nd International Conference on Pattern Recognition*, pages 2269–2274, 2014.
- [69] M. Sezgin and B. Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic Imaging*, 13(1):146–165, 2007.
- [70] Y. Shi, R. Lai, and A. Toga. Conformal mapping via metric optimization with application for cortical label fusion. *Proceedings of the conference Information Processing in Medical Imaging*, 23:244–255, 2013.
- [71] A. J. Silver, I. Janecka, J. Wazen, S. K. Hilal, and J. N. Rutledge. Complicated cholesteatomas: Ct findings in inner ear complications of middle ear cholesteatomas. *Radiology*, 164(1):47–51, 1987.
- [72] A. Sobiecki, A. Jalba, and A. Telea. Comparison of curve and surface skeletonization methods for voxel shapes. *Pattern Recognition Letters*, 47:147 – 156, 2014.
- [73] P. Soille and L. Vincent. Determining watersheds in digital pictures via flooding simulations. *Visual Communications and Image Processing '90*, 1360:240–250, 1990.
- [74] J. Stewart. The histopathology of mastoiditis. *Proceedings of the Royal Society of Medicine*, 21(10):1743–1758, 1928.
- [75] M. M. Sula, B. L. Njaa, and M. E. Payton. Histologic characterization of the cat middle ear. *Veterinary Pathology*, 51(5):951–967, 2014.
- [76] B. Svensson. *A Multidimensional Filtering Framework with Applications to Local Structure Analysis and Image Enhancement*. Linköping Studies in Science and Technology. Dissertations No. 1171, Linköping University, Sweden, 2008.
- [77] J. Swarts, B. C. Doyle, C. Alper, and W. Doyle. Surface area-volume relationships for the mastoid air cell system and tympanum in adult humans:

- Implications for mastoid function. *Acta oto-laryngologica*, 130(11):1230–1236, 2010.
- [78] J. Swarts, S. Folay, C. Alper, W. Doyle, and J. William. Mastoid geometry in a cross-section of humans from infancy through early adulthood with a confirmed history of otitis media. *International journal of pediatric otorhinolaryngology*, 76(1):137–141, 2011.
 - [79] J. D. Swarts, B. Doyle, C. Alper, and W. Doyle. Surface area-volume relationships for the mastoid air cell system and tympanum in adult humans: Implications for mastoid function. *Acta Oto-Laryngologica*, 130(11):1230–1236, 2011.
 - [80] J. D. Swarts, B. Doyle, and W. Doyle. Surface area-volume relationships for the mastoid air cell system in adult humans. *The Journal of Laryngology and Otology*, 125(6):580–584, 2011.
 - [81] J. Swartz and L. Loevner. *Imaging of the Temporal Bone*. Thieme Publishers Series, 2011. ISBN 9781588906496.
 - [82] A. Tagliasacchi, T. Delame, M. Spagnuolo, N. Amenta, and A. Telea. 3d skeletons: A state-of-the-art report. *Computer Graphics Forum*, 35(2):573 – 597, 2016.
 - [83] A. Telser, J. Young, and K. Baldwin. *Elsevier’s Integrated Histology*. Mosby/Elsevier. Inc. Philadelphia, 2007. ISBN 0323033881.
 - [84] N. Todd, R. Pitts, I. Braun, and H. Heindel. Mastoid size determined with lateral radiographs and computerized tomography. *Acta Otolaryngologica*, 103:226–231, 1987.
 - [85] M. Tos. *Manual of Middle Ear Surgery. Vol.2: Mastoid Surgery and Reconstructive Procedures*. Georg Thieme Verlag, 1995. ISBN 0865775893.
 - [86] C. Virapongse, M. Sarwar, S. Bhimani, C. Sasaki, and R. Shapiro. Computed tomography of temporal bone pneumatization: 1. Normal pattern and morphology. In *American Journal of Roentgenology*, volume 145, pages 473–481. 1985.
 - [87] J. Vlassenbroeck, M. D. B. Masschaele, V. Cnudde, L. V. Hoorebeke, and P. Jacobs. Software tools for quantification of x-ray microtomography at the UGCT. *Nuclear Instruments and Methods in Physics Research*, A(580):442–445, 2007.
 - [88] J. Vrabec, S. Champion, J. G. R. Johnson, and G. Chaljub. 3D CT 534 imaging method for measuring temporal bone aeration. *Acta Otolaryngol*, 122:831–835, 2002.
 - [89] H. Wang, C. Northrop, B. Burgess, M. Liberman, and S. Merchant. Three-dimensional virtual model of the human temporal bone: A stand-alone, downloadable teaching tool. *Otology and neurotology: official publication of the American Otological Society, American Neurotology Society and European Academy of Otology and Neurotology*, 27(4):452–457, 2006.
 - [90] J. Wu. *Advances in K-means Clustering: A Data Mining Thinking*. Springer Publishing Company, Incorporated, 2012. ISBN 9783642298066.
 - [91] J. Yi and J. B. Ra. A locally adaptive region growing algorithm for vascular segmentation. *Int. J. Imaging Systems and Technology*, 13:208–214, 2003.
 - [92] D. Zhang and G. Lu. Review of shape representation and description techniques. 37:1–19, 2004.

-
- [93] J. Zhang, C.-H. Yan, C.-K. Chui, and S.-H. Ong. Fast segmentation of bone in CT images using 3d adaptive thresholding. *Computers in Biology and Medicine*, 40(2):231–236, 2010.
 - [94] C. Zollikofer and J. Weissmann. A morphogenetic model of cranial pneumatization based on the invasive tissue hypothesis. *The Anatomical Record, WILEY-LISS, INC.*, 291:1446–1454, 2008.

Papers

The articles associated with this thesis have been removed for copyright reasons. For more details about these see:

[http://urn.kb.se/resolve? urn:nbn:se:liu:diva-137288](http://urn.kb.se/resolve?urn:nbn:se:liu:diva-137288)