



***Towards Automated Human Identification
Using Dental X-ray Images***

by

Jan-Vidar Ølberg

Supervisor

Associate Professor, Morten Goodwin, Ph.D

This master's thesis is carried out as a part of the education at the University of Agder and is therefore approved as a part of this education. However, this does not imply that the University answers for the methods that are used or the conclusions that are drawn.

The University of Agder
Faculty of Engineering and Science
Department of Information and Communication Technology

Grimstad, May 26, 2015

Abstract

Systems for automated human identification from dental X-ray images can be used to greatly reduce the necessary effort spent today by dental forensics experts.

In this work a new methodology is proposed to create a system for automated dental X-ray identification. The methodology includes both state-of-the-art methods and a novel method for separating a dental X-ray image into individual teeth. The novel method is based on lowest cost pathfinding and is shown to achieve comparable results to the state-of-the-art. In experiments it is able to separate 88.7% of the teeth in the test images correctly.

The identification system extracts tooth and dental work contours from the dental X-ray images and uses the Hausdorff-distance measure for ranking persons. The results of testing the system on a new data set show that the new method for dental X-ray separation functions well as a component in a functional identification system and that the methodology on the whole can be used to identify persons with comparable accuracy to related work. In 86% of cases, the correct person is ranked highest. This accuracy increases to 94% when the five highest ranked images are considered.

Due to small distances in similarity between highest ranked individuals, doubts are raised concerning the scalability of the method. This is seen as a matter of expansion, such as refining features, rather than redesign. The conclusion is that the proposed methodology, including the path-based method of separation, performs well enough to be worth consideration when designing an automated dental identification system.

Preface

This thesis was written for the IKT590 Master Thesis course part of the ICT Master programme at University of Agder, Norway.

I would like to thank my supervisor associate professor Morten Goodwin, Ph.D for presenting me with the problem and supporting me throughout. I would also like to thank Jan Tore Olsen, from Arendal Tannlegesenter, who helped with dentistry domain insight and critically for providing images and helping to create and verify the data set.

Last, but not least, thanks to professor Ole-Christoffer Granmo, Ph.D and the others part of group supervision for advice and support during the project period.

Contents

| | |
|-------------------------------------------------|-----------|
| Preface | I |
| Contents | II |
| List of Figures | IV |
| List of Tables | VI |
| 1 Introduction | 1 |
| 1.1 Problem Statement | 2 |
| 1.2 Motivation | 3 |
| 1.3 Thesis Outline | 4 |
| 1.4 Assumptions and Limitations | 4 |
| 1.5 Definitions | 5 |
| 2 Theoretical Background | 6 |
| 2.1 The Dental X-ray Image | 6 |
| 2.2 State-of-the-Art | 7 |
| 2.2.1 Enhancement | 8 |
| 2.2.2 Separation | 9 |
| 2.2.3 Feature Extraction and Matching | 13 |
| 2.3 Methods and Techniques | 15 |
| 2.3.1 Gamma Transformation | 15 |

CONTENTS

| | | |
|----------|---------------------------------------------|-----------|
| 2.3.2 | Grayscale Morphology | 16 |
| 2.3.3 | Adaptive Thresholding | 19 |
| 2.3.4 | Region Growing | 20 |
| 2.3.5 | Boundary Following | 21 |
| 2.3.6 | Integral Intensity Projection | 22 |
| 2.3.7 | Dijkstra's Algorithm | 23 |
| 2.3.8 | Canny Edge Detection | 24 |
| 2.3.9 | Active Contour Without Edges | 26 |
| 2.3.10 | Bi-Directional Hausdorff Distance | 27 |
| 3 | Approach | 28 |
| 3.1 | System Overview | 28 |
| 3.2 | Feature Extraction | 31 |
| 3.2.1 | Dental Work Contours | 32 |
| 3.2.2 | Enhancement | 36 |
| 3.2.3 | Tooth Contours | 41 |
| 3.3 | Matching | 52 |
| 4 | Results and Discussion | 55 |
| 4.1 | Separation | 55 |
| 4.1.1 | Jaw Separation | 55 |
| 4.1.2 | Teeth Separation | 59 |
| 4.2 | Image Retrieval and Ranking | 61 |
| 5 | Conclusion | 69 |
| 6 | Future Work | 70 |
| | Bibliography | 72 |

List of Figures

| | | |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.1 | The layers of a bitewing dental X-ray image. | 7 |
| 2.2 | The gap between upper and lower jaw (thick horizontal line) and the gaps between teeth (thinner vertical lines). | 11 |
| 2.3 | The result of opening a dental X-ray using a circular structuring element. | 18 |
| 2.4 | The result of closing a dental X-ray using a circular structuring element. | 18 |
| 3.1 | Process overview. | 30 |
| 3.2 | The dental work extraction process. | 33 |
| 3.3 | On the left is a dental X-ray image. On the right is the result of gamma transformation. | 34 |
| 3.4 | On the left is the result of thresholding and edge detection of the gamma transformed image. On the right is the result of region growing of background. | 35 |
| 3.5 | Example of extracted dental work contours. | 36 |
| 3.6 | The enhancement process. | 37 |
| 3.7 | The result of top-hat filtering a dental X-ray using a circular structuring element. | 38 |
| 3.8 | The result of bottom-hat filtering a dental X-ray using a circular structuring element. | 38 |
| 3.9 | Top: Original jaw image. Bottom: Enhanced jaw image. | 40 |

LIST OF FIGURES

| | | |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.10 | An example of the path costs of an image. The brighter red areas are more costly than the darker blue areas. | 42 |
| 3.11 | The horizontal intensity sums normalized by width. | 44 |
| 3.12 | Teeth separation process. | 46 |
| 3.13 | Top: Lines drawn from start to end points. Bottom: Paths used to separate teeth. | 47 |
| 3.14 | The vertical intensity sums normalized by height. | 48 |
| 3.15 | Steps during contour extraction. | 49 |
| 3.16 | Example of extracted tooth contours. | 51 |
| 3.17 | Some calculated contour distances. These are before taking the logarithm of the value. | 53 |
| 4.1 | Correct jaw separation using path-based method. | 56 |
| 4.2 | Incorrect jaw separation using path-based method. | 57 |
| 4.3 | Correct jaw separation using integral-based method. | 58 |
| 4.4 | Incorrect jaw separation using integral-based method. Notice how the "strips" are misplaced. | 58 |
| 4.5 | Example of teeth separation[1] which is deemed successful. However it appears that some parts of the teeth near the edges have been cut off. | 60 |
| 4.6 | Results of different dental work feature weights. | 63 |
| 4.7 | This person was ranked 40th of 67 in the query results. | 66 |
| 4.8 | This person was ranked 46th of 67 in the query results. | 66 |
| 4.9 | This person was ranked 65th of 67 in the query results. | 67 |

List of Tables

| | | |
|-----|-----------------------------------------------------------------------------------------------------------|----|
| 4.1 | Results of teeth separation. | 59 |
| 4.2 | State-of-the-art results of teeth separation. | 61 |
| 4.3 | Results of using either tooth contour, dental work contour or both as features during ranking. | 62 |
| 4.4 | Overview of results reported in related research. | 64 |

Chapter 1

Introduction

The swift and correct identification of deceased persons can hold great value for those left behind. There are situations such as accidents, criminal acts or natural disasters, where victims can be left in a state that makes identification by documents or by visual means difficult. In some cases, family or friends of victims can be used for visual identification, but this is not always appropriate or even possible. Thus, biometric indicators such as fingerprints, DNA sampling and dental records become the only possibility for identification. Given ante-mortem (before death) and post-mortem (after death) dental X-ray images, an expert in forensic dentistry can compare the characteristics of the teeth and determine whether or not they are from the same person. However this is a time-consuming and focused effort that requires expert knowledge.

Automated identification systems, such as those based on dental X-ray records can reduce the work required for identification significantly. Traditionally the number of possible individuals has been reduced by using information found on the person or the geographical area the person was found. By using features inherent in the teeth of a person, it is possible to rank persons by the similarity of stored dental records to a query image. This allows for a much broader search

with less human intervention.

The Automated Dental Identification System (ADIS) was proposed at the turn of the millennium[2, 3] and much work has been done since then to create and improve solutions to the challenges[4] involved when building such a system. The work described here includes using both state-of-the-art and novel methods to create a new methodology for automated dental X-ray image identification, and testing it on a new data set.

The following contributions are made as part of the work:

- Identification of persons by a new methodology applied to a new data set.
- Introducing a novel method for bitewing dental X-ray separation.

The new separation method has the benefit of being able to recover from poor starting conditions, making it robust in cases where the teeth in the image are severely rotated or oddly shaped.

1.1 Problem Statement

The ultimate goal of this work is to create an ADIS using new methods and methods from related research. As part of the path towards this goal a few hypotheses will be tested. These can be listed as follows

1. *It is possible to identify the persons in the data set using only tooth and dental work contour information extracted from the dental X-ray images.*

The identification by dental X-ray component of this hypothesis has been verified by earlier research. The purpose of testing this hypothesis is to see whether it is true for this data set and to have a base to verify further hypotheses.

The particular combination of techniques and methods applied in this particular work has not been used to solve this problem before. This is a further reason to find how the whole system performs with regards to ranking and retrieval accuracy.

2. Pathfinding can be used effectively for separating teeth and jaws in dental X-ray images.

As will be seen in the related works section of this text, the separation of dental X-ray images into smaller images each containing a single tooth is a significant problem during feature extraction from such images. Separation of objects in images by path-finding algorithms is not a new idea but it has not been applied to dental X-ray images to the knowledge of this author. The work described here aims to find whether this is a viable solution to the dental X-ray separation problem in particular.

1.2 Motivation

For those whose lives are put on hold when a loved one goes missing identification can provide closure. It is also important in criminal cases to bring those responsible to justice and to finalize state records in cases where a person has been missing.

In areas affected by natural or man-made disasters, such as earthquakes, tsunamis or wars, authorities can be left with a large number of unidentified victims. It is important to keep records of those who have been affected, both for personal and public matters. However, the identification effort affects the recovery and healing process by imposing large amounts of work on the societies affected which takes up already sparse and limited resources. The use of automated identification can aid by freeing resources to be used elsewhere.

Biometric indicators such as fingerprints and DNA are subject to decay over

time, while teeth are some of the most resilient tissues of the human body and their placement means that they are in many cases intact even when other indicators are missing or hard to extract. Further, dental records are good candidates for identification due to the widespread use of X-ray images for normal dental health examinations. This means that there is a large potential store of existing images which, without underestimating political and privacy issues, could be used to form a large database for an identification system.

Though there is already a significant body of work in the field, these systems have not been widely adopted yet. Hopefully, by finding better methods and less complex methodologies for the systems, this can change.

1.3 Thesis Outline

Following the introduction, in chapter 2, is a section describing the previous work in the field. Section 2.3 afterwards gives an in-depth explanation of the techniques and methods that have been used. The combination of methods and techniques and the reasoning behind is presented in chapter 3. The testing of both the new method and the methodology as a whole, with discussion, can be found in chapter 4. After results are presented, a conclusion is made with a few thoughts on further work wrapping up the thesis.

1.4 Assumptions and Limitations

The data set available does not contain images taken post-mortem. However, the work described in this text assumes that the subject material of the images would remain unchanged to a large degree even after death. This assumption is backed up by an expert with experience in forensic dentistry.

The images are from adults only and as such the methods are not tested on

images of children's teeth, which can be very different and thus can be seen as a different problem in itself.

Most techniques and methods were implemented in the Java programming language and the test system as a whole was created as a collection of Java applications and libraries. The library BoofCV¹ was used for Canny Edge Detection and an implementation in Python² was used for the MorphACWE method. The SQLite-jdbc bindings³ were used for database handling. The results presented herein rely on the correctness of these libraries and implementations.

1.5 Definitions

Throughout the text some terms will be used repeatedly to refer to a specific concept. The most pertinent terms are explained in the following section.

An *image* is a two-dimensional grid of related points, called *pixels*. The images discussed in this text have a single channel (gray) with a scale from 0 to 255. Each pixel has such a value referred to as the pixel *intensity*. The perceived brightness of the pixels in the image is based on the intensity values of the pixels.

A *dental record* or just *record* is a single image or collection of images, extracted features and other information that is related to a persons identity. An example would be when discussing ADIS database records, which refers to stored information about a person and the ID of that person.

A *bitewing* dental X-ray image is an X-ray image showing the side of teeth, both upper and lower jaw, of a person. This is different from a periapical image, showing only part of one jaw, or a panoramic image showing the whole set of teeth, centered around the front teeth.

¹<http://www.boofcv.org/>

²<https://github.com/pmneila/morphsnakes>

³<https://bitbucket.org/xerial/sqlite-jdbc>

Chapter 2

Theoretical Background

This chapter is split into three parts. First, the bitewing dental X-ray image, the only type in the data set used in this work, is described as an introduction to the domain. Then, an overview of the field is presented to show the extent of previous research. Finally, the methods applied in the proposed solution are explained.

2.1 The Dental X-ray Image

Consider a bitewing dental X-ray image, see figure 2.1. It can be said to consist of three layers (from darkest to brightest): background, bone areas, and teeth. The teeth layer is the most prominent layer, often separated from the background by sharp boundaries around the contours of each tooth. The bone layer contains the regions at the lower and upper parts of the image, and between the roots of the teeth. This layer can be hard to separate from the teeth layer due to similar intensity and often unclear boundaries between the roots of the tooth and the jawbone. Finally, the background consists of the regions between teeth, where only flesh is left to absorb X-rays.




Figure 2.1: The layers of a bitewing dental X-ray image.

Dental works show up as part of a tooth but with a higher intensity than the surrounding area. These works can be fillings, implants, bridges etc. with fillings being by far the most common. The brightness of the dental work in the radiograph depends on the material used to create the dental work. Traditional metal-based materials such as amalgam appear very bright, while modern filling materials such as composite resin have a density much closer to that of the natural tooth and thus are barely brighter than the surrounding tooth.

2.2 State-of-the-Art

The field of forensic dentistry has been studied and applied for decades[5], but digital forensic dentistry has a much shorter history. Creating a fully automated computer-based system for human identification based on dental X-ray images has only been a focus of research for a little over a decade. Nassar et al.[2], Fahmy et al.[3] and Abdel-Mottaleb et al.[4] outline the general components of what they define as an Automated Dental Identification System (ADIS) and the most important challenges of developing such a system.

The initial research toward defining the ADIS[3, 4] states that the system should take as input a dental record consisting of one or more images and other pertinent information. The system should as output produce a candidate list of most likely candidates or some form of negative result if there are no acceptable matches.

The rest of this section is split into three separate parts, all describing research that can be related to the problems encountered in this work; enhancement, bitewing image separation, feature extraction and matching. The field has grown to include dental registration and tooth classification, topics that are not included in this section. After each topic there is a summary and conclusion of the most popular or effective methods.

2.2.1 Enhancement

Enhancement is a vital step in the processing of X-ray images due to often poor, uneven exposure and unclear region boundaries[6]. The different methodologies benefit from different enhancement methods and there are many that have been proposed. The purpose of enhancement of dental X-ray images is to enhance the contrast, particularly between the different layers of an image[3].

Morphological gray-scale reconstruction[7] consisting of bottom- and top-hat filtering of dental radiographs is used by several authors[1, 3, 6, 8, 9, 10, 11, 12] for enhancement of bitewing dental X-ray images. Such filters are implemented by using morphological operations, using the characteristics of the neighbours of a pixels to set the value of that pixel. The bright regions of the image are made more prominent by adding the result of a top-hat filtering operation. Conversely by subtracting the bottom-hat filtering result from the image, the darker parts are further removed. These operations normally lead to an image with increased contrast between the teeth-layer and other layers.

Lin & Lai et al. have introduced, and improved upon, a method of enhance-

ment based on adaptive local contrast stretching, adaptive morphology, and homomorphic filtering.[13, 14, 15]

Pushparaj et al.[16] suggest using Butterworth high-pass filter together with homomorphic filtering, a frequency-space solution to the problem. Shamsafar[17] uses gamma transformation to increase the difference between dark areas, such as background, and the bright dental work areas.

The method of enhancement needs to be chosen to benefit the other components in a system. For a system using image separation when extracting features, the morphological filtering method is reported to give good results.

2.2.2 Separation

In the interest of feature extraction it is useful to separate the X-ray image into smaller images such as those containing individual teeth.

Separating teeth from background and separating dental works from teeth both often involve some form of image segmentation, splitting the layers of the X-ray image. In the field of image processing the terms segmentation and separation are used to describe both the separation of objects in the image and the segmentation of layers. In this text, separation means splitting an image into several images, while segmentation will be used to describe the assignment of pixels into layers. Segmentation and separation are often both the most difficult and the most important tasks of a machine vision system[18, p.689].

Most of the separation techniques focusing on dental X-ray images have three steps: first segmenting the teeth layer from both the background and bone layers, then separating the upper and lower jaws, and finally each tooth.

Some authors[4, 19] propose a segmentation method based on masking the image by the result of iterative thresholding followed by adaptive thresholding. This method has been used extensively by other authors[8, 11, 14, 20, 21, 22, 23,

24]. The iterative thresholding step bases the initial threshold value on the average of the edge pixels detected by Canny edge detection. Since the edges are where the intensity is between two extremes, the background and the teeth/bone, this threshold results in few iterations with good separation of these two layers. The result from the iterative thresholding is used as a mask to remove pixels from the original image, this image is then adaptively thresholded.

Some have used only part of the method such as iterative thresholding[15], or adaptive thresholding[6, 10], after enhancement. A version of this method is used by Zhou & Abdel-Mottaleb[10] for segmenting both the teeth and the bones between teeth.

A texture-based method using clustering and a fuzzy inference system is described by Lin & Lai[13]. It is focused on using textural information to segment the image into background, teeth and bone areas.

The most popular method is based on iterative and adaptive thresholding, to mask away the pixels that are unwanted, thus segmenting the teeth from the background and bone.

Jaw and Teeth Separation

The separation into upper and lower jaw and further into individual teeth is a common step in feature extraction from dental X-ray images.

In a bitewing dental X-ray image, there is a natural separating gap or curve between the upper and lower teeth, shown in figure 2.2. Most techniques for separating the image try to discover this gap and assign pixels on either side to separate regions. The most obvious candidates for separating teeth are the low intensity background pixels between the teeth. The methods for separating teeth focus mostly on locating these lines and extracting the pixels between two such separating lines for creating the tooth region.




Figure 2.2: The gap between upper and lower jaw (thick horizontal line) and the gaps between teeth (thinner vertical lines).

Several methodologies for separating the jaws and teeth have been suggested. Some authors bypass the problem entirely by requiring a manually marked region around the desired tooth[25, 26], or suggest that the user add or delete lines to properly segment teeth[27].

A texture-based method using wavelet transform and morphology is introduced by Said et al.[9]. However, the explained results are not clear, with what seems to be a very forgiving measure of success, by counting as failure only when no teeth at all are properly separated in an image.

One of the earliest fully automatic and most cited approaches is that of integral intensity projection for both jaw and teeth separation[4]. Some have also used it to separate the bones between teeth[6, 10]. In the integral intensity projection method, the separation of upper and lower jaw is performed by locating the horizontal line with the lowest sum of intensities. In a proper bitewing X-ray image there is only one such line but it is not guaranteed to be straight. Jain & Chen[27] finds the gap valley by taking an estimate from the user and adjusting by a probability based on the intensity sum around the selected point. This is done strip-wise

to account for significant local intensity changes, ending with a spline representation that allows a curved separation of the jaws. Nomir & Abdel-Mottaleb[19] suggest a piecewise separation starting with the straight line of lowest intensity, splitting it into pieces and adjusting according to the lowest local sum of intensity. This process can have issues in cases where the dark area between the teeth is part of the local strip, and has lower intensity than the gap between teeth.

Integral intensity projection methods depend on enhancement and segmentation of the teeth layer in order to clearly make out the low-intensity areas between objects. To select the correct separating lines using this method, it is necessary to identify the projections that are between teeth. This leads to the problem of how to decide which of the lowest vertical sums are suitable for separating teeth and which are part of the same valley. Though there are many who cite this approach, few give good explanations of how to consistently find the correct valleys for the teeth separation lines. Abdel-Mottaleb et al.[4] use a threshold to find the valleys in the diagram of the vertical sums. Jain et al.[27] states that the valleys are used, but not specifically how these valleys are found, leading to the assumption that they are either detected by threshold or manually specified. Further, erroneous or missing lines are left up to the user to handle manually.

Abdel-Mottaleb & Zhou[6] use snakes for separating jaws and teeth, with initial lines determined by integral projection. Al-Sherif suggests the application of seam-carving to the problem of separating both jaws and teeth[24]. As the authors state, the method relies heavily on the enhancement step but this seems to be the case for other methods as well.

An interesting example of path-based jaw separation in a 3D mesh is presented by Le et al.[28]. Though this is not the inspiration for the approach introduced in this work, it is an example of path-finding applied to image and mesh splitting.

One important point to note is that when it comes to separation of dental X-ray images, integral intensity projection has been a popular technique to find the separating lines between jaws and teeth.

2.2.3 Feature Extraction and Matching

The images the system receives need to be stored in a format that allows for efficient and precise retrieval. This means extracting features that describe teeth and dental works accurately enough to compare the unique characteristics of each person.

A bitewing dental X-ray in grayscale contains two main types of information: shape and intensity. Shape is defined by the changes in intensity that form objects, while intensity is the changes within an object or other region of the image.

Features that have been used to identify persons from their dental X-rays fall into these two categories, with further separations within each category. The most frequently used shape-based features are tooth contours and dental work contours. Some authors[12, 14] have focused on numbering and classifying present and missing teeth. The work also introduced methods for detecting broken or missing teeth.

Abdel-Mottaleb et al.[4] introduce a method based on salient points taken from high curvature areas of the contour. These points are then related to other salient points by distance and the signature vector. A process of optimizing a transformation to find the best fit is also included. For matching the authors propose an equation focusing on the differences between vector components, with a division by the mean of one of the vectors.

In the work of Jain et al.[25, 27] a user specified region and point inside the tooth is used to find the contour of the tooth, by performing a radial scan and maximizing the difference between an inner intensity and outer intensity for each point along the contour. An additional step to align contours is also suggested, where a transformation is optimized for matching distance.

Other authors[8] use connected-component analysis, Fourier descriptors and Bayesian classification to classify teeth. Tooth contours are extracted using connected-

component analysis. Fourier descriptors are found for the contour result, with the first coefficient discarded and all others divided by the second coefficient to become scale and translation invariant. The spatial arrangement of the classified teeth is also used to improve results.

Gaussian mixture models have also been used[29] to find tooth contours. The authors then propose a new shape registration algorithm to find the correspondence between contours. This correspondence is then the main component of the matching algorithm. The article introduces new methods both for calculating distance between tooth contours and dental work contours.

Evolving contours have also been used as feature extraction methods. Zhou & Abdel-Mottaleb[6] apply the active contour method known as snakes to the contour extraction problem. They further separate these contours into crown and root sections for more accurate matching later. Jain & Chen[30] use a directional snake to extract the contour of teeth, after segmentation. Hofer & Marana[31] and Shamsafar[17] apply a similar snake algorithm to improving the contours of dental works. Similarly, Shah et al.[32] uses active contour without edges to find the contours of teeth. The authors claim to show better results when compared to the snakes method.

The problem of evaluating the similarity between two dental records can be quite difficult. It is a problem of comparing two sets of possibly unequal size with often different types of elements.

The work of Nassar & Ammar[33] shows that a neural network approach is viable. The approach uses two levels of comparison with different resolutions, which increases efficiency.

The Levenshtein distance is used for the matching stage by Hofer & Marana[31] with a new algorithm for calculating the edit distance between dental works. The method relies on location information in a panoramic radiograph and would require work to be applied for bitewing images.

Several authors[6, 15, 21] use the Hausdorff-distance for similarity measurement between contours. This approach seems to be the most effective reported. Weighted sum-of-squared distance has also been used[26]. Nomir & Abdel-Mottaleb[23] suggest using the hierarchical Chamfer distance, a multi-resolution technique, for matching images. The benefit of this method is more efficient reduction of the search space. The same authors also propose a novel matching distance[19], and a fusion of different matching techniques that apply several different feature representations and distance calculations[22]. The distance calculation methods include Euclidean and Absolute distance between point vectors[21].

Some of the most successful ranking methodologies[6, 15] use the Hausdorff distance or some variant of evolutionary contours such as active contour without edges or snakes.

2.3 Methods and Techniques

The following section provides detailed explanations of the methods that are used as part of the thesis work. Parts of, or the whole, section can be skipped if the methods are already familiar.

2.3.1 Gamma Transformation

It is often useful to increase the contrast between dark and bright pixels in an image. This can be done by exponentially transforming the intensity for each pixel. The gamma transformation is based on a power-law function with the exponent constant γ . The definition below is borrowed from Gonzales & Woods[18, p.111].

Definition 1 (Gamma Transformation).

Given an input r , the transformed value s can be found by the function:

$$s = cr^\gamma$$

where c and γ are both constants.

When the constant γ is less than zero the function will emphasize values that are towards the lower end of the intensity scale. If the value is above zero the transform will emphasize those values that are towards the upper end of the intensity scale.

When applying this transformation, with a γ value larger than 1, to a dental X-ray image it will bring out the bright dental work areas and suppress the darker areas. When γ is 1 the function becomes an identity function with the output equal to the input, given a c value of 1.

Digital images normally have a limited scale for the intensity values. This means that the result of a transformation such as the gamma transformation will need to be either cropped or scaled to fit between the maximum and minimum values after transformation.

2.3.2 Grayscale Morphology

Mathematical morphology, in relation to image processing, is a theory and set of techniques that define the results of applying a structuring element in particular ways to sets of image pixels. This section explains the concepts and techniques of grayscale morphology that are most relevant for understanding the applications within the solution described later.[18, p.665]

A typical morphological operation consists of two parts, the input image and a structuring element. The structuring element is most often a symmetrical shape. For the sake of implementation this element is often a two-dimensional array with

the element values set and the rest left empty.

The structuring element is passed over the image by centering around each pixel once and performing some specified operation which in most cases will change the intensity value of this pixel. The pass can be seen as a spatial convolution of the image by the structuring element.

All morphological operations are based on some property of the pixels contained within the structuring element. The most basic operation in grayscale morphology, erosion, assigns the smallest intensity value within the structuring element to the pixel at the center. This has the effect of shrinking all objects in the image, but will also darken it.

The opposite of erosion is the dilation operation. The same process is performed but it assigns the maximum intensity value of the pixels contained within the structuring element to the center pixel. This will grow objects out at every edge where there are changes in intensity, brightening the image at the same time.

By combining erosion and dilation the closing and opening morphological operations can be defined. To perform an opening of an image the erosion operation is followed by a dilation of the result. Since both operations are performed with the same structuring element, some parts of the edges around objects will be eroded away and not reconstructed by the dilation. This leads to a resulting image where edges are smoother, with widened gaps, but where small or complex objects may have been blurred or removed completely.




Figure 2.3: The result of opening a dental X-ray using a circular structuring element.

The closing of an image is the dilation by a structuring element followed by an erosion of the result by the same structuring element. This operation will result in an image where edges are smoothed and where small gaps and holes have been filled.




Figure 2.4: The result of closing a dental X-ray using a circular structuring element.

2.3.3 Adaptive Thresholding

A thresholding technique segments the pixels of an image into two different layers based on the intensity value and a specified threshold value. Adaptive window-based thresholding is used to threshold images where a global threshold is not appropriate, such as when background illumination differs relative to the foreground. The description of this technique is from the work of Nomir & Abdel-Mottaleb[19].

The technique begins by finding the average intensity value within a specified window around the pixel being thresholded. The threshold is then calculated as this average value multiplied with a coefficient. If the intensity of the pixel is greater than or equal to this threshold it is included in the foreground layer, if not it is set as part of the background layer.

Definition 2 (Adaptive Threshold).

Given an image intensity function $f(x, y)$, the adaptive threshold $t(x, y)$ for a pixel (x, y) is calculated as follows:

$$t(x, y) = \frac{1}{N} \sum_{i=y_0}^{y_1} \sum_{j=x_0}^{x_1} f(x+i, y+j)$$

where y_0, y_1 and x_0, x_1 are the start and end points of the window, relative to x and y , and where N is the number of non-zero pixels within that window.

Further the thresholding function is defined as:

$$g(x, y) = \begin{cases} 1 & : f(x, y) \geq c_0 t(x, y) \\ 0 & : f(x, y) < c_0 t(x, y) \end{cases}$$

where c_0 is a constant coefficient.

An output from $g(x, y)$ of 1 means the pixel is part of the foreground and a 0 means it is part of the background.

The coefficient to the threshold, c_0 , can be adjusted to move the threshold around the adaptive value. For instance, setting this coefficient less than 1 will include more pixels in the foreground, while setting it larger than 1 will include more in the background.

2.3.4 Region Growing

The region growing algorithm grows a region outward from a specified starting seed pixel by including pixels that are deemed similar by some criteria. These criteria can be anything, both local such as pixel intensity and color or regional such as total region size or distance from starting pixel.

The steps of the algorithm are as follows[18, p.764]:

1. Start by specifying the initial starting seed pixels. Each seed is added to a unique region set.
2. For each seed, add the neighbours that are not already part of a region set to the set of pixels to be considered.
3. Go through each pixel in the set of pixels to be considered and add it to the set of region pixels if it fulfills the criteria for that set. If the pixel was added to the region set, add the neighbours that are not already part of a region to the set of pixels to be considered.
4. Repeat the previous step until the set of pixels to be considered is empty.

When the algorithm finishes the result is a collection of region sets. Each set contains the pixels that are connected to the seed and fulfill the criteria for the set.

2.3.5 Boundary Following

The boundary following algorithm is used to find an ordered set of the pixels that lie at the boundary of an image region. It is a simple algorithm based on the Moore neighbourhood, the eight closest pixels, to the boundary pixels.

The principle of the algorithm is to make sure that the following algorithm always approaches boundary pixels from the outside. This is achieved by examining neighbours in a clockwise direction and always proceeding from pixels that are known to be part of the boundary.

Following the description by Gonzales & Woods [18, p.796] the algorithm works as follows:

1. Find the uppermost leftmost pixel in the image that is not part of the background, denoted b_0 . The westward neighbour of this pixel is then denoted c_0 . Since b_0 is the first leftmost pixel, the neighbour to the west of it must be a background pixel. Next, the first of the eight neighbours of b_0 , starting from c_0 and going in a clockwise direction, that is not part of the background is labelled b_1 . The background pixel immediately preceding b_1 is labelled c_1 .
2. Initialize the iterative steps by setting $b = b_1$ and $c = c_1$.
3. Find the eight neighbours $n_1, n_2, n_3, \dots, n_8$ of b , going clockwise around b starting at c . Find the first non-background pixel n_k for use in the next step.
4. Set $b = n_k$ and $c = n_{k-1}$.
5. Repeat steps 3 and 4 until $b = b_0$ and the next boundary pixel is the same as b_1 .

Perhaps the intuitive stopping condition would be when the starting pixel is reached again, this is insufficient however, since there are cases where the same

pixel lies at different points of the boundary. An additional condition that the outward direction from the starting pixel should be the same as during the start is therefore introduced by seeing that both b_0 and b_1 are the same.

When the algorithm completes, the result is an ordered list of pixels that lie at the boundary of the region, starting at b_0 and going all the way around to this point.

2.3.6 Integral Intensity Projection

A useful technique for finding objects and also gaps between objects in an image is integral intensity projection[19, 34]. The intensity projection is the sum of the intensity for all pixels encountered along a line going from one side of an image to the other.

For an image represented by the intensity function $f(x, y)$ the *vertical* intensity projection is defined as:

Definition 3 (Vertical Intensity Projection).

$$V(x) = \sum_{y=y_1}^{y_2} f(x, y)$$

Similarly the *horizontal* intensity projection is defined as:

Definition 4 (Horizontal Intensity Projection).

$$H(y) = \sum_{x=x_1}^{x_2} f(x, y)$$

The sum can be normalized by dividing the resulting sum by the maximum intensity value per pixel and by either height for vertical and width for horizontal.

Sometimes the goal is to find the lowest intensity sum at every point along an axis when allowing for rotation within some range. The projection is then

calculated as before, but the intensity for a point is the lowest sum found when rotating the image using the point as anchor. This is known as intensity projection with rotation.

Each direction is defined according to the axis the sum is calculated along. Thus for each point along an axis there will be a unique sum associated with that point. This is interesting because it means the diagram of the intensity projection for an axis can show both the extend of the object, along the axis the sums are calculated along, and the shape of the object along the other axis.

2.3.7 Dijkstra's Algorithm

Dijkstra's algorithm[35] is an understandable and elegant solution to the problem of finding the shortest, lowest-cost path between two vertices in a graph.

The principle of the algorithm is that in knowing the minimal path from a source vertex to a target vertex, one has to already know the shortest path to every vertex lying along said path. Thus by finding the minimal paths from P to each node R by increasing cost from P , until Q is reached, the shortest path from P to Q is found.

Begin with three sets of vertices, the set A of visited vertexes (those whose shortest path from P is known), B the set of all unvisited vertices that are connected to one of the vertices in A , and C the set containing the rest of the vertices. In addition to these sets, a map of the cost from P to each visited node also has to be kept.

1. When the algorithm begins, the cost of travel from P to each node is set to infinity and the cost of travel from P to P is set to 0. Before proceeding, P is removed from C and added to B .
2. The vertex in B with the lowest cost, denoted R , is the next to be considered. For all unvisited neighbours of R , calculate the updated cost of that vertex.

This calculation is the cost to reach R plus the cost from R to that neighbour. If the new cost for a neighbour is less than the existing one, that is if it is cheaper to go to this vertex via R , then it is updated to the new lower cost.

3. Place all unvisited neighbours of R , that are currently in set C , into set B . Finally, since R has now been visited, it is placed in set A .
4. If Q is in set A or if set C is empty, the algorithm stops. If not, continue from step 2.

When the algorithm stops, the cost associated with vertex Q is the cost of the lowest-cost path from P to Q . Finding the vertexes along this path can be achieved by storing the source vertex each time the cost of a vertex is updated. The list of vertices would then be obtained by going from Q , via each source vertex, until P is reached.

The cost of moving from one vertex to a connected vertex can be defined by a cost function. This function can depend on any available information, such as the intensity of pixels in an image.

2.3.8 Canny Edge Detection

Edge detectors are methods designed to find pixels in an image where the intensity changes abruptly. These are called edge pixels, and collections of such pixels will normally form edge segments.

One of the classic algorithms for edge detection in images is the one proposed by John Canny in his 1986 paper "*A Computational Approach to Edge Detection*"[36]. In the paper Canny proposes three criteria for evaluating an edge detectors performance. These are:

- Low error rate - All true edges in an image should be marked and no false edges should be marked.

- Well localized edge points - The marked points should lie as close to the true edge as possible.
- Single edge point response - The detector should not mark the same true edge point more than once.

The paper concludes that a good approximation to the optimal edge detector that satisfies the criteria mentioned in one dimension, is the first derivative of a Gaussian. The algorithm is based on this approximation, which is extended to two dimensional image processing by using a Gaussian smoothing filter together with the gradient magnitude and direction at each pixel.

The image processing version of the algorithm is based on the gradient, since this is the extension of the first derivative to multiple dimensions.

The Canny Edge detector algorithm can be summarized as follows[18, p.719]:

1. First, apply a Gaussian smoothing filter to all pixels.
2. Find the gradient magnitude, $M(x, y) = \sqrt{g_x^2 + g_y^2}$, and the gradient direction $\alpha(x, y) = \tan^{-1}\left(\frac{g_y}{g_x}\right)$ for every pixel in the image. High magnitude pixels lie between areas of different intensity, with ridges formed where the intensity changes gradually to high intensity and then to a lower intensity.
3. Remove "ridge" pixels from $M(x, y)$ to thin the edges. This is done by comparing the magnitude of a pixel and its two neighbours along the edge direction, as calculated based on the gradient direction $\alpha(x, y)$, of the pixel under consideration. If the magnitude is less than that of one of its neighbours, the pixel is set to zero. This is called nonmaxima suppression. The resulting image $g_N(x, y)$ contains some false edge pixels.
4. The false edge pixels are removed from $g_N(x, y)$ by thresholding. This is often done by hysteresis thresholding, combining the results of applying both a lower and an upper threshold to form the final result.

5. Apply some connectivity technique to set linking pixels to form connected edges in the image.

The result of the algorithm is an image where the only set pixels are those that are found to be part of an edge.

2.3.9 Active Contour Without Edges

Contour evolution methods have been applied for image segmentation problems for many years, with active contour without edges (ACWE)[37] being a prominent example of this type of method.

The principle of the active contour methods is that a curve can evolve based on differential operators to find a minimal solution. By expressing contour evolution as changes in partial differential equations, the contour can be evolved without the need for re-parametrization as would be the case with a more traditional curve. This allows the contour to split and merge into different contours if necessary to segment multiple objects in an image.

The contour is evolved based on an energy that consists of the sum of an external energy and an internal energy. The external energy can be thought of as an outside pressure that keeps the contour short and smooth, while the internal energy is an inside pressure that tries to grow the contour to include all parts of the object. Both of these terms are balanced by coefficients that determine the strengths of the pressures.

A feature of the active contour without edges method is that terms can be defined that maximise the difference in mean intensity between the object contained by the contour and the rest of the image[37]. Thus the contour can be made to both use the shape, intensity and gradients that describe the contained object.

For more formal, mathematical explanations of these methods see the original article by Chan & Vese[37] (ACWE) or the more recent article by Marquez-Neila

et al.[38] (MorphACWE).

2.3.10 Bi-Directional Hausdorff Distance

Hausdorff distance is a metric for calculating the distance between two sets of points. Given two such sets, A and B , the Hausdorff distance between these sets is the shortest distance to any point in B from the point in A furthest from any point in B . The distance is called bi-directional when the maximum of the distance both from A to B and from B to A is the result.

Using the definition in [39] this can be written as:

Definition 5 (Bi-Directional Hausdorff Distance).

$$H(A, B) = \max(h(A, B), h(B, A))$$

where the Hausdorff distance is

$$h(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\|$$

If $H(A, B)$ has value d then the distance from each point in A to some point in B must be less than or equal to this value. This means that the Hausdorff distance between two sets is a measure of the point of largest deviation.

Chapter 3

Approach

This chapter will explain the details of the components of the proposed system, from enhancement and image separation through feature extraction and matching. It does not delve too deeply into implementation details, instead focusing on the methodology in terms of algorithms and techniques.

3.1 System Overview

The process of going from a dental X-ray image to features that can be stored in a database and compared effectively consists of several steps. See figure 3.1 for an overview of the steps proposed in this text.

There are two possible tasks the system can perform when an image is input to it, either the image is to be used for training, i.e. to build the database, or it is to be used to find a matching record from the database.

In both these cases, the features need to be extracted from the image. The difference between the two tasks come at the end of the process when the system decides what to do with the features. When training the system, that is when

performing database entries, the final step is to store the extracted information in the database. When performing retrieval the final step is to compare the extracted feature values with those records already stored in the database.

The following list describes the different steps that the system goes through both when building the database and when matching a record. The term *separation* is used when referring to the act of splitting an image into smaller size images. This is not to be confused with the often used term segmentation which has been used to refer to the same thing, but is also a general term in image processing referring to the act of assigning pixels of an image into different regions often based on layering.

1. The user inputs the bitewing X-ray images into the system, specifying which side of the mouth each is from. The images need to be correctly rotated with lower jaw in the bottom part of the image. If training the system, a person ID must also be provided to be associated with the record.
2. Dental works are extracted from the image without separation. The contours of these are represented as point vectors and stored for later use.
3. The image is separated into two regions, one containing the upper jaw and one containing the lower jaw.
4. Each jaw image is enhanced and then separated into images each containing a single tooth. The tooth images are taken from the original unenhanced image.
5. Each tooth image is enhanced and the contour of the tooth is extracted. This contour is stored as a point vector for use in the next step.
6. If the system is training, the next step is to store each tooth and dental work contour vector in the database. If the system is retrieving then the next step is to compare each stored person record to the input record, tooth by tooth and dental work by dental work.

7. If the system is retrieving images, the next step is to produce a ranked list of person records based on the comparison scores from the previous step.

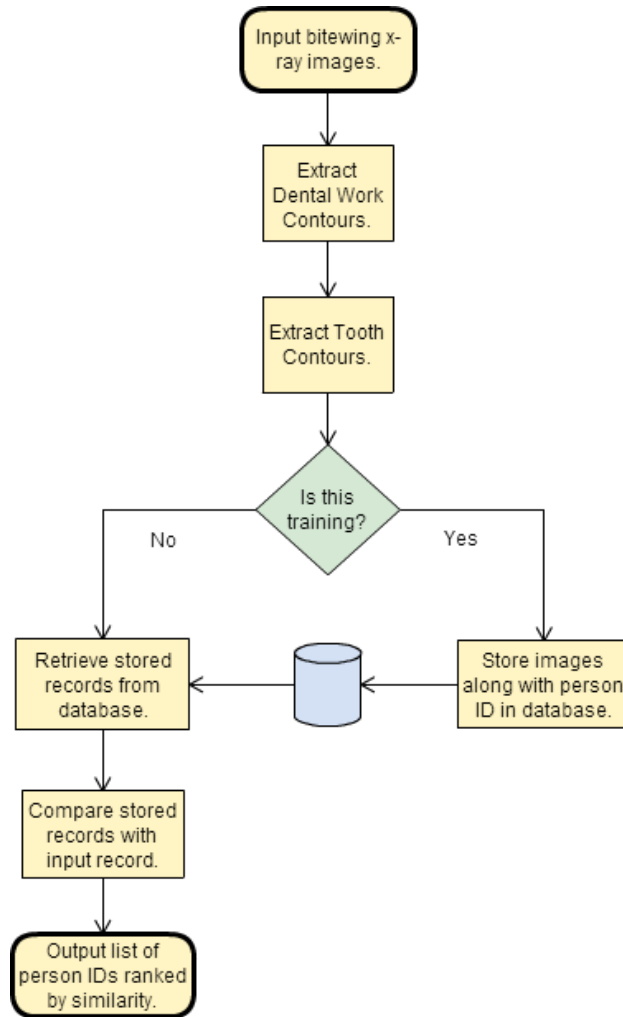


Figure 3.1: Process overview.

For a system of identification to be of use as a tool for forensics, it has to present the result to an expert who can verify the results and testify in a court of law. The value of the system is in reducing the search space for such experts in

coming to a justifiable conclusion based on the available data.

Since the outcome of the identification process often deals with matters of law and potentially significant emotional distress, it is very important to not allow any false positives. For this reason, the results of identification are presented as a percentage of inclusion within a certain number of highest ranked entries.

3.2 Feature Extraction

The content of a bitewing dental X-ray image includes information of shape, position and subject material density. There is no color information and depth information is based on the thickness and density of the materials affecting the attenuation of the X-rays.

The most obvious and most utilized features of these images are the shapes and sizes of teeth and dental works. As seen in related work, dental works are very important for identification whereas the shape and size of the teeth are becoming more and more relevant as dental health improves. Other possible features include missing teeth, distance between teeth, root canals and teeth implants.[5]

Features such as missing teeth, root canals or teeth implants all depend on an existence principle, which is hard to guarantee without a very robust feature extraction method. For instance, the feature extraction may be unable to identify a missing tooth if the center of the image is offset significantly or if a part of the image is very dark due to image quality. Similarly a root canal could be obscured in an image due to image quality and exposure. In either case, it is difficult to use the feature to increase classification confidence without ensuring that those that are extracted are very likely to be correct. Due to both time limitations and the described issues, these features are not proposed as part of this solution.

The contours of teeth, as seen in the state-of-the-art, have been used successfully as an identification feature. However, these are also subject to both gradual

change over time, due to aging and erosion, and sudden change if broken, pulled or restored.

The work described in this text focuses on tooth and dental work contours as features in the matching process. Each feature has a separate extraction process and methodology. Since dental works sometimes spread across teeth, they are extracted from the complete image without separation. Tooth contours, however, require separating the image into single tooth images before contour extraction.

3.2.1 Dental Work Contours

Dental works, when present, are the most recognizable and unique characteristics in a dental X-ray image[5]. This makes shape information from dental works a useful feature for identification.

The dental work extraction is focused on fillings that have a significant difference in brightness from the tooth. These are isolated from the rest of the image by the process outlined in figure 3.2.

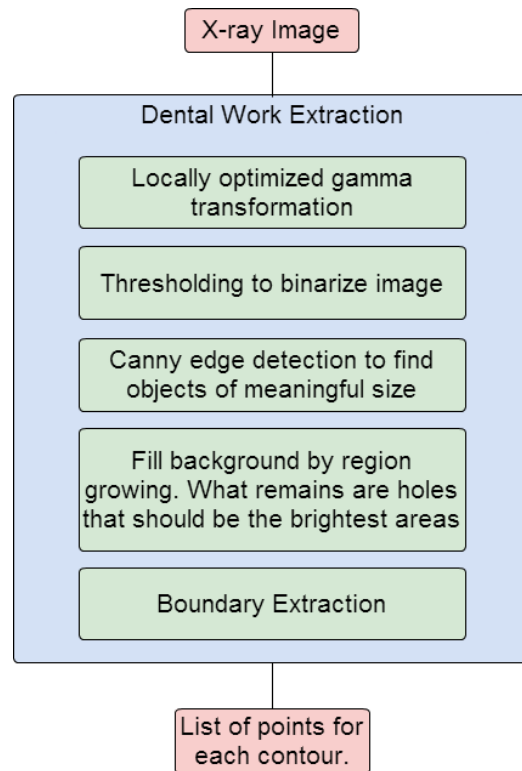


Figure 3.2: The dental work extraction process.

The first step, gamma transformation, is meant to increase this difference in brightness by using a power law transform to exponentially increase the intensity of each pixel. The work of Shamsafar [17] shows that this technique can be used to extract dental works. The work focuses on panoramic dental X-rays, but states that it should be possible to use for bitewing images.

Figure 3.3 shows the result of gamma transformation and a subsequent scaling of the result to fit the 256-value scale. Notice how the darker areas have been suppressed and the brighter areas, including dental works, have been retained.




Figure 3.3: On the left is a dental X-ray image. On the right is the result of gamma transformation.

An issue with Shamsafar's approach is that there is no dynamic selection of the gamma value, which will very likely lead to poor results on images with different bright to dark area ratios. The gamma value used during this transformation is highly dependent on the intensity of the image and the ratio of dental work pixels to other pixels. To select an appropriate gamma value for an image another step is introduced. This is a simple local optimization of a single variable, the gamma value, over a range of possible values. The best gamma value for an image is deemed to be the one that maximizes the average intensity of the pixels remaining after thresholding.

After the gamma transformation, the image is thresholded so the intensity of all pixels below the average intensity value are set to zero. The resulting image contains all those pixels that were part of a bright area, most often dental work, with the rest of the pixels at zero intensity.

The thresholded image may contain some noise and unconnected pixels. Some of this can be filtered out by using an edge detector such as the one proposed by Canny. Figure 3.4 shows the result of edge detection with subsequent morphological dilation. This step has the added benefit of inverting the values contained

inside dental works, which is necessary for the next steps. The pixels remaining in the image are all found to be part of an edge, excluding single pixel noise.




Figure 3.4: On the left is the result of thresholding and edge detection of the gamma transformed image. On the right is the result of region growing of background.

The next step is to find the pixels included in each dental work region. To ensure that only pixels within dental works are empty, the background region is grown from the darkest pixel found near the center of the image. In figure 3.4 the right-most image shows the result of region growing with the initial seed being the darkest pixel found in a window in the middle of the image. Passing over this image, from left to right and top to bottom, a new region is grown from each encountered zero-pixel. Since each region will grow to include all connected pixels, only one region will be grown for each set of connected pixels.

This process can lead to some false regions being included, most importantly edges of the image that have not been properly exposed by X-rays and very small regions due to high intensity noise. These are filtered out by two simple criteria: more than a certain percentage of the dental work must not lie at the edge of the image and the pixel count must be larger than some percentage of the total image pixel count.

For each accepted region, boundary extraction is used to obtain a list of points along the contour of the dental work. To simplify later comparison, the points are centered around the origin before storage. This is done by subtracting the x and y values of the region centroid from each point. The remaining list of points can be stored and matched numerically with other contours.




Figure 3.5: Example of extracted dental work contours.

Figure 3.5 shows a typical example of contours extracted from an image. Notice that the large round dental work does not have an extracted contour. In fact, the method is able to extract the contour, but it is incorrect due to the bright edges of the image being included. It is therefore excluded from the final record.

3.2.2 Enhancement

The tooth contour extraction process includes an enhancement step. This is designed to increase the contrast between dark areas, such as background and bone, and bright areas such as dental works and teeth. This is most important for the

intensity projection steps that are part of both jaw and teeth separation.

Figure 3.6 shows the different techniques, based on morphological filtering and adaptive thresholding, used during the enhancement process.

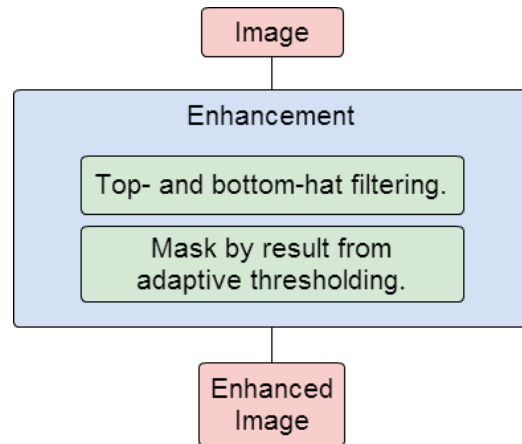


Figure 3.6: The enhancement process.

The morphological operations, top-hat and bottom-hat filtering are first used to obtain two images. Such filters are commonly used to enhance images by removing uneven brightness[7] and more specifically as a an enhancement method for dental X-rays as proposed by Zhou and Abdel-Mottaleb [6].

An image is top-hat filtered by subtracting the result of a morphological opening from the original image. What remains in each pixel then is the difference in intensity introduced by the opening operation. Since opening will spread intensity out from brighter areas, this operation will lead to an overall darkening of the image.




Figure 3.7: The result of top-hat filtering a dental X-ray using a circular structuring element.




Figure 3.8: The result of bottom-hat filtering a dental X-ray using a circular structuring element.

The result of top-hat filtering is added to the original image and the result of bottom-hat filtering is subtracted from the same image. Equation 3.1 shows this. Since the top-hat filter result has higher intensity values for pixels part of the bright areas, adding this will increase the brightness of these areas. The bottom-hat result, on the other hand, has higher intensity values for pixels that are part of dark areas. Subtracting this will reduce the brightness even further in these areas.

$$result = original + tophat - bottomhat \quad (3.1)$$

The morphological operations are performed using a circular structuring element due to it being a relatively easy shape to implement and the generally rounded features in dental X-ray images. The radius of the structuring element and the number of iterations of each morphological operation are both parameters that can be changed to yield different results to the separation step later on. Generally, increasing these will give a higher intensity result.

After morphological filtering the image appears much more evenly illuminated. A further enhancement can be made to benefit the next step, namely adaptive thresholding.

The purpose behind the thresholding step is to find those areas that already have low intensity and set their intensity to zero. This leads to more pronounced peaks and valleys in the intensity projection, a vital feature when locating points that lie in between teeth. An iterative thresholding step was considered, but it tends to converge on a threshold that keeps the bone pixels in the image and thus is not appropriate for this particular system.

If the threshold value is the same for all pixels, the technique is considered static. Static thresholding is not useful in cases where the image is unevenly illuminated or where the average brightness is very different from region to region. For dental X-ray images, thinner teeth appear darker in the image. The thickness of bone also varies and will result in different levels of intensity across the image. For these reasons, segmenting teeth from bone and background in bitewing X-ray images requires a locally adaptive thresholding method.

By considering only a window around the pixel being thresholded, the value can be adapted to the particular regional differences that characterize the pixels neighbourhood. Since the bones in one part of the image can have a higher absolute intensity than a tooth in another part of the image, the best thresholding

results are obtained with an adaptive window that is big enough to include both parts of the teeth around the bones and the background, but not so big as to include large differences in illumination. The threshold found is multiplied by a factor of 0.9 to avoid excluding too many pixels. This was suggested by Nomir & Abdel-Mottaleb[19].

The output from the adaptive thresholding step is used to mask the image obtained after the morphological filtering step. This means that all the pixels that have zero intensity after thresholding are set to zero in the output while all other pixels retain the intensity value as output from the morphological filtering step. This is to keep the differences in intensity between pixels that are necessary for the final steps.



Figure 3.9: Top: Original jaw image. Bottom: Enhanced jaw image.

Figure 3.9 shows a jaw image before and after the enhancement step. The contrast between dark and bright areas has been increased for the whole image.

Notice how the areas between the teeth and how darker parts have been threshold masked to zero intensity.

3.2.3 Tooth Contours

Recall the difference in X-ray image brightness between traditional metal-based and modern composite-resin based dental works. This is unfortunate for identification purposes since it means that as fillings that show brightly and in strong contrast to the tooth are being replaced with modern materials they become much harder to extract. For this reason and the fact that dental care has become so advanced that fewer and fewer people will have dental work performed in the future, tooth contours will become increasingly important[4].

In a bitewing dental X-ray image the outside of the row of teeth is almost orthogonal to the direction of the X-rays, due to the positioning of the bitepiece during image capture. This means that, for each side of the mouth, the same outline of the teeth will be shown whenever a bitewing X-ray is taken. The bitepiece allows relatively little rotation of the exposure plate which means that it is not necessary to account for skew or shear of the image.

The tooth consists of a crown and a root part, where the root is the part of the tooth inside the jaw bones. The contours of the teeth are the pixels in the image that are part of an intensity gradient change that lies at the border between background and crown and jawbone and root.

Path-Based Image Separation

The following section explains how path-finding has been applied in this work to solve the problem of automatically separating dental X-ray images. Lowest cost graph search is a well-studied problem with some very elegant algorithms, such as Dijkstra's, which are accessible and understandable.

An image can be expressed as a graph by adding a vertex for each pixel and then edges between a pixel vertex and each Moore neighbour. The problem of finding dark areas between teeth and jaws can then be formulated as a lowest cost graph search problem.

To define the lowest cost path in an image based on intensity, the cost function needs to take this pixel intensity into account. Thus, for this work the following definition of the cost function is used. Given a graph with a related image intensity function $f(x, y)$, the cost of travelling from the vertex at (x, y) to one of its neighbours is:

$$g(x, y) = 1 + \ln(1 + f(x + a, y + b)) \quad (3.2)$$

In equation 3.2 a and b have values 0, 1 or -1, depending on the direction to the neighbouring pixel. Figure 3.10 shows a visualization of the different path-costs that are calculated for pixels as part of a graph.




Figure 3.10: An example of the path costs of an image. The brighter red areas are more costly than the darker blue areas.

The method can be configured to behave in a specific way by tweaking the path-cost function. The most basic cost would be to use the grayscale intensity value of the pixel. This will produce paths that are not limited by length, and thus

will not tend to be straight. They will however seek to lie in the lowest intensity areas of the image.

Basing the path cost only on intensity can cause problems in instances where there is a dark empty border around the image. To solve this issue, a constant term is added to the cost function. By including a constant term the algorithm will prefer shorter paths over longer. This may sometimes have the adverse effect of causing it to cut through "rough terrain", such as the inside of a tooth if the separation between the jaws has a significant curvature.

Notice the use of the natural logarithm. This function has the desirable property that it increases sharply and levels out for larger values. Translated to the path-cost problem this means that very bright areas, such as dental works, do not have too large an impact on the path selection. Since these areas normally lie at the top of the crown the path would sometimes prefer to pass through the root of the tooth instead of passing through a high intensity bridge at the crown, if the cost was linear. Both the square root and the logarithm were tested, with very similar results, but the natural logarithm was chosen due to slower growth putting more emphasis on values near the lower end of the intensity scale.

The path-based method uses the intensity value and does not directly benefit from thresholding, unlike some other methods like integral intensity projection. This is because it takes into account the intensity value in its cost calculations and will use differences in intensity to find the lowest intensity path. When thresholding, all pixels over the threshold are given the maximum intensity value, which in many cases can mean that bones between teeth are inadvertently given the same path-cost as teeth.

Jaw Separation

The jaw separation problem is solved by path-finding by locating a start and an end point that both lie at the position of the gap valley, at the left and the right

side of the image. Horizontal intensity projection is used to find the darkest area in both the left side and the right side of the image.

By only calculating part of the width in towards the center, the method can allow for images where the separating line is severely rotated. In these cases, the horizontal intensity sum for a point can be lower than the gap valley.

Since the bitewing dental X-ray always shows two rows of teeth, the upper and lower jaw, the separating line between these rows will consist of background pixels. The lowest-cost path between the start and end point will pass through these pixels in almost all cases where the image is not severely underexposed or the person is missing the majority of the teeth in one or both jaws.




Figure 3.11: The horizontal intensity sums normalized by width.

The path-based method does benefit from enhancement that increases contrast

between background and teeth, but tests have shown that for jaw separation, the method performs as well on enhanced as on raw images for this data set. This moves the need for enhancement further down the pipeline, to smaller images, a possible performance benefit.

Teeth Separation

After obtaining an image containing either the lower jaw or the upper jaw the next step is to separate it into a set of images each containing a single tooth.

The teeth separation process has much in common with the jaw separation process. The most significant difference however is that there is more than one separating line, which are less pronounced in the intensity projection diagram.

The process begins by enhancing the jaw images output by the previous jaw separation process. The path-based method of image separation is used also during this phase of feature extraction. Dark areas between teeth are then located using intensity projection and by using these locations the separating lines are found by the path-based method. These lowest-cost paths going between teeth, from left to right, are used to separate the image by using the maximum and minimum points of each path and the path to its right.

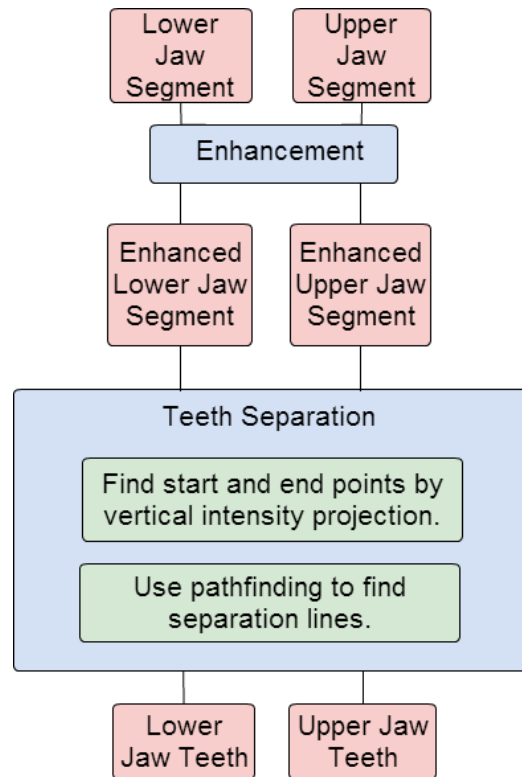


Figure 3.12: Teeth separation process.

Each separating path is found by first obtaining a start and end point by vertical intensity projection with rotation, and then finding the lowest-cost path between these points.




Figure 3.13: Top: Lines drawn from start to end points. Bottom: Paths used to separate teeth.

The intensity projection diagram seen in figure 3.14 is of an enhanced jaw segment. The intensity sums have been normalized to remove differences due to image height and intensity scale. This is useful when comparing images of different height or width. It is also useful in cases where the relative shape differences within the image are more important than the absolute shape differences. There are four peaks, and four valleys which coincide with the four teeth present in the image.




Figure 3.14: The vertical intensity sums normalized by height.

In the case at hand, when the goal is to find dark areas separating objects, the difference between peaks and valleys in the intensity diagram is more important than the height of the peaks. To find the starting points to use for path-finding, those valleys that are more than a specified percentage lower than the height of the previous peak are deemed to be between teeth.

To avoid multiple points in the same valley a minimum distance between points, in effect a minimum tooth width, is used as a parameter. Since there are often partial teeth at the ends of the image, a starting point at each side of the image is also included.

The gaps between teeth are in many cases not vertical, due to both widening at the crown, natural rotation of the tooth or image rotation. For this reason, the

vertical intensity projection finds the lowest sum with rotation. The angle of the lowest sum is therefore used to decide where the end point should be placed, aiming at that angle from the start point and ending where this line reaches the other side of the image. The upper half of figure 3.13 illustrates this.

Tooth Contour Extraction

After jaw and teeth separation, the result is a collection of separate images each with a centered tooth. These are not useful as features until the contour if the tooth can be represented numerically. This is the purpose of the final part of the tooth contour extraction process. Figure 3.15 shows the different parts of this step.

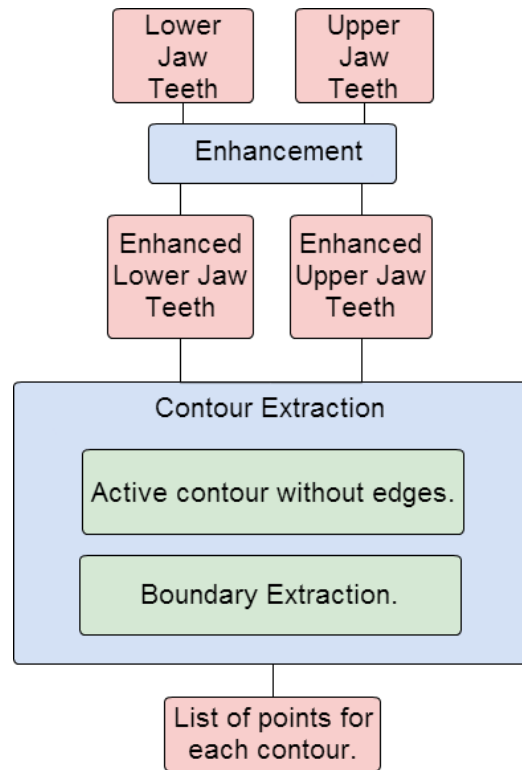


Figure 3.15: Steps during contour extraction.

First, the pixels contained within the tooth shape are found by the use of the morphological active contour without edges (MorphACWE) method[38]. Active Contour Without Edges (ACWE) has been applied to the problem before[32] with good results. MorphACWE is an alternative definition of the ACWE method which is faster and simpler due to the exchange of numerical methods with morphological operators.

By initializing the active contour as a circle in the middle of the image it is guaranteed to begin in a high intensity area, given a correctly separated tooth image. This initialization has been shown to work for this type of problem[32]. From there, the active contour can evolve according to the external and internal forces determined by set parameters and the image gradient.

The MorphACWE energy functional relies on both the gradient points at the contour and the intensity and size of the area inside the contour. Some teeth contain dental works whose brightness is significantly different from the surrounding tooth material. This can lead to the active contour coming to rest at these gradient points before reaching the actual tooth contour due to the difference in average intensity. The opposite problem of bone material having a density very similar to the root can cause the active contour to spread into the surrounding bone and further into visible parts of the neighbouring tooth.

These problems can be avoided by using the same enhancement as for jaw separation, albeit with different parameter values. Using multiple iterations of the morphological filtering step and a larger structuring element, the intensity of the material inside the tooth can be increased to lie near the top of the intensity scale along with the dental work areas. When the tooth image is enhanced such that the brightness of pulp, enamel and dental work are all pushed to the top of the intensity scale the internal average intensity is much more even also when the area includes the lower areas of the tooth and dental works. This will help to lead the active contour out towards the actual tooth contour. Further, when the threshold masking of the enhancement removes some of the bone material at the sides of

the roots, the gradient becomes more pronounced which strengthens the term of the active contour that relies on the gradient.

The output of the active contour without edges algorithm is an image where all the pixels outside the contour are set to zero and those inside are set to one. From there, the same boundary extraction process as used during dental work extraction can be used to find an ordered list of points along the contour.




Figure 3.16: Example of extracted tooth contours.

In figure 3.16 the contours extracted from an image can be seen. The contours are representative for most of the teeth, but at the edges of the crown, where there are differences in depth and thickness, the method sometimes fails to find the true contours. The same can be seen at the part of the roots within the jawbone. The sometimes jagged edges are not the true contours of the teeth. However, though there is room for improvement, the method can successfully extract contours that can be used as a feature in the matching step.

3.3 Matching

The most important utility of the system is identification of a person by the ranking of stored record against a presented dental X-ray image. The storage of extracted feature values is necessary to facilitate this. The specifics of the database will not be discussed in this text, though it is worth noting that any stored features need to somehow be associated with the person ID in order to be used for identification. Information regarding the side of the mouth the image is from is also stored in the database.

The matching phase is when the stored database records are compared with the query images. These records consist of point-vectors for each contour extracted from an image. Producing an ordered list of records based on similarity between images requires a method for comparing the point-vectors. For this purpose, bi-directional Hausdorff distance is used for comparing both dental work contours and tooth contours. This similarity measure has been shown to be well-suited for use in image comparison[39] and has been used for dental contour matching[6].

When comparing two sets of points belonging to boundaries of objects, the sequence of the points is irrelevant when using Hausdorff since the points are implicitly ordered by distance during the calculation. This is very beneficial when it is difficult to align starting points of boundaries, as is the case with tooth and dental work contours.

A dental record consists of many teeth and dental works, each with a contour. To find how similar two dental records are, it is necessary to aggregate the individual contour similarity values to find a total similarity value.

For each contour in the query record, a ranked list of the distance to each relevant contour and the related person ID is created. A person will then be given a total similarity score based on the sum of the distances of their highest ranked contour in these lists. Some authors[20, 22] have used majority vote based on the number of teeth from one person, but this was found to rely too heavily on the

number of and the quality of correctly extracted contours.

Since the user specifies the side of the mouth the image is taken on, the record can be compared only with the same side records from the database. This halves the necessary number of comparisons. Further, the system assumes that any input image is correctly rotated with the lower jaw in the bottom and the upper jaw in the top. This fact is also used to reduce the necessary similarity calculations by only calculating similarity between teeth from the same jaw.




Figure 3.17: Some calculated contour distances. These are before taking the logarithm of the value.

One issue with the matching algorithm is the way it punishes persons that have teeth or dental works that do not match well with anything in the database. This would be fine if the feature extraction method was perfect and the dissimilarity was due to actual differences and not due to extraction difficulties. In that case it would be possible to say that a tooth or dental work with no acceptable match was new and could be discarded during the dissimilarity accumulation phase. It does, however, happen that the database record or the matching record is incomplete due to a failure of the feature extraction method.

As mentioned earlier, some contours are not correct representations of the source tooth. If contours in the database record have been extracted correctly,

then this contour will be found to be very dissimilar to the correct contour. The same is true if the database contour is incorrect. The score for any particular person is the accumulated score of the most similar teeth and dental works against the query image. A single poorly extracted contour therefore has the potential to unfairly punish a record that is otherwise very similar.

A small change to the dissimilarity calculation is therefore proposed to reduce the impact of very dissimilar contours. The same principle as in the path-cost function can be used here, taking the natural logarithm of the dissimilarity between contours, before adding it to the total sum for a record. The logarithmic increase in distance puts less emphasis on those contours that are very dissimilar.

Chapter 4

Results and Discussion

In this chapter the results of testing the explained approach on the data set are presented. First, the path-based method is tested for jaw and teeth separation then the entire system is used to rank images for identification.

4.1 Separation

This work introduces a new method for separating jaws and teeth in dental X-ray images. This section outlines the results from experiments devised to test the efficacy of this method, and verify or dispose hypothesis 2 from section 1.1.

4.1.1 Jaw Separation

To test the performance of the path-based method for the problem of jaw separation it was used to separate 100 images. Two images, one from each side of the mouth of each person, were selected for use. To avoid bias due to X-ray equipment settings and hardware, the images were selected from random dates so that

some persons had older records and some had newer.

It was not possible within the time limits to implement and test every state-of-the-art method on the particular data set used in this text. Instead, the most popular method[4], based on integral intensity projection, was tested on the same images for comparison.

The results of the jaw separation of the described test images using the path-based method were very good. Of 100 images, 98 were correctly separated and 2 were incorrectly separated. See figures 4.1 and 4.2 for examples of correct and incorrect jaw separation using the path-based method.




Figure 4.1: Correct jaw separation using path-based method.

In the cases where the method fails it can be attributed to missing center teeth or very similar intensity where the crowns of the upper and lower teeth meet, making it very hard to detect the gap. Either problem can cause a misplaced start or end point for the path, because there is a lower intensity horizontal sum in an area that is not part of the actual gap.




Figure 4.2: Incorrect jaw separation using path-based method.

A possible solution to this would be to select a lowest cost path from several created from the lowest starting points. If the starting point is on the wrong side of the teeth, the path will have to cross a higher intensity area at some point, which will increase the cost significantly.

The intensity projection method was able to separate 92 of the 100 images correctly. It fails to segment some images due to the darkest part of the strip being found inside of a tooth. This can happen when the jaws are at an angle in the image or when the crowns of the teeth are touching. The strip technique with locally adjusted height suffers from the problem that it is possible to find a lower intensity line above the crown of the teeth in some cases.

Although the intensity projection method allows the separation to be curved, it is limited, and in some cases where teeth are missing it can be initialized at the wrong point and fail to separate properly due to this limitation. If thresholding is performed on the image before using the strip method, this problem can be avoided when the tooth is darker, but not when there is a gap between teeth.




Figure 4.3: Correct jaw separation using integral-based method.




Figure 4.4: Incorrect jaw separation using integral-based method. Notice how the "strips" are misplaced.

Notice the difference between figure 4.1 and figure 4.4. The figures show the same images split by both methods. These images illustrate the point that the strip-based integral intensity method can fail when the image is significantly rotated. The path-based method performs better when separating such images since the path can diverge from the gap valley and since it can have two unrelated start and end points (two valleys).

4.1.2 Teeth Separation

Similarly to the jaw separation problem, an experiment was designed to test the performance of the path-based teeth separation method. The data set for this experiment was selected in the same way as for the jaw separation experiment, with randomly chosen images, two per person. Unsuccessfully separated images from the jaw separation step were not included.

The evaluation of the results is based on the following definition of a correctly separated tooth image. Each tooth, including those partially visible at the edges of the image, must be in its own image with the complete tooth in the center of the image.

The results of the experiments, seen in table 4.1 show that the path-based method outperforms the integral intensity method for this data set. There are several cases where both methods fail due to the gaps between teeth being difficult to find in the vertical intensity sum diagram.

| | Correctly Separated Upper | Correctly Separated Lower |
|--------------------------------------|--------------------------------------|--------------------------------------|
| Integral Intensity Method [4] | 270/336 - 80.3% | 248/306 - 81.3% |
| Path-based Method | 270/336 - 89.3% | 248/306 - 88.2% |

Table 4.1: Results of teeth separation.

The methods differ most when individual teeth or the whole jaw is rotated in the image. In this case, the path-based method has an advantage by being able to use two unrelated start and end points. When either of these are wrongly placed, inside a tooth, the resulting path will necessarily cut through the tooth to find the darker, lower-cost path between the teeth. In these cases, the resulting separation will contain the crown and most of the tooth, which means the method is able to recover to some degree even with poor initialization. This is different from straight line approaches[4, 19] reported which can not recover from a situation where the gaps between teeth are not clear from the intensity diagram.

CHAPTER 4. RESULTS AND DISCUSSION

Comparing results with those reported by other authors is difficult, since different data sets are used. It may also be unfair since teeth separation evaluation is largely a manual effort relying on different criteria for different people. For instance, in figure 4.5 the authors show examples of what are supposed successful separations, but the images show that edges of some teeth have been cut off. This reveals both different criteria for acceptable errors while separating teeth and precisely the issues the path-based method can help solve with its adaptive capability.




Figure 4.5: Example of teeth separation[1] which is deemed successful. However it appears that some parts of the teeth near the edges have been cut off.

A shallow comparison can be made, by looking at some reported results, to determine whether or not the path-based approach is viable. In table 4.2 those authors who have reported teeth separation ratios are listed. Taking these numbers at face value, the path-based results of 89.3% and 88.2% for upper and lower jaw, respectively, can be considered to be favourably comparable to the state-of-the-art.

Notice also that Abdel-Mottaleb et al.[4] report a higher percentage of success while using the integral intensity method than was found for this data set. This can be explained due to reasons such as evaluation criteria differences or data set differences.

| | Correctly Separated Upper | Correctly Separated Lower |
|---------------------------------------|--------------------------------------|--------------------------------------|
| Abdel-Mottaleb et al. 2003 [4] | 169/195 - 85% | 149/181 - 81% |
| Nomir et al. 2005 [19] | 329/391 - 84% | 293/361 - 81% |
| Al-Sherif 2012 [24] | 1604/1833 - 87.5% | 1422/1692 - 84% |

Table 4.2: State-of-the-art results of teeth separation.

The results of experiments show that this method has proven to be competitive and it succeeds in enough cases to allow for further testing of actual matching and identification experiments.

4.2 Image Retrieval and Ranking

The system will produce a ranked list of all relevant images in the database by their distance to the query image. A relevant image is considered to be one that is of the same side of the mouth. Measuring the performance of the identification step is done by finding the first occurrence of an image from the same person as the query image.

The retrieval accuracy is always relative to the top-n retrievals. If the accuracy is referring to top-1, only the first image is considered. This is the absolute accuracy of the system and the optimal system would have a top-1 accuracy of 100%. If, however, the accuracy refers to top-5 a correct retrieval will rank the correct image within the first five. The reason for reporting the results in this way is to give an idea of how many images an expert would have to examine, on average, before finding a match. This is a good measure of the time-saving value the au-

CHAPTER 4. RESULTS AND DISCUSSION

tomated dental identification system can have as a tool for forensic dentists and other investigators.

Experiments were run to see how the system performs for identification purposes. All experiments, unless stated otherwise, were performed on a data set of two, older, images per person in the database and two, newer, images for query. There were a total of 67 persons in the data set, and all were part of the queries. This makes a total of 134 queries; 67 right images and 67 left images. The older images were considered ante-mortem images, and their features were stored in the database. The newer images were considered post-mortem and used as query images. This method of splitting the data set has been done before[31] and consultations with an expert led to the conclusion that this would be a suitable data set for testing.

The best results are obtained when using both features for matching. This is in large part due to those images that do not have any dental works, or where the method fails to extract the dental works. These images are impossible to rank correctly when using only dental work contours, but by using tooth contours a more representative ranking based on more information can be achieved.

| | Dental Work Contours Only | Tooth Contours Only | Dental Work and Tooth Contours |
|---------------|----------------------------------|----------------------------|---------------------------------------|
| Top-1 | 77.61% | 50.51% | 86.56% |
| Top-5 | 88.05% | 65.67% | 94.03% |
| Top-10 | 91.79% | 72.38% | 94.03% |
| Top-20 | 93.28% | 80.59% | 97.01% |

Table 4.3: Results of using either tooth contour, dental work contour or both as features during ranking.

Table 4.3 shows the results when using both features, and also demonstrates that dental work contours alone are better than tooth contours alone. This seems to corroborate that dental works are more unique than tooth contours. Using tooth contours, which are in some cases incorrectly extracted, can introduce noise to the

distance calculations but the results show that this is outweighed by the benefit of having more information to match.

The best results are obtained when weighting the dental works during distance calculation. The graph in figure 4.6 shows the effect of using different weight factor values for the dental work contour distance. The accuracy is raised by increasing the weight of the dental work feature, and thus the importance of the dental work contour distance values. The best overall result is obtained when the dental work contours are weighted twice as much as the tooth contours.




Figure 4.6: Results of different dental work feature weights.

There can be several reasons why the accuracy increases when dental work contours are weighted twice as much as tooth contours. It can be seen as an indication that dental work contours are more unique and that the contours are

CHAPTER 4. RESULTS AND DISCUSSION

more precisely extracted. It could be because similar tooth contours from other persons introduce noise. It can also be that tooth contours are more prone to disappear either because teeth are missing or because of image offset.

Comparison with other reported methodologies is, as with teeth separation, difficult to do fairly without implementing and testing them on the same data set. Such a data set was not easily obtained and the implementation of other methodologies was outside the scope of the work.

The comparison of the accuracies seen in table 4.3 with reported results should be done with care due to the differences in data sets and number of queries. A rough comparison is made, however, to see how the proposed methodology performs relative to the state-of-the-art.

Table 4.4 shows some of the reported results of dental X-ray image ranking. The PM/AM image ratio refers to the number of query images to the total number of images in the database. Most authors use 30-50 query images while testing.

| | Top-1 Correct | Top-1 Accuracy | Top-5 Correct | Top-5 Accuracy | PM/AM Image Ratio |
|----------------------------------------|----------------------|-----------------------|----------------------|-----------------------|--------------------------|
| Jain et al. 2004 [25] | 25/38 | 65.8% | 34/38 | 89.5% | 38/130 |
| Nomir & Abdel-Mottaleb [19] | 32/43 | 74.4% | 38/43 | 88.4% | 43/117 |
| Zhou & Abdel-Mottaleb [6] | 33/40 | 82.5% | 38/40 | 95% | 40/102 |
| Nomir & Abdel-Mottaleb [20] | 42/50 | 84% | 50/50 | 100% | 50/162 |
| Hofer & Marana [31] | 19/22 | 86.3% | 20/22 | 90.9% | 22/46 |
| Nomir & Abdel-Mottaleb [21] | 43/50 | 86% | 50/50 | 100% | 50/162 |
| Nomir & Abdel-Mottaleb [23] | 42/50 | 84% | 50/50 | 100% | 50/187 |
| Lin et al. [15] | 33/35 | 94.3% | 35/35 | 100% | 35/93 |

Table 4.4: Overview of results reported in related research.

The results obtained in this work can be favourably compared to those reported in earlier research, even against more complex methodologies. It is worth noting that the best reported result[15] introduces a contour alignment method to significantly improve results. This could most likely be used to improve the

proposed methodology as well.

Where the proposed methodology seems to struggle is the inclusion of all queries in the top results. Several authors report 100% inclusion in the top-5 retrievals, while the proposed methodology only manages to reach 94%. Without using the same data set it is not possible to conclude whether this actually is an inherent weakness in the methodology or due to differences in the number of difficult to rank images. What follows is a more detailed analysis of these particular images and some discussion of the reasons why the system fails to retrieve them.

Detailed Analysis

Some of the images which the system failed to rank are analyzed in this section. The purpose of the analysis is to identify the cases of failure and the reasons why it is so. The figures show database image and query image for a person. The rightmost image is the query image, which is always newer than the database image.

By manually analyzing the results for each query some categories of failure were identified. The identified categories were: poor image quality, translation offset, and subject matter change.

In figure 4.7 an example of subject matter change as well as translation offset can be seen. This image is ranked very low due to the two smallest dental work contours in the leftmost image being very different from the merged dental work in the rightmost image. In addition to the difference, the work has been fused in the more recent image so both the smaller contours are scored poorly against the large new contour.




Figure 4.7: This person was ranked 40th of 67 in the query results.

Figure 4.8 illustrates a person with no extractable dental work contours and difficult to extract tooth contours in the rightmost image. Additionally the teeth in one image are offset relative to the other, which causes different contours to be part of the records. Notice in particular the large tooth to the left in the rightmost image. When the images are compared these contours will penalize since there are no extracted contours that can be matched well.




Figure 4.8: This person was ranked 46th of 67 in the query results.

The final example, in figure 4.9 shows an example of poor image quality and a difficult subject due to few available features and significant change between the images. Though the person definitely has characteristic teeth, the method of aggregating scores used in this system will cause images such as these to be

penalized more than they should.



Figure 4.9: This person was ranked 65th of 67 in the query results.

The examples above show that the system is not at the moment very change-resistant. A possible solution to these issues in the future can be to add missing tooth detection, tooth classification and registration to the system[14]. This has both the advantage of reducing the necessary calculations for matching and allowing the system to exclude contours that are not present in one image during comparison to give a more representative similarity score.

A Note on Scalability

Though the results show that the system is capable of identifying a person by the use of dental X-ray images, it also raises some questions of whether it is possible to increase the size of the database and still see the same results.

The difference in similarity in the ranked list is only 12.6% between the two highest ranked images and only increases to 18.9% between the highest ranked and the 20th highest ranked image. This small difference suggests that the results will not scale well to a larger database. If the distances and rate of occurrences are representative of dental X-rays in general, that would mean that a top-10 result could scale to a result in the thousands and beyond very quickly.

It is also hard to argue that a database of only 67 possible persons, though not among the smallest reported in the research (see table 4.4), can be truly representative of an actual automated dental identification system database. If such a system is to be used regionally or nationally, perhaps even globally, then the features extracted for a particular individual must truly be unique.

Since the problem is identification, each possible person represents a new class that must be distinguished from all other classes by its features. This means that a good match should be significantly better than all other matches in order to allow scalability. This is not the case with the current extracted features and such an implementation would require further refinement of the extracted contours, a better alignment algorithm and also new features.

Chapter 5

Conclusion

The problem of human identification by using dental X-ray images has been studied in this work. Particularly a novel method for the problem of separating a dental X-ray image into individual teeth has been introduced as part of verifying hypothesis 2 from section 1.1, stating that a path-based method can be used as a dental X-ray separation method.

The new dental X-ray separation method is able to correctly separate 88.7% of the teeth in tests. The methodology for image ranking is able to rank the correct person first in 86% of the cases when tested on a database of 67 persons. This shows the veracity of hypothesis 1 from section 1.1. However, more work is needed show how the performance of a system using this methodology is affected when the number of individuals in the database increases.

The work has been successful in verifying the two initial hypotheses leading to the conclusion that using the new methodology, including the path-based separation method, is a viable option when building an automated dental identification system.

Chapter 6

Future Work

This chapter suggests both possible verifications of and improvements to the proposed system and some directions of further research.

The work described in this text was all performed with post-mortem identification in mind. Unfortunately, the data set only contains ante-mortem images. In this light, it would be interesting to see how the method performs on actual post-mortem dental X-ray images. Testing the system on some of the other data sets used in the field would also be valuable due to the comparisons that could be made with existing research.

The introduced method for jaw and tooth separation should be tested more thoroughly against other methods as is done by Al-Sherif et al.[24], preferably all on the same data set.

There are several improvements that can be made to the matching step. At the moment, a lot of essentially redundant calculations are made to compare teeth and dental work. This is because there is currently no way to know whether it makes sense to look for similarity between two teeth. There is, however, exclusive information in the image which could be used to cut the search space into much smaller pieces. For instance, as reported in some of the related work it is possible

to classify premolars and molars, meaning molars could be compared only to other molars and premolars only to other premolars. As some writers have shown, it is possible to perform a complete dental registration of all the teeth visible. With a robust dental registration step each teeth need only be compared with same position teeth in the database.[12, 24, 40]

Sometimes a tooth may be touching another tooth or the bitepiece, which can cause the active contour algorithm to grow outside the tooth contour. These could be identified by detecting these bridges and cutting off the growth automatically. A method for trimming parts of the contour based on concavity was proposed by Shah et al.[32]. Adding an alignment[15] or transformation[19] step during matching would also be interesting.

The brightness of dental X-ray images actually represents the thickness and density of the material it had to penetrate[26]. It would be interesting to see if this characteristic can be used as a feature, such as looking at the thickness of the teeth or bone at different points.

When dentists perform further work on a tooth, the dental work contours may merge or change in parts. When this happens, it would cause the Hausdorff-distance to increase significantly when compared to the older separate works. Looking into how this affects the results over time and testing other distance measures is necessary to ensure that the system is capable of handling change over time.

Bibliography

- [1] E. Said, D. Nassar, G. Fahmy, and H. Ammar, "Teeth Segmentation in Digitized Dental X-ray Films Using Mathematical Morphology," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 178–189, Jun. 2006.
- [2] H. H. A. Diaa Eldin M. Nassar, "A Prototype Automated Dental Identification System (ADIS)." 2003.
- [3] G. Fahmy, D. Nassar, E. Haj-Said, H. Chen, O. Nomir, J. Zhou, R. Howell, H. H. Ammar, M. Abdel-Mottaleb, and A. K. Jain, "Towards an Automated Dental Identification System (ADIS)," in *Biometric Authentication*, ser. Lecture Notes in Computer Science, D. Zhang and A. K. Jain, Eds. Springer Berlin Heidelberg, Jan. 2004, no. 3072, pp. 789–796.
- [4] M. Abdel-Mottaleb, O. Nomir, D. Nassar, G. Fahmy, and H. Ammar, "Challenges of Developing an Automated Dental Identification System," in *2003 IEEE 46th Midwest Symposium on Circuits and Systems*, vol. 1, Dec. 2003, pp. 411–414 Vol. 1.
- [5] I. A. Pretty and D. Sweet, "Forensic Dentistry: A Look at Forensic Dentistry Part 1: The Role of Teeth in the Determination of Human Identity," *British Dental Journal*, vol. 190, no. 7, pp. 359–366, Apr. 2001.
- [6] J. Zhou and M. Abdel-Mottaleb, "A Content-based System for Human Identification Based on Bitewing Dental X-ray Images," *Pattern Recognition*, vol. 38, no. 11, pp. 2132–2142, Nov. 2005.
- [7] L. Vincent, "Morphological Grayscale Reconstruction in Image Analysis: Applications and Efficient Algorithms," *IEEE Transactions on Image Processing*, vol. 2, pp. 176–201, 1993.

BIBLIOGRAPHY

- [8] M. Mahoor and M. Abdel-Mottaleb, "Automatic Classification of Teeth in Bitewing Dental Images," vol. 5, Oct. 2004, pp. 3475–3478 Vol. 5.
- [9] E. Said, G. F. Fahmy, D. Nassar, and H. Ammar, "Dental X-ray Image Segmentation," vol. 5404, 2004, pp. 409–417. [Online]. Available: <http://dx.doi.org/10.1117/12.541658>
- [10] J. Zhou and M. Abdel-Mottaleb, "Automatic Human Identification Based on Dental X-ray Images," vol. 5404, 2004, pp. 373–380.
- [11] M. H. Mahoor and M. Abdel-Mottaleb, "Classification and Numbering of Teeth in Dental Bitewing Images," *Pattern Recognition*, vol. 38, no. 4, pp. 577–586, Apr. 2005.
- [12] F. Aeini and F. Mahmoudi, "Classification and Numbering of Posterior Teeth in Bitewing Dental Images," in *2010 3rd International Conference on Advanced Computer Theory and Engineering (ICACTE)*, vol. 6, Aug. 2010, pp. V6–66–V6–72.
- [13] Y. H. Lai and P. L. Lin, "Effective Segmentation for Dental X-Ray Images Using Texture-Based Fuzzy Inference System," in *Advanced Concepts for Intelligent Vision Systems*, ser. Lecture Notes in Computer Science, J. Blanc-Talon, S. Bourennane, W. Philips, D. Popescu, and P. Scheunders, Eds. Springer Berlin Heidelberg, 2008, no. 5259.
- [14] P. L. Lin, Y. H. Lai, and P. W. Huang, "An Effective Classification and Numbering System for Dental Bitewing Radiographs Using Teeth Region and Contour Information," *Pattern Recognition*, vol. 43, no. 4, pp. 1380–1392, Apr. 2010.
- [15] P.-L. Lin, Y.-H. Lai, and P.-W. Huang, "Dental Biometrics: Human Identification Based on Teeth and Dental Works in Bitewing Radiographs," *Pattern Recognition*, vol. 45, no. 3, pp. 934–946, Mar. 2012.
- [16] V. Pushparaj, U. Gurunathan, and B. Arumugam, "An Effective Dental Shape Extraction Algorithm Using Contour Information and Matching by Mahalanobis Distance," *Journal of Digital Imaging*, vol. 26, no. 2, Apr. 2013.

BIBLIOGRAPHY

- [17] F. Shamsafar, “A New Feature Extraction Method From Dental X-ray Images for Human Identification,” in *2013 8th Iranian Conference on Machine Vision and Image Processing (MVIP)*, Sep. 2013, pp. 397–402.
- [18] R. C. Gonzalez and R. E. Woods, *Digital Image Processing, 3rd*, 2008.
- [19] O. Nomir and M. Abdel-Mottaleb, “A System for Human Identification From X-ray Dental Radiographs,” *Pattern Recognition*, vol. 38, no. 8, Aug. 2005.
- [20] ———, “Hierarchical Dental X-Ray Radiographs Matching,” in *2006 IEEE International Conference on Image Processing*, Oct. 2006, pp. 2677–2680.
- [21] ———, “Human Identification From Dental X-Ray Images Based on the Shape and Appearance of the Teeth,” *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 2, pp. 188–197, Jun. 2007.
- [22] ———, “Fusion of Matching Algorithms for Human Identification Using Dental X-Ray Radiographs,” *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 2, pp. 223–233, Jun. 2008.
- [23] ———, “Hierarchical Contour Matching for Dental X-ray Radiographs,” *Pattern Recognition*, vol. 41, no. 1, pp. 130–138, Jan. 2008.
- [24] N. Al-sherif, G. Guo, and H. Ammar, “A New Approach to Teeth Segmentation,” in *2012 IEEE International Symposium on Multimedia (ISM)*, Dec. 2012, pp. 145–148.
- [25] A. K. Jain, H. Chen, and S. Minut, “Dental Biometrics: Human Identification Using Dental Radiographs,” in *Proc. of 4th Int’l Conf. on Audio- and Video-Based Biometric Person Authentication (AVBPA)*, 2003, pp. 429–437.
- [26] M. Omanovic and J. J. Orchard, “Image Registration-Based Approach to Ranking Dental X-Ray Images for Human Forensic Identification,” *Canadian Society of Forensic Science Journal*, vol. 41, no. 3, pp. 125–134, Jan. 2008.
- [27] A. K. Jain and H. Chen, “Matching of Dental X-ray Images for Human Identification,” *Pattern Recognition*, vol. 37, no. 7, pp. 1519–1532, Jul. 2004.

BIBLIOGRAPHY

- [28] B. H. Le, Z. Deng, J. Xia, Y.-B. Chang, and X. Zhou, "An Interactive Geometric Technique for Upper and Lower Teeth Segmentation," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2009*, ser. Lecture Notes in Computer Science, G.-Z. Yang, D. Hawkes, D. Rueckert, A. Noble, and C. Taylor, Eds. Springer Berlin Heidelberg, 2009, no. 5762, pp. 968–975.
- [29] H. Chen and A. Jain, "Dental Biometrics: Alignment and Matching of Dental Radiographs," in *Seventh IEEE Workshops on Application of Computer Vision, 2005. WACV/MOTIONS '05 Volume 1*, vol. 1, Jan. 2005, pp. 316–321.
- [30] ———, "Tooth Contour Extraction for Matching Dental Radiographs," in *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004*, vol. 3, Aug. 2004, pp. 522–525 Vol.3.
- [31] M. Hofer and A. Marana, "Dental Biometrics: Human Identification Based On Dental Work Information," in *XX Brazilian Symposium on Computer Graphics and Image Processing, 2007. SIBGRAPI 2007*, Oct. 2007.
- [32] S. Shah, A. Abaza, A. Ross, and H. Ammar, "Automatic Tooth Segmentation Using Active Contour Without Edges," in *Biometric Consortium Conference, 2006 Biometrics Symposium: Special Session on Research at the*, Sep. 2006, pp. 1–6.
- [33] D. E. M. Nassar and H. H. Ammar, "A Neural Network System for Matching Dental Radiographs," *Pattern Recognition*, vol. 40, no. 1, pp. 65–79, Jan. 2007.
- [34] R. Brunelli and T. Poggio, "Face Recognition: Features versus Templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 10, pp. 1042–1052, Oct. 1993.
- [35] E. W. Dijkstra, "A Note on Two Problems in Connexion With Graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, Dec. 1959.
- [36] J. Canny, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.

BIBLIOGRAPHY

- [37] T. Chan and L. Vese, “Active Contours Without Edges,” *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, Feb. 2001.
- [38] P. Marquez-Neila, L. Baumela, and L. Alvarez, “A Morphological Approach to Curvature-Based Evolution of Curves and Surfaces,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 1, pp. 2–17, Jan. 2014.
- [39] D. Huttenlocher, G. Klanderman, and W. Rucklidge, “Comparing Images Using the Hausdorff Distance,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 9, pp. 850–863, Sep. 1993.
- [40] D. Nassar, A. Abaza, X. Li, and H. Ammar, “Automatic Construction of Dental Charts for Postmortem Identification,” *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 2, pp. 234–246, Jun. 2008.