

Modelling Wear Patterns on Footwear Outsoles

A DISSERTATION PRESENTED BY

XAVIER S. FRANCIS

TO

THE SCHOOL OF COMPUTING & INFORMATION TECHNOLOGY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

MASTER OF COMPUTING

UNITEC INSTITUTE OF TECHNOLOGY
AUCKLAND, NEW ZEALAND

MAY 2019

©2019—XAVIER FRANCIS
ALL RIGHTS RESERVED.

Modelling Wear Patterns on Footwear Outsoles

ADVISORS

PRINCIPAL	ASSOCIATE	ASSOCIATE
Dr Hamid Sharifzadeh	Angus Newton	Dr Nilufar Baghaei

ABSTRACT

The outsoles of footwear develop nicks, cuts, and tears via repeated exposure to the abrasive forces that occur between the outsole and the ground. These abrasions result in the formation of characteristics unique to the outsole and the individual wearing them; additionally resulting in the degradation of the outsole design imprinted by the manufacturer. The combination of these characteristics allow the forensic scientist to uniquely identify the individual to whom it belongs. Quite often a period of time can elapse between the discovery of a shoeprint at the crime scene and the identification of a suspect. In these instances, the forensic scientist must rely on their training and expertise—developed through years of experience and study—to determine if the crime scene shoeprint matches the outsole of the suspect’s shoe.

This work introduces a computational framework capable of modelling wear patterns on the outsoles of footwear. This model is able to predict the evolution of the wear pattern after an arbitrary time period given in weeks. We introduce an additional model capable of reconstructing the outsole back to its original state on a given week. This framework—built on convolutional neural networks—provides an objective point of reference for forensic scientists in their evaluation of outsole wear patterns.

Publications

Learning Wear Patterns on Footwear Outsoles Using Convolutional Neural Networks.

Accepted in the *IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, Rotorua, New Zealand—by X. Francis, H. Sharifzadeh, A. Newton, N. Baghaei, & S. Varastehpour (August 2019).

Feature Enhancement and Denoising of a Forensic Shoeprint Dataset for Tracking Wear-And-Tear Effects. Submitted to the *IEEE Global Conference on Signal and Information Processing*, Ottawa, Canada—by X. Francis, H. Sharifzadeh, A. Newton, N. Baghaei, & S. Varastehpour (November 2019).

Modelling Wear Patterns on Shoeprints Using End-To-End Deep Learning. Submitted to the *IEEE Transactions on Pattern Analysis and Machine Intelligence*—by X. Francis, H. Sharifzadeh, A. Newton, N. Baghaei, & S. Varastehpour (May 2019).

Contents

0	INTRODUCTION	1
1	COMPUTATIONAL FORENSICS	5
1.1	Shoeprint classification	6
1.2	Shoeprint wear	10
1.3	Machine learning approaches	13
2	PREPARING THE DATASET	15
2.1	Issues	16
2.2	Prior work	21
2.3	Proposed denoising methodology	22
2.4	Evaluations	25
3	MODELLING WEAR PATTERNS	28
3.1	Related work	29
3.2	Proposed modelling methodology	29
3.3	Experiments	33
4	CONCLUSION	39
	APPENDIX A MODEL EXPERIMENTS & SOURCE CODE	41
	BIBLIOGRAPHY	50

FOR MY MOTHER.

Acknowledgements

To my advisor, Hamid—for his guidance, for his faith in my abilities, and for the opportunity to work on this project,

To my co-advisors, Angus and Nilufar—for their invaluable insights and assistance,

To my parents, Mary and Francis—for their conviction that my education was their best investment, and

To all those who came before me—whatever their contributions to the field, great or small.

I stand on the shoulders of giants.

Thank you.



Introduction

Forensic science sees the application of scientific methods to aid in the practice of criminal investigations. The earliest applications of the field can be found as far back as the 13th century in China [1]. Practitioners—known as forensic scientists—collect, process, and analyse evidence and other materials found at the scene of a crime to aid in the criminal investigation. Over the years, forensic science has brought many other sciences into its fold to be applied to criminology; spawning numerous sub-disciplines that specialise in art, botany, chemistry, archaeology, accounting, digital forensics, and much more. With the human element, forensic scientists consider human behaviour and biometric signatures such as DNA [2], hair [3], blood [4], fingerprints [5], etc.

Forensic podiatry is the sub-discipline that deals with the application of podiatric knowledge to crime scene investigations [6]. This practice aims to answer any questions regarding matters of the foot or footwear to forensic science and so studies aspects like bare footprints, footwear, and gait analysis. Some of the most common types of evidence found at crime scenes are marks and prints formed by the footwear worn by the criminal(s); their study being recorded as far back as 1786 [7]. These marks are imprints formed by the outside sole (outsole) of the footwear when it comes into contact with a surface. Imprints may be three-dimensional, when present in surfaces like sand or snow—or more commonly—two-dimensional impressions found on carpet, tiles, and other flooring. To capture shoeprint evidence, a cast is taken if the impression is in a 3-dimensional

setting, else an image is obtained through gel or photography. The image of the shoeprint is first digitally enhanced, removing noise and other artefacts of scanning; and then compared against a database of known prints and prints collected from other crime scenes to find a match [8]. Such shoeprints are capable of uniquely linking a suspect to a crime by the skilled forensic practitioner as they are as strong as other kinds of impression evidence like fingerprints [9]. They also serve to link different crime scenes of the same offender [10].

Despite their uniquely identifiable nature and the frequency of their discovery at crime scenes, shoeprints are not often used as evidence in a criminal trial, due to the variation in quality of scene-of-crime prints. Often, the retrieved crime scene print can be incomplete or degraded. A study out of the Netherlands found that in 1993, only 500 out of 14,000 shoeprints were positively identified [11]. The variations in quality can arise from the surface on which footwear pressure is applied, or from the walking patterns of the wearer. Nevertheless, shoeprint evidence have proven to be an effective method of linking suspects to crime scenes and the study of footwear has earned a place under the forensic science umbrella.

The characteristics that describe shoeprints are of two types. *Class characteristics* are those features of a shoeprint that are capable of linking it to the manufacturer, brand, and model of the shoe [12]. They describe the geometric patterns of the outsole imprinted by the manufacturer for aesthetic and practical purposes. Since there are a large number of shoe manufacturers, models, and designs, class characteristics serve to narrow down the search space to identify a given suspect's shoeprint. *Individualising characteristics* are those characteristics that arise on the outsole as a function of wear [13]. They consist of nicks, cuts, abrasions, punctures, tears, and other types of marks that form on the outsole as a result of the wearing process. The position, configuration, and orientation of each of these defects—combined—are capable of uniquely identifying the shoe on which they form. The position of a defect is its location relative to the perimeter, tread patterns and the like. Shapes are defined by width, height, and area. The rotation of a characteristic helps to separate it from other similarly shaped defects. The combination of class and individual characteristics are analysed by a shoeprint examiner to uniquely identify a suspect's shoe and to attribute it to a crime scene shoeprint.

The *wear pattern* is the sum total of individualising characteristics present on the shoeprint. This is unique to an item of footwear and is formed by erosion of the outsole through repeated use. The environment the shoes are worn in, the height and weight of the person wearing them, the frequency and nature of use are all factors that influence wear. To analyse these wear patterns, forensic scientists rely on their knowledge and skills—gained through years of study and experience. Quite often a gap of weeks or months can elapse before a suspect is identified and his/her shoeprints are

obtained. Insights into how wear characteristics develop over time can help the forensic investigator develop a better ability to interpret them. Pattern recognition can be subjective and an experienced examiner may more accurately identify and assess patterns in evidence [14]; although attempts have been made to formalise their interpretation [15–18]. Computer aided identification and analysis can help alleviate inherent human biases as well as intra-analyst variation in identifying features or characteristics in the evidence by complementing the forensic examiner’s skillsets.

Due to the rate of outsole patterns being designed and manufactured, any archive of shoeprints will quickly go out of date, if not diligently maintained. This compounds the difficulty of identifying class characteristics of crime scene prints. To assist the forensic examiner in such tasks, computational methods have been developed, driven by the interdisciplinary field of *computational forensics*—the application of computational models, analysis, algorithms, and simulations to assist in the identification and analysis of forensic evidence [19]. Earlier approaches relied upon statistical and mathematical methods to form quantitative analyses of evidence. Criticisms of forensic testimony for lacking a scientific basis—among other reasons—have led to the development of computer aided methods in assisting forensic investigations [20]. These methods can objectively analyse and identify the class characteristics of a given shoeprint, providing metadata—like make and model—exponentially faster than an individual can look it up.

This thesis contributes to the literature by introducing a computational model and framework for studying and analysing wear patterns on footwear outsoles. We present a methodology that leverages deep neural networks (DNNs) to model the wear pattern captured in a unique dataset of shoeprints. We hope that our models can serve as an empirical point of reference for the forensic science community; additionally providing a platform for novice forensic scientists to hone their skills with.

* * *

The rest of this dissertation is organised as follows:

- In Chapter 1 we review the field of computational forensics; surveying the work in the domains of shoeprint classification and shoeprint wear, with a special focus on machine learning applications in the field.
- Chapter 2 introduces a unique dataset collected for the purpose of studying wear patterns. We describe the challenges faced in preparing this dataset for predictive modelling and detail a novel denoising methodology that satisfies our unique constraints.

- Chapter 3 presents the core of the thesis—computational models of shoeprint wear. We describe the mathematical algorithms that power our methodology and introduce two models built using this method—the first capable of predicting wear patterns on outsoles and the second able to reconstruct the outsole to its original state.
- Finally, in Chapter 4, we summarise the ideas of the thesis and point to directions of future work.

Chance favors the prepared mind.

Louis Pasteur

1

Computational Forensics

The forensic examiner's role in a criminal investigation is to identify, preserve, and analyse all evidence of crime in order to—(i) eliminate items that are not relevant to the investigation, (ii) identify evidence that may link a suspect to the crime, (iii) provide proof of an individual's involvement in a crime, and (iv) to provide written reports or testimony as required.

The responsibilities of the forensic *footwear* examiner are—(i) to identify class characteristics of a given shoeprint, by comparing it against a large set of known prints to discover details like manufacturer, make, and model and (ii) to consider the individualising characteristics of the print to assign the print to an owner. The computational methods developed to assist the examiner in these tasks fall under three domains:

1. Shoeprint classification—systems that match class characteristics in a query image to those in a database of reference images,
2. Shoeprint wear—those studies that look at the formation of accidental characteristics on outsoles over a period of time, and
3. Machine learning—research that sees ML methodologies applied to the analysis of shoeprints.

1.1 Shoeprint classification

Researchers have to work with several challenges when designing systems of shoeprint recognition. High quality reference images found in reference databases are taken by the shoeprint examiner in laboratory conditions, while the query prints that need to be matched are scene-of-crime images (socs). Crime scene impressions are often incomplete, noisy, and/or distorted. Automated methods have to consider these variables to be a robust recognition system. At a high level, all the methods surveyed in this section follow a similar pattern—first a data pre-processing stage is conducted, where the input image is cleaned, denoised, and rotated, as required. Then key-points or *features* of the shoeprint are computed and compared against the same metric for each shoeprint in the database, in the hopes of finding a match for class characteristics. Over the years, shoeprint retrieval systems have developed more efficient ways of pre-processing, feature extraction, and similarity measurement. We begin by looking at classification methods that are partially automated.

1.1.1 Semi-automated methods

The earliest methods used a semi-automated approach and appeared in the mid 90s [21–23]. By manually encoding the obtained print and reference images using shape primitives (e.g. geometric shapes, patterns etc.), these systems were able to find a match [24, 25]. These methods required significant manual intervention, were time-consuming in performance, and were prone to error because different analysts may encode the primitives in different ways. The practice of shoeprint classification was new and immature, but early results instilled confidence that computational methods could prove effective in aiding the forensic examiner.

1.1.2 Automated methods

One of the earliest automatic methods was presented by Geradts and Keijzer [11] who developed a database of reference shoeprints (labeled ‘REBEZO’), in addition to their proposed method which computed Fourier features of print segments and used an artificial neural network for classification. Alexander et al. [26] devised a method where shoeprint images are decomposed into fractals and the coefficients are matched against a database of prints. The fractal decomposition of a database image with minimal changes to the image under test is considered to be a match. Mean square noise error was used to compute this change. In later work [27], they test the system’s robustness in dealing with rotation and translation.

de Chazal et al. [28] discussed a technique that used discrete Fourier transformations, and calculated the coefficient of power spectral density (PSD). Pattern matching was done using 2D correlation coefficient similarity measure. Their tests showed the system returning a positive match

in the first result 65 percent of the time. However, their model works regardless of the spatial positioning of the print. Zhang and Allinson [29] represented shoeprint features using edge-direction histograms, and computed 1D discrete Fourier transforms on the normalised histogram. The similarity measure used was the Euclidean distance. Su et al. [30] introduce a method of thresholding shoeprint patterns in noisy images by extending a general model out of non-local mean filtering.

Pavlou and Allison [31] presented a technique that uses local image features, where maximally stable extremal region (MSER) feature detectors are encoded using scale invariant feature transform (SIFT) descriptors to obtain a match, using a Gaussian weighted similarity metric. Pavlou and Allison on this work by encoding a codebook of histograms representations for each shoe pattern and clustering them using k-means to facilitate fast indexing [32]. Ghouti et al. [33] proposed using directional filter banks (DFBs) to create a condensed representation of a shoeprint they called ‘ShoeHash’. This method would encapsulate both local and global details of the print and used a normalised Euclidean distance to assess similarity. Su et al. [34] have utilised a hybrid pattern and topological spectra—using the Euler number—and computed similarity using the normalised value of these two measures. In Crookes et al. [35], the authors propose a feature detection method they call ‘enhanced local image feature’ which combines an automatic Laplace-based scale selection with the Harris corner detector. SIFT descriptors are used to represent features. They match two images based on their descriptors using a nearest-neighbour search.

Gueham et al. [36] evaluated an advanced correlation filter—OTSDF—for classification. The method processes partials and distorted prints. They propose an alternate method [37] where feature extraction is done using the the Fourier-Mellin transform (which involves a log-polar mapping and a Fourier transform). They use a 2D correlation method to find similarities between shoeprints. Their method was tested on a dataset including degraded images and proved to be efficient. In Al-Garni and Hamiane [38], Hu moment invariants are used to classify shoeprints. They opted to compute similarity with the most common distance measures seen in the literature—City-block, Euclidean, Canberra, and Correlation. The method performs perfectly with rotated images, and well with lower resolution images, but performance drops drastically on noisy query images. Xiao and She [39] further the work done with PSD by pairing it with Zernike moments—the Zernike method being used to handle distorted prints, and the similarity measured by the correlation coefficient of PSD.

Jing et al. [40] performed feature extraction based on the direction of patterns present in the print (vertical lines, circular shapes, geometries etc.). They capture patterns on three levels—using co-occurrence matrices, directional masks, and local and global Fourier transforms. Using these feature vectors, they compute a similarity using the sum-of-absolute-difference. In Nibouche et

al. [41], Harris points and SIFT descriptors are combined in a method of feature extraction. The method was compared to de Chazal et al. [28] and found to be superior. This method is designed to work with partials, and is rotation and noise resistant. Matches are found using the RANSAC algorithm.

Dardi et al. [42] noted that most of the literature that came before them were tested on artificially generated crime scene prints, and developed a texture based method that computed the Mahalanobis distance for a print as the feature descriptor, which they tested on real crime scene prints. The authors go on to test this system in subsequent work [43, 44], comparing the performance using synthetic and real crime-scene prints against de Chazal et al. [28] and Gueham et al. [36]. Wang et al. [45] have used wavelets as an edge detector and neural networks for recognition. In Patil and Kulkarni [46], shoeprint rotation estimation is handled using Radon transform, features are computed using Gabor transform, and matched with the Euclidean distance. Additionally, it is rotation and intensity invariant and performs well on partials. Pei et al. [47] extract features based on texture and geometry by combining odd and even Gabor filters. Texture features are used to retrieve matching prints, and the geometry features are used as a similarity metric. This method is robust against noise and partial prints.

In Tang et al. [8, 48], they proposed a method of feature extraction using an iterative variant of the standard Hough transform to detect lines and circles, and a modified randomized Hough transform to detect ellipses. Once the features are extracted, they are represented with an attribute relational graph. They also introduce the ‘footwear print distance’ measurement. Prints are clustered using recurring patterns and similarity assessed using a cumulative match score. This end-to-end system is claimed to be distortion tolerant as well as being tolerant of partial prints. Li et al. [49] use SIFT to construct scale-spaces and detect points of interest in an object. Similarity by way of cross correlation is measured for each of the key-points on a shoeprint. The method is invariant to illumination, rotation, translation, and scale. Hasegawa and Tabbone [50] decompose prints into connected components and use the histogram radon transform (HRT) descriptor. Similarity is measured by the mean of local similarity measures. HRT has the advantage of being robust to geometric transformations of the components.

Wei et al. [51] developed a system that relies on SIFT descriptors to construct features out of local extrema; after constructing different scale spaces to detect these extrema. Similarity is measured using cross-correlation. Kong et al. [52] have combined Gabor and Zernike filters as feature descriptors, and use a normalised correlation coefficient score as a matching metric. Li et al. [53] formulated an algorithm using the Gabor transform, by computing the histogram of the integral. The histograms are also used to measure similarity. Min and Qi’s classification algorithm [54] con-

structs a feature circle from detected feature points using the RANSAC algorithm; feature points being detected with a discrete Fourier transform. Almaadeed et al. [55] use the SIFT descriptor to gain rotation invariance and achieve scale invariance using both Harris and Hessian detectors. Combining the two detectors outperforms similar algorithms and delivers partial matching that is noise and distortion resilient.

Li and Wang [56] propose an automated algorithm to position a print image along the vertical axis, by using vertex angle and secondary positioning. Gwo and Wei [57] propose another method of shoeprint alignment, by computing the core point. Here, contour point detection and curve fitting are used to detect a shoeprint. Once the core point is found, the print is partitioned into regions from which Zernike moments are computed to develop a pattern description. Shoeprints are matched by calculating the Euclidean distance of the pattern. The paper goes on to discuss variations of Zernike methods, and optimisations for better results. Alizadeh and Kose [58] divide shoeprints into two blocks and compute their sparse representation for feature extraction. Two dictionaries are also computed for the reference images; with L_1 minimisation being used as features of the test image. Using a cumulative match score they assess whether the method returns a match.

Kortylewski et al. [59] developed an unsupervised method of shoeprint retrieval, utilising periodic patterns and local Fourier transforms as features. Their method was able to identify outsoles with alternative patterns and was designed to perform in unconstrained noise conditions. Wang et al. [60] use manifold ranking to bridge the distance between image content and semantic information. To rank images, they consider three factors—(i) the similarity of features between query and database images, (ii) the relationship between every two images in the dataset, and (iii) an opinion score assigned to every crime scene print by the investigator. The normalised image is decomposed using Haar Wavelets and the PSD of each wavelet band is computed to find the similarity. They claim the performance to be far better than similar methods introduced in prior literature. This work builds on prior research done by the authors [61].

Some of the methods seen above do not account for variations in query images, such as distortions and noise. Others are not designed to work with partial prints, as are often found at crime scenes. Comprehensively evaluating each method is a challenge, especially without a standardised dataset. Luostarinen and Lehmussola [62] attempted an evaluation of eight of the popular methods—testing their performance against partial, rotated, and noisy prints to find RANSAC-based methods the best performers. Richetelli et al. [63] evaluated phase only correlation (POC), Fourier-Mellin transform (FMT), and SIFT + RANSAC algorithms to find that POC outperformed the others.

1.2 Shoeprint wear

In Fruchtenicht et al. [64], wear is defined as “a continuous alteration of class and accidental characteristics that can result in an individual appearance.” In other words, individualising characteristics that appear on the outsole of a shoe—which forensic examiners rely on to uniquely identify a shoeprint—arise as a function of the *wear process*. These characteristics are unique due to the nature of the shape and size of the wearer’s feet as well as their biomechanics. Bodziak defined wear [9] as “the erosion of the outsole due to abrasive forces that occur between the outsole and the ground.” Evaluating the potential of accidental characteristics as evidence requires some intuition into how wear patterns form on the outsole—the rate at which they appear, disappear, and are retained [65]. Studies on individualising characteristics have looked into the interdependence of their features [66], the development of tools to evaluate their rarity [67], and more; but the focus of this review are studies that consider wear formation over time. First, we look at studies that account for wear manually.

1.2.1 Manual methods

Cassidy [68] undertook 3 research projects to look at accidental characteristics in detail using manual examination. Project 1 observed two controlled groups and looked at the evidential value of general wear and the chance of accidental marks re-occurring at the same location on another shoe. The odds of wear being duplicated were found to decrease in proportion to the length of time that a pair of shoes were worn. Project 2 looked at accidental characteristics of 4 different heel patterns over 6 months to see if the accidentals could re-occur in the same spot. The author calculated a 1 in 60 chance of an accidental characteristic being duplicated in this study; noting that the project was conducted in a controlled environment that favoured the duplication of these characteristics and that real world samplings would differ. Project 3 looked at the durability of accidental patterns on a single pair of brand new rubber heels over the course of 68 days. Impressions were taken on the first three days and at one week intervals thereafter. Roughly 33% of the characteristics were observed to last through the duration of the study; with quite a few only lasting a very short time. The author’s conclusion is that footwear evidence should not be discarded after a couple of weeks for fear that they will be of little value, and that positive identification can be made much later.

Wyatt et al. [69] collected 54 sets of shoeprints from different individuals as well as details like shoe size, type, age etc. from each wearer. Another set of prints were collected from the same volunteers and shoes after a gap of 2 months. Three separate analysts then analyzed the shoes for—(i) individualising characteristics that appeared in the first batch of outsole impressions but not in the second, (ii) those that appeared in the second set of prints but not in the first, (iii) those that ap-

peared in both sets of prints, and (iv) general observed wear patterns. Based on the experience of the analysts, 22 out of the 54 sets could be positively traced back to the origin. Some characteristics from the first set were found to have either worn away, or to have been occluded by new characteristics formed in the second set. The majority of shoeprints were found to have developed wear patterns and new individualising characteristics in the span of two months. They note that characteristics present in the first set were very likely to be observed in the second.

Adair et al. [70] studied the accidental patterns that formed on hiking boots worn over a 3.5 hour, 11.27 kilometre hiking trip. The six participants wore a pair of boots while ascending Mount Bierstadt in Colorado and another pair for the descent. All boots observed in the study were brand new. Each outsole was found to have sufficient characteristics to allow for individualisation, and to be differentiated from the other outsides in the study. By using the same style of boots, for the same duration, in the same environmental conditions, and over the same walking path, they eliminated many of the variables that contribute to wear pattern formation, leaving only the individual's walking mannerisms and the random form in which outsides make contact with the topography of the terrain; thereby supporting the generally agreed upon hypothesis that accidental pattern formation on footwear is generated randomly.

Moorthy and Chelliah [71] investigate the wear that develops on 2 sets of shoeprints from 2 individuals over the course of 3 months. Impressions were collected from both subjects on days 1, 45, 60, and 90. The analysts then looked at the ball and heel of the prints to observe the development of accidental marks. An appreciable amount of wear was in the prints, with many accidentals increasing in length over time. New characteristics were discovered in the later shoeprints, while ones that were seen in early impressions disappeared over time; thus contributing to making that footwear unique.

1.2.2 Computational methods

Petraco et al. [72] applied statistical and pattern recognition methods to the study of accidental characteristics. Five pairs of brand new shoes of the same make and model are worn by an individual for a length of 30 days for each shoe. With all shoes being worn by the same individual, it is expected that wear patterns will form around the same regions, since many of the variables affecting wear are eliminated in this study. For each pair of shoes, 15 prints were recorded over 30 days. Only the location and quantity of accidental characteristics were looked at, disregarding shape and size. These characteristics were represented as feature vectors where each component showed the number of accidentals on a shoe on a given day. Patterns were compared by reducing the feature vectors using principal component analysis (PCA) and measuring the distance in PC-space using the

MLG-LCA distance metric. Using these statistical techniques, the study aimed to assign an unknown pattern to a known shoeprint with the most similar accidental patterns. PCA was applied to all patterns from days 1 to 30, as well as segmented datasets from the later portion of the study in order to look at the wear effect in detail. As expected, shoeprints taken from near the end of the study had accumulated the most accidental characteristics. By plotting the patterns on a 3D plot, it was apparent that accidental marks became more distinct as time passed. Patterns formed by the same shoe tended to cluster closer together. Length of wear and the number of accidentals accounted for were the two factors which most influenced classification rates. The authors claim that their method could be even more successful if it accounted for physical characteristics of the wear patterns.

Sheets et al. [73] embarked on a study to understand the appearance and disappearance of accidentals on shoe soles over time, by studying the wear effect—which they define as “the removal of the original texture on the outsole of the shoe.” The authors purchased 11 pairs of shoes and introduced artificial cuts onto fixed locations on the outsoles, in an attempt to mimic randomly acquired characteristics. These shoes were then worn by volunteers over a 7 week period, in the same manner in which they wore their daily shoes. The outsoles were then scanned at the 2, 4, and 7 week milestones and digitally processed to analyse them. A feature vector method was used to capture wear and accidental characteristics. This allowed them to capture information about the size and location of characteristics as well as the wear. A semi-automated method to process the feature vectors was implemented. Software was written to align the print to a grid. A human operator would then record the percentage of observed characteristics present in each grid cell. This resulted in two feature vectors of 200 numbers each, which captured the characteristics and wear of the shoe. However, they made no attempt to record the location, orientation, or shape of a characteristic. To study these characteristics, they utilised a multivariate statistical algorithm—PCA—which allowed them to summarize and visualise the data on a 2D plot. The study has several findings—the authors conclude that most wear occurred in the heel of the shoes, with some found in the ball of the foot. The lack of substantial change seen over seven weeks lends to the possibility of linking shoes to crime scene impressions even after several weeks. They were surprised at the lack of appearance of new characteristics, which was expected based on previous studies, and theorise that this may be due to the urban environment in which these shoes were worn. They suggest that worn shoes may acquire characteristics more readily than new shoes, as used in the study. They note the effectiveness of feature vectors in representing the data and capturing characteristics, allowing for detailed multivariate analysis, as well as being capable of describing total net wear. The paper concludes with the suggestion that capturing the class characteristics in addition to the above would facilitate automated methods of classification and comparison.

1.3 Machine learning approaches

Applying computational power to manual work affords the user the ability to automate tedious tasks. In the case of work that requires attention to fine-grained details—as in shoeprint matching—computational methods can greatly enhance the examiner’s skillsets by bringing to light details that he or she may have missed. Machine learning (ML) methodologies have the proven ability to capture information that a human might miss, as well as being well-suited to automation. ML has been successfully applied to many areas of science—often delivering state-of-the-art results—in areas as diverse as predicting cardiovascular risk [74], identifying craters on the moon [75], and in detecting earthquakes [76].

Geradts and Keijzer [11] pioneered the application of computational methods to footwear evidence. They were one of the first to attempt automated shoeprint classification; collaborating with the Norwegian police to develop a database of shoeprint images for the task in 1992. A single-layer feedforward neural network was used by them for the classification task, using features computed by the Fourier transform. Sun et al. [77] approached shoeprint analysis from the perspective of forensic data mining; clustering shoeprints using k-means and expectation maximisation (EM). Their objective was to analyse the results of each algorithm and to visualise the results. Wang et al. [45] applied wavelets as edge detectors to capture accidental characteristics in four categories—triangles, circles, ellipses, and irregulars. Feature vectors were constructed using the angle, length, and region of these characteristics. A competitive neural network was then trained using fuzzy rules to automatically judge which shape an input belonged to.

Ramakrishnan and Srihari [78] used conditional random fields to extract shoeprint pattern foreground from the background of images, their method performing better than the baselines—Otsu thresholding and neural networks. Kortylewski and Vetter [79] extend hierarchical compositional models (HCM) with a statistical framework. The parameters of the model were determined by a greedy EM-style clustering algorithm introduced by the authors. This pattern model is formulated in a fully probabilistic manner. The authors claim state-of-the-art results on shoeprint classification. In addition, the model is said to be resistant to distortions in the image. Most recently, Kong et al. [80] have developed a method of shoeprint recognition using deep convolutional neural networks (CNNs). They detail the training of a ‘Siamese’ network model using both reference database prints and scene-of-crime images. They also propose a multi-channel normalised cross-correlation measure of similarity. Zhang et al. [81] were also among the first to adapt CNNs to shoeprint retrieval systems.

* * *

The study of shoeprints in the field of computational forensics has been heavily biased towards shoeprint retrieval systems. While this is an extremely valuable and challenging area of study, shoeprint wear is an equally important area to be considered, and yet has seen relatively little interest. In a criminal investigation, there can be a gap of weeks or months before a suspect is identified, and their shoeprints obtained. Insight into how these characteristics develop over time can help investigators and forensic scientists develop a better ability to interpret them. The first step in designing a study of wear patterns must necessarily be the collection of data to be studied. The challenge of limiting the variables that influence wear, and the extended timeframe required to capture sufficient data may be some of the reasons that the research in this field is sparse.

In the next chapter, we introduce a unique dataset explicitly collected for this purpose—by capturing the life and wear of a pair of shoes.

Truth... is much too complicated to allow anything but approximations.

John von Neumann

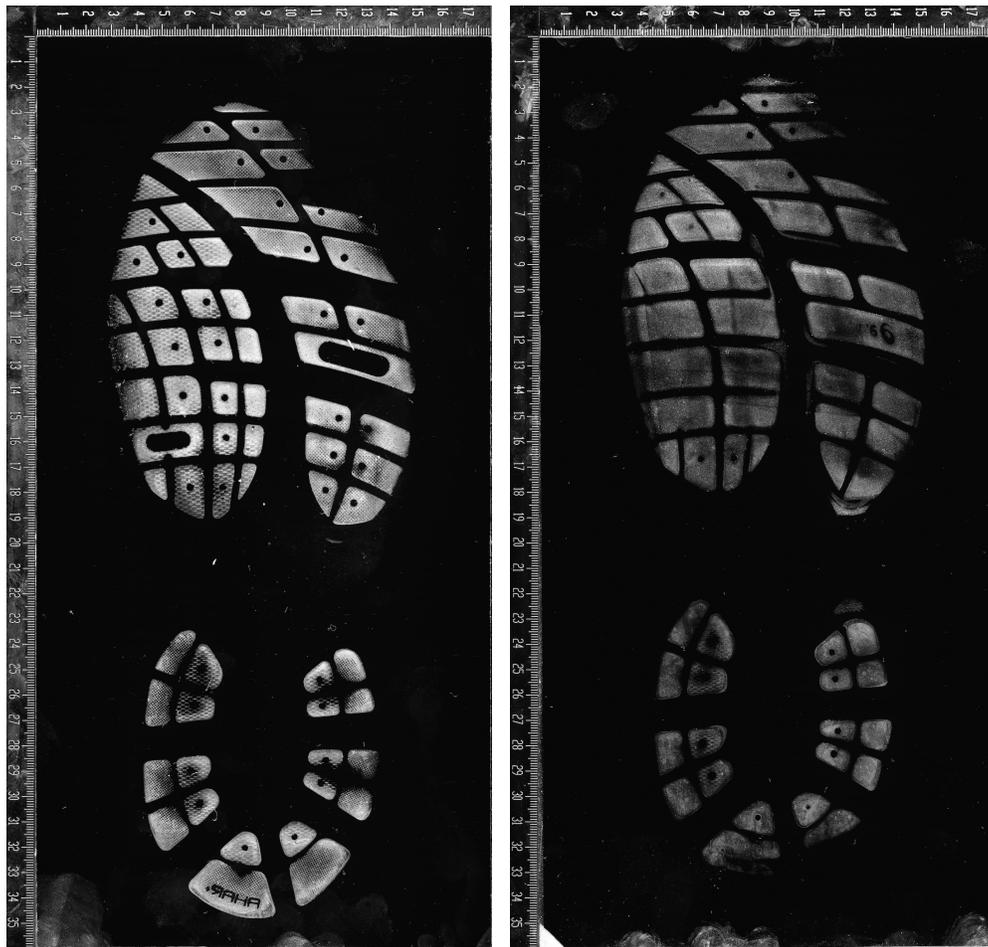
2

Preparing the Dataset

Class characteristics can exclude members of a suspect population but cannot uniquely identify a shoe by themselves. Characteristics acquired through wear can reveal information that strongly connects an individual to a crime scene. Evaluating the wear process requires some background information about the rate at which accidental characteristics are acquired, the extent to which they are retained, and how often they occur. To this end, a data collection project was initiated by the *Institute of Environmental Science and Research* (ESR), New Zealand.

As the *sole* provider of forensic services to the New Zealand Police, ESR has extensive experience with work of this nature. In collaboration with their forensics department, a year-long data collection project was undertaken while attempting to limit the variables that influence wear—the product of which was a series of images that captures the wear and life of a single pair of shoes.

A pair of *Asics*-brand men's sneakers were worn by the forensic scientist every day over the course of a year in an urban environment. The outsoles consisted of approximately 63 *block features* of varying size and shape. Impressions of both outsoles were captured every fortnight; yielding 52 impressions including week 0 (regrettably, impressions were not made for week 4). These impressions were captured using the gel-lift method, using *BVDA*-brand gel-lifters, and then scanned into high-resolution digital negatives in the `.tiff` file format. Each file is a 256-level grayscale image. Two sample images from this dataset are shown in Figure 2.1.



(a) Right outsole on week 2.

(b) Right outsole on week 44.

Figure 2.1: Original images from our dataset.

2.1 Issues

While every effort was made to remain consistent during the course of recording these impressions, many forms of unwanted features were captured in addition to the shoeprint. These included air bubbles, fingerprints, dust, debris overlapping the shoeprint, ghosting of the impression, and areas of missing detail. The observations given below were made while studying the dataset:

- Debris consists of different sizes/materials—some appear to be fibres caught on outsole, and others are larger unidentified artefacts.

- Some debris appear to have been dislodged in the process of imprinting; leaving areas of the outsole missing in detail.
- All of the images have varying levels of (apparent) fingerprint noise. In the worst cases (e.g. week 0), the fingerprint partially occludes the outsole impression.
- Certain regions of the outsole (e.g. ridges above right side of heel) appear in a few of the prints but are missing/weak in others.
- All the images have some noise in the form of debris, but they are most notable in the prints from weeks 0–30.

Table 2.1 provides a comprehensive list of the issues that were deemed to be potentially problematic to the performance of the model, accompanied by the comments of the forensic examiner where relevant, and figures to aid the reader’s comprehension.

Table 2.1: A list of issues found with the dataset.

Image	Issue	Remark
0L, 0R	Air bubbles	-
	Fingerprint overlapping outsole (Figure 2.2a)	-
2L, 2R	Air bubbles	-
4L	Missing detail (Figure 2.2b)	-
8L, 8R	Air bubbles	-
	Fingerprints	-
14R	Missing detail (striations/spots)	-
18R	Unknown impression above print (Figure 2.2c)	-
	Air bubble overlapping print (Figure 2.2d)	-
20L, 20R	Missing detail (striations/spots) (Figure 2.2e)	Regular symmetrical circular shading within block feature in forefoot area. Seen in some other block features. Can only think that these are bubble impressions (symmetrical shape is the major clue) caused by trapped air bubbles between gel and coversheet, which have pressed down on the gel-lift creating the slight difference in tone.

Table 2.1 continued

22L	Ghosting (Figure 2.2f)	Ghosting over block features at toe region. Yes, repositioning is the reason for the doubling of the impression in this region. Despite best efforts, training and experience, it is still possible to experience slight shifting of the shoe during the printing process, especially at either end of the shoe.
	Missing detail (striations)	-
22R	Bleeding	-
	Striations	-
24R	Unknown (Figure 2.2g)	Very thin, straight line features converging at ~right angle between the forefoot and heel patterns. This has the appearance of the edge of the corner of the cover sheet, which has come into contact with the gel lift.
26L	Unknown / Ghosting (Figure 2.2h)	Ghosting over and between block features on the toe region. Possibly light contact between the toes/front region of the sole and the gel? Not clear so I'm uncertain.
28L	Missing detail	-
28R	Bottom region of heel missing	-
30L, 30R	Overlapping fingerprint	-
32L	Regions of low intensity throughout print	-
32R	Horizontal striation (Figure 2.2i)	Two straight-edged bands running across width of forefoot, producing a darker region than surrounding block features. Uncertain about cause of feature. Possibly underlying features beneath the gel when print was made (unlikely as floor is flat linoleum)?
	Low intensity regions	-
34L, 34R	Heel impression very weak; lots of detail missing	-
	Striations	-
	Regions of low intensity	-
34R	Missing detail	-
	Unknown mark (Figure 2.2j)	Large, bright irregularly-shaped features within block features in heel and forefoot, with accompanying smaller features. Appears to be a reflective surface contamination on the shoe sole? Possibly debris of some description? Uncertain.
	Striations	-
36L		

Table 2.1 continued

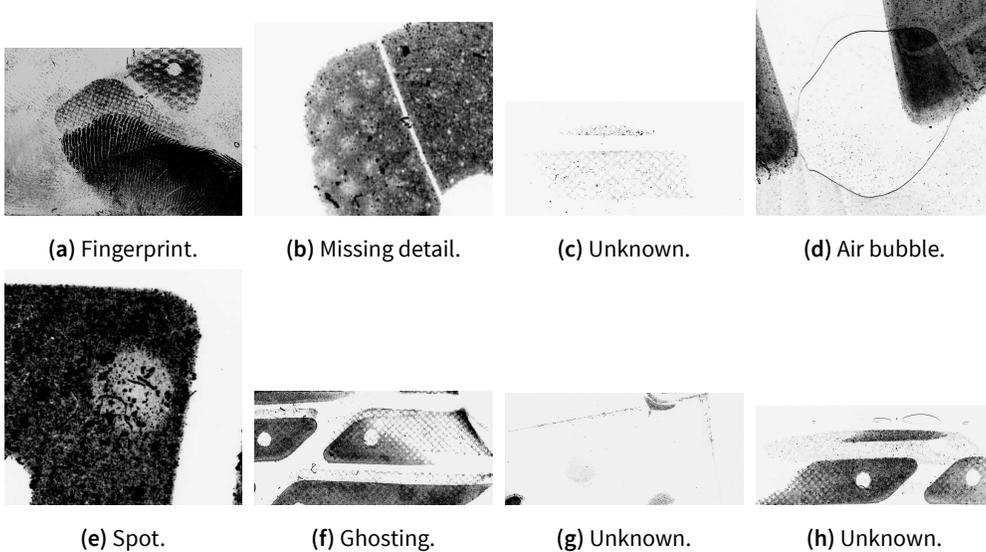
	Unknown mark (Figure 2.2k)	Amorphous, irregularly-shaped feature in a block feature in the heel area, with a particulate texture that has a different appearance than the surrounding block feature. Also encircled by a halo. Suggest this is some sort of substance stuck to the heel.
	Bottom region of heel missing	-
36R	Striations	-
38L	Striations	-
38R	Missing detail	-
	Areas of low intensity	-
40L	Striations	-
	Low intensity regions	-
40R	Unknown (Figure 2.2l)	Small circular feature within a block feature in the forefoot area. Has a smooth texture. Looks to be a compression of the block feature with a circular shape. Therefore, possibly an air bubble impression on the shoeprint?
	Striations (Figure 2.2m)	A difference in intensity separated by a sudden linear border that spans two neighbouring block features in the rear edge of the forefoot. Uncertain of cause, except perhaps a re-laying of the cover sheet?
42L	Unknown (Figure 2.2n)	As for 24R, very thin, straight line features converging at \sim right angle between the forefoot and heel patterns. This has the appearance of the edge of the corner of the cover sheet, which has come into contact with the gel lift.
	Bottom region of heel missing	-
42R	Bottom region of heel missing	-
46R	Unknown (Figure 2.2o)	Elliptical shaped feature in a block feature at the rear of the forefoot. Indistinct margins and texture has similar appearance to surrounding block. Possibly a temporary slight compression of the sole material due to an air bubble or other object? Uncertain.
50L	Left region of heel missing	-
52L	Unknown (Figure 2.2p)	Irregularly-shaped darker features towards the rear of the forefoot. Margins are indistinct. Could be a crease or wrinkle in the gel at the time the impression was made?

Along with possible explanations for the issues, the following general comment was provided:

Interesting to read your note regarding the amount of debris in weeks 0–30. The difference would coincide (week 32) with when I started to tape lift the shoe soles prior to printing. I finally looked at the scanned images that were being sent back to me and noticed the large amount of debris on the soles. Knowing that this would impact on the subsequent analysis of the data, I started cleaning the soles with the tape lifts.

Given this information, we speculate that the tape lift method might be responsible for the inconsistent appearances in intensity (striations and spots) observed in the latter half of the dataset. Perhaps some residue from the tape was left behind on the outsole and then transferred to the gel while imprinting.

Out of the issues identified above, the most egregious was determined to be the debris which appeared to consist of fibres and other objects that were transferred from the outsole and onto the gel in the process of imprinting. The debris was particularly problematic due to it obscuring regions of interest in the shoeprint. For the purpose of our research we refer to these unwanted features that obscure the object of interest as *noise*. To correct this noise, we returned to the literature and surveyed the denoising methods presented there.



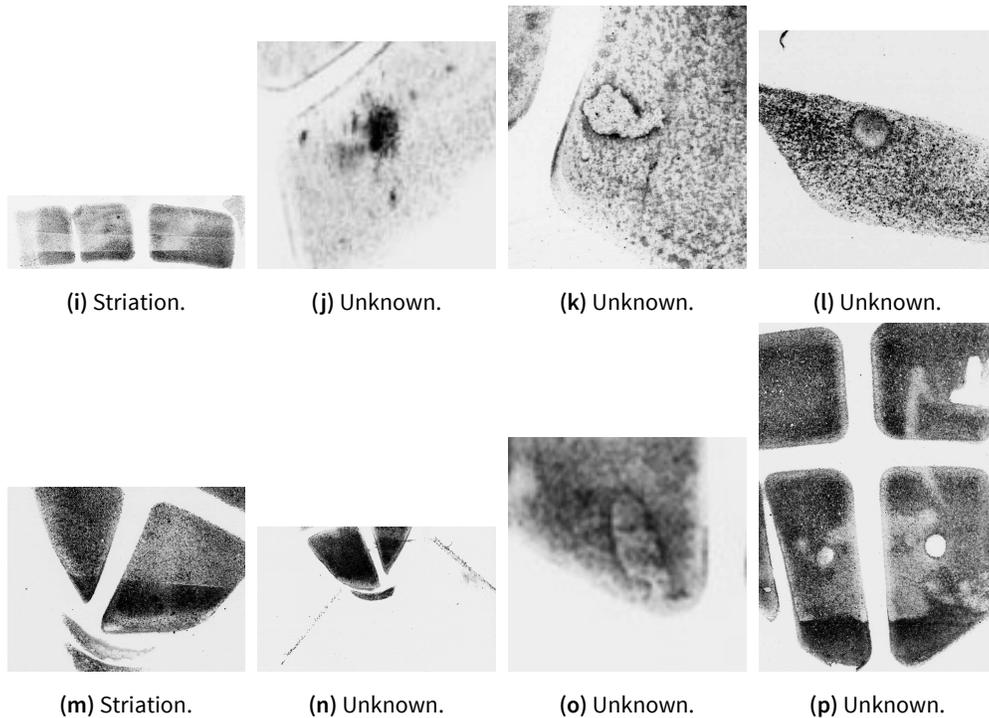


Figure 2.2: Crops of the dataset revealing specific instances of the issues detailed in Table 2.1.

2.2 Prior work

In the following, we briefly survey the existing methods that address noise found on shoeprints. It should be noted that these methods have been developed to analyse impressions or photographs of shoeprints taken from the crime scene, and then compare them against reference shoeprints in a database. The datasets used for reference shoeprints typically consist of a single instance of a given model of shoe. In such scenarios, finer details of the prints can be discarded. In contrast, our dataset consists of high-resolution impressions of one shoe taken in the laboratory every fortnight—thus capturing the development of low-level characteristics on the outsole.

In Kong et al. [52], their pre-processing steps for photographs taken from the crime scene include manually rotating and cropping the shoeprint out of the photograph; then binarising the image and applying median filters and morphological operators which delivers the input to their feature extraction algorithm. Srihari and Tang [8] discuss the various approaches to enhancing image quality including thresholding and edge detection; they compare the performance of popular methods in each approach and evaluate the results. Li et al. [53] developed a preliminary pre-processing

stage as part of their shoeprint retrieval algorithm. They smooth the crime scene image with a low-pass filter, perform an adaptive binarisation, followed by morphological operators to arrive at the final image. Li and Wang [56] have used noise filtering, contrast enhancement, and edge detection as preliminary steps to their shoeprint positioning algorithm. Alizadeh and Kose [58] eliminate noise by median filtering followed by thresholding, rotation, and scaling of the shoeprint. Wang et al. [60] have detailed a manifold ranking method for shoeprint retrieval which involves thresholding, rescaling, and rotating as pre-processing steps. Zhang and Allinson [29] have smoothed noisy figures using anisotropic filters. Dardi et al. [44] have employed histogram equalisation and edge detection as part of their shoeprint retrieval pipeline. Jing et al. [40] transform colour figures into grayscale, apply block-based noise removal using a Gaussian low-pass filter, extract edges using a Sobel filter, and finally apply principal component transform to rotate the figures. Luostarinen and Lehmußola [62] conducted an independent empirical evaluation of automatic shoeprint retrieval methods and concluded that novel pre-processing methods were still desirable.

Additionally, we attempted to employ state-of-the-art image denoising filters such as BM3D [82] and non-local means [83]. We found these methods to deliver no discernable improvement in the image quality of our wear-and-tear dataset.

The survey of the existing literature did not yield any suitable methods to fit our purposes. Applying these methods to our dataset resulted in a clearer denoised image, while also destroying the low-level features of the outsole that characterise the wear-and-tear pattern. These features are essential for the task of studying and modelling the wear pattern. The existing methods have been designed as pre-processing methods that fit into larger shoeprint retrieval systems and denoise shoeprints at the cost of image clarity and definition. Our principal concern in developing a denoising method for this dataset was to efficiently mitigate noise, while maintaining high-level features of the outsole and the low-level wear patterns. The denoising method we developed to satisfy these unique constraints is described next.

2.3 Proposed denoising methodology

Starting with the digital negatives, each of the 52 figures were processed individually to be denoised. In the first step, the complement of the negative was taken. The image was then examined to determine if its intensity needed to be adjusted. Due to variations in the amount of pressure applied while capturing the impression, some of the figures are of a weaker intensity than others. After analysis of the image and its histogram, the intensity of the image was scaled to the full 0–255 grayscale spectrum, if required. A binary mask of this image was then created using the adaptive

thresholding algorithm, using a sensitivity level parameterised by a . Due to the variation in overall intensity levels among the impressions, a slightly aggressive value of a ensured that the maximum amount of the shoeprint was captured, but this came with the cost of capturing much of the noise present in the background of the image. To deal with the additional noise in the binary mask, noisy regions were marked using region-of-interest (ROI) polygons and then filtered out. This step ensured that the shoeprint mask remained intact while background noise was filtered.

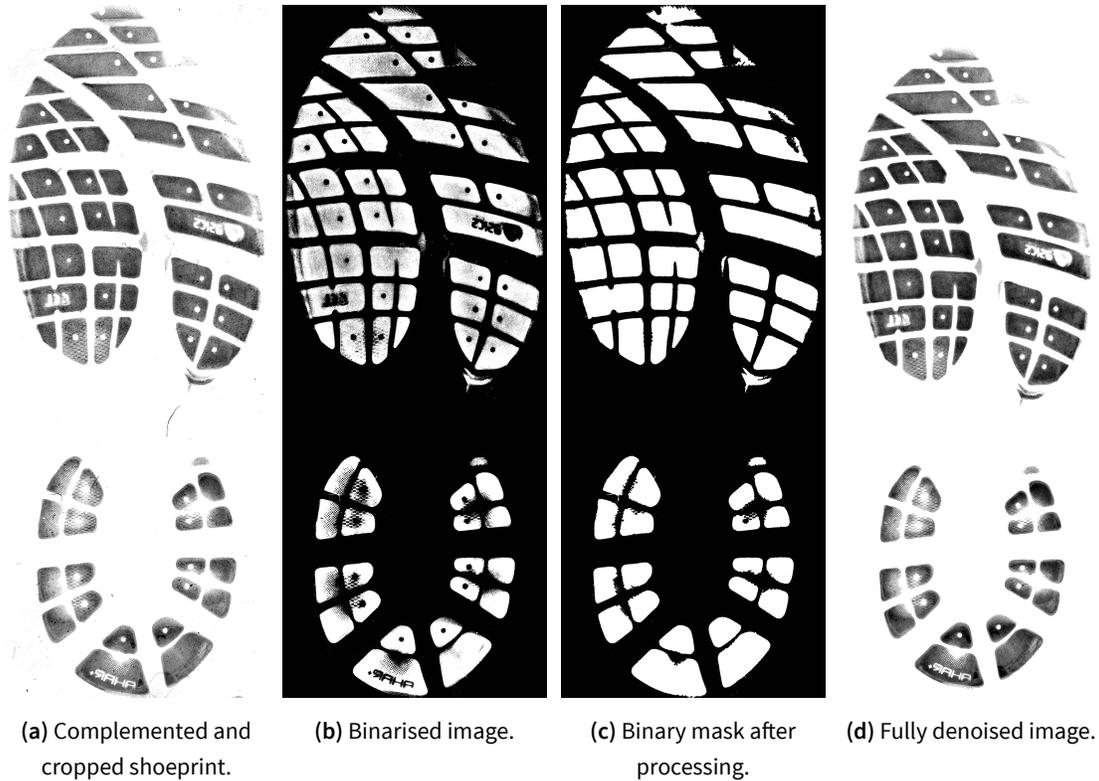


Figure 2.3: Print of the right shoe on week 18 through various stages of the proposed denoising method.

Using a disk-shaped structuring element with a logical neighborhood of β , the mask was then dilated to fill in the block regions of the outsole. Holes inside connected components were also filled in. This dilation process was repeated 4–7 times, as deemed necessary. Understandably, this process also tends to inflate background noise. The next step was therefore to erode the mask with a square-shaped structuring element using a logical neighborhood of size γ . Erosion was performed 3–5 times, depending upon the image. The shape and size of structuring elements were determined via experimentation. At this stage of the process, the logical mask clearly defines the pixels of the

original image that comprises the shoeprint. Using this mask we set the background of the image to be pure white, leaving the shoeprint untouched. The outputs of these steps are given in Figure 2.3.

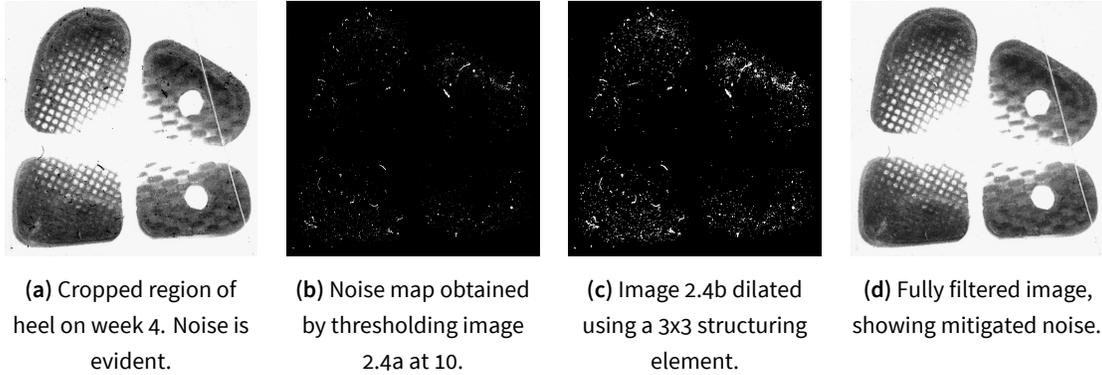


Figure 2.4: Intermediate stages of noise filtering highlighted using left shoe on week 4.

The convex hull of each block feature was computed and converted into its own binary mask. This allowed for each block feature to be processed in isolation. The next step was to denoise the image of the debris contaminants. After an exhaustive search of methods to identify the noise in a given region of the image, the best method was determined to be a simple thresholding to obtain a noise map defining pixels of a block feature as noise, parameterised by δ . Once the noise map was obtained, we dilated it using a disk-shaped structuring element of size ϵ . This step allows us to capture the borders of noise patterns while filtering. Using the noise map as a ROI mask, the block feature was then filtered using an averaging filter of size ζ . This process is then repeated for each block feature present in the outsole, to deliver a shoeprint impression free of obstructions. Figure 2.4 illustrates the steps involved in denoising. Each image is written back to disk as an uncompressed `.tif` file.

Table 2.2: Parameters from our method and the values optimised for our dataset.

Parameter	Value
α	0.6
β	7×7
γ	3×3
δ	10
ϵ	3×3
ζ	50×50

Despite the great care taken in the denoising process, the automated steps can sometimes leave noise present on the binary mask. Regions of the background may also be regarded as foreground in the thresholding process. Therefore we touch-up each image the dataset by manually examination, removing any noise that remained after filtering. Our semi-automated method was efficient enough that manual touch-ups were only required for less than 5% of the dataset. Finally, the dataset was padded to a uniform 13750×5500 size and registration was performed to align all of the prints. A block diagram outlining these steps is shown in Figure 2.5. The parameters we discovered to best fit this methodology with our dataset are given in Table 2.2.

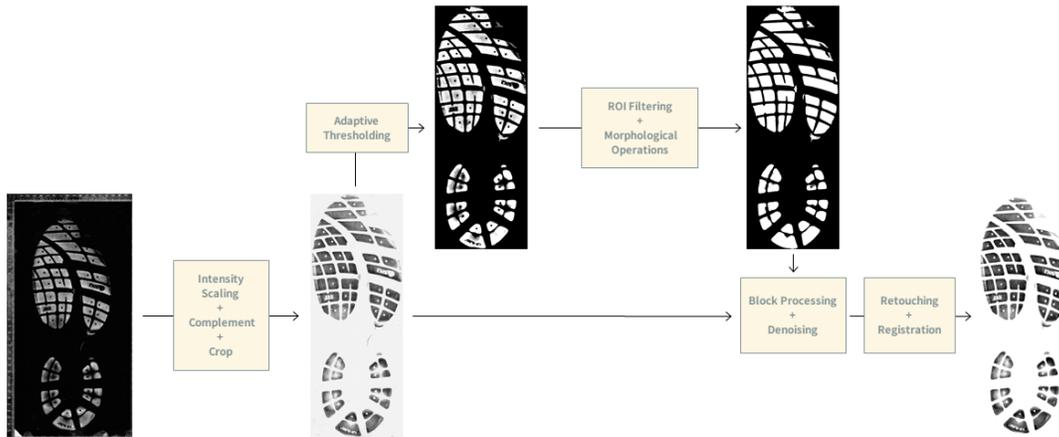


Figure 2.5: A block diagram outlining the processing steps in the proposed method.

2.4 Evaluations

A subjective assessment of the post-processed dataset determined it to be sufficiently adequate for our purposes. The proposed denoising method identified the noisy pixels and altered it to be near inconspicuous—leaving the outsole regions that show wear-and-tear untouched.

To perform an objective evaluation, we compared our method against the baselines of median filtering and adaptive Wiener filtering, using the metric of *Structural Similarity Index* (SSIM) [84], defined in (2.1).

$$SSIM(f, g) = l(f, g)c(f, g)s(f, g), \quad (2.1)$$

where

$$\begin{cases} l(f, g) = \frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}, \\ c(f, g) = \frac{2\sigma_f\sigma_g + C_2}{\sigma_f^2 + \sigma_g^2 + C_2}, \\ s(f, g) = \frac{\sigma_{fg} + C_3}{\sigma_f\sigma_g + C_3}, \end{cases}$$

wherein l , c , and s denote the luminance, contrast, and structure comparison functions respectively. The term f denotes the noisy image and g the denoised image. μ and σ denote mean and standard deviation of image luminance and contrast, respectively. σ_{fg} is the covariance between f and g . C_1 , C_2 , and C_3 are positive constants employed to avoid a null denominator.

The SSIM index is a positive value $\in [0, 1]$ where 0 denotes no correlation and 1 denotes $f = g$. From the results obtained after processing all figures in our dataset, we ascertain that our denoising method can improve on the SSIM score of median filtering by up to 0.0465. The results are given in Table 2.3.

Table 2.3: Mean and standard deviation of SSIM scores computed over all 52 images.

	Proposed Method	Median Filtering	Wiener Filtering
Mean	0.9143	0.8678	0.8868
STD	0.0281	0.0219	0.0250

$$PSNR(f, g) = 10 \log_{10}(255^2 / MSE(f, g)) \quad (2.2)$$

Additionally, we compared *Peak Signal-to-Noise Ratio* (PSNR, defined in (2.2)) scores against the same baselines. Once again, f and g denote the noisy and denoised figures, respectively, and MSE is the mean squared error—i.e. MSE between f and g . Results are given in Table 2.4.

Table 2.4: Mean and standard deviation of PSNR scores computed over all 52 images.

	Proposed Method	Median Filtering	Wiener Filtering
Mean	20.9197	20.4064	20.7873
STD	1.8926	2.0890	1.9608

PSNR scores obtained using our method see an average improvement of 0.1324 over the baseline methods. Furthermore, a subjective examination of the processed figures using the proposed method compared against the baseline filters shows that our method is much more efficient at maintaining the finer details we desire.

* * *

Having surveyed the related literature and found little suited to the task of denoising a footwear impression while maintaining low-level characteristics, we designed a denoising method to fit the unique requirements of our dataset of shoeprints. This dataset captures the evolution of wear characteristics and can be used for their study. In the next chapter, we discuss the method we use to model wear, our CNN architecture, and the results of our experiments.

All models are wrong, but some are useful.

George Box

3

Modelling Wear Patterns

How does the outsole of a pair of shoes change over time? The forensic examiner’s interpretation of the shoeprint and its admissability as evidence is built through their years of experience in studying shoeprints and the individualising characteristics that contribute to the wear pattern. Such knowledge is notoriously hard to quantify. Deep learning models have made large strides in developing computational representations of domains with these traits—including in the field of forensics—where they have been successfully applied to diverse problems like the recognition of finger vein patterns [85], orientation in iris images [86], and touchless palmprint recognition [87].

In this chapter we present a CNN architecture that performs a pixel-wise prediction of footwear outsole appearance. Our core contributions presented here are as follows—(i) we apply a CNN model to learn and predict the wear formation on our dataset of shoeprints, and (ii) an alternate model that is able to reconstruct the outsole back to its original state on a given week within a time-frame of one year.

In the following sections we first survey the related literature in the domains of deep learning and forensics & shoeprints. We proceed to detail our methodology for modelling, describe the experiments done with our models, and finally, we analyse the results of our experiments.

3.1 Related work

Using our dataset of 52 shoeprints, described in Chapter 2, we wish to learn a model of the wear pattern captured within. Once trained, this model should be capable of extrapolating the wear pattern when given a new shoeprint as input. Fundamentally, we approach this as an image-to-image regression task.

The literature contains many successful applications of deep learning to these types of dense prediction tasks; such as image in-painting [88, 89], super-resolution [90], denoising [91], and image recovery from compressed representations [92]. Deep neural networks and their convolutional variants have established state-of-the-art performance over nearly all facets of computer vision tasks. One of the primary advantages of using DNNs is their ability to learn end-to-end mappings without the use of image priors, or the explicit engineering of features.

Our dataset shows the life and wear of a pair of shoes through impressions captured at evenly spaced intervals of time. To the best of our knowledge, this is the first time such a dataset has been used in the literature of deep learning. A closely related problem is video frame generation/prediction [93] that involves operations on inputs in the spatial domain, while simultaneously capturing correlations in the temporal domain. Notably, in video frame prediction, one has access to an extensive amount of data by using each frame in the video sequence as a datapoint. Finn et al. [94] use a combination of convolutional and LSTM layers to model pixel motion and optical flow. They introduce a dataset with 1.5 million video frames and a model that predicts video sequences up to 1 second in the future. Our dataset is significantly smaller in size.

3.2 Proposed modelling methodology

3.2.1 Convolutional neural networks

In this section we provide a brief mathematical background to the operation of the convolutional neural network which forms the basis of our modelling method. We start by describing a feedforward neural network (more details can be found in Goodfellow et al. [95], Chapters 6 and 9).

A learning algorithm in machine learning is essentially a mathematical function that learns a mapping between inputs x , and outputs y . Neurons are a machine learning paradigm loosely modelled after the operation of a biological brain. The artificial neuron takes the mathematical form shown in (3.1).

$$y = f\left(\sum_i w_i x_i + b\right), \quad (3.1)$$

wherein x_i represents a set of inputs that are parameterised by a set of weights w_i , and added to a bias offset b .

This linear combination is fed into what is typically a non-linear *activation function*, $f(z)$ —where $z = \sum_i w_i x_i + b$. Depending on the task or architecture, $f(z)$ is most commonly either the rectified linear unit ($\text{ReLU}, f(z) = \max(0, z)$), or the sigmoidal function ($f(z) = 1/(1 + \exp(-z))$). The task in ML is to learn the right values for the weights w_i and biases b_i that fully approximate a set of y values when given a set of known x_i values.

Stacking neurons into a sequential architecture forms an artificial neural network (ANN), which is capable of modelling deeper representations in more complex problem domains. ANNs consist of sequential layers of neurons where each neuron is independent of each other, with its own set of weights and biases. A neural network is then trained by feeding inputs x into one end, obtaining predictions y from the other end, computing the error of the predictions against the known values (a.k.a. *ground-truth*) by a pre-defined loss function, and finally propagating the error signal backwards through the network using the backpropagation algorithm [96]. The training process involves finding the optimal set of weights for each neuron so that the network best approximates the ground-truth, and is performed automatically through a convex optimiser like gradient descent.

Computer vision tasks involve inputs that are typically images—matrices that define pixel intensities, and the target outputs are either a class label for tasks such as object detection and image classification; or another image in the case of problems like denoising, in-painting, etc. In the latter, both inputs and outputs are two-dimensional matrices, therefore we define the neuron in the i th row and j th column of a layer as:

$$y_{ij} = f\left(\sum_k \sum_l w_{ij,kl} x_{kl} + b_{ij}\right), \quad (3.2)$$

where x_{kl} represents pixel intensities of the input image in the NN's first layer. ANNs have no limitations on $w_{ij,kl}$, allowing weights of a neuron to vary independently of other neurons. For images, this implies the loss of spatial context from one layer to the next.

Convolutional neural networks vary from ANNs in that neurons are connected locally to the output volume of the previous layer [97], performing discrete convolutions in each layer between the input volume and a *filter* or kernel matrix, parameterised by the weights. Convolutional layers embed spatial context of images directly into the architecture of the network. Additionally, sharing the filter weights between neurons requires only a small number of them to be stored in memory, greatly simplifying the training process. This weight sharing also exploits the property that weights useful to one section of an image might also be useful to another. Traditional CNN architectures

make use of pooling layers which discard non-essential information in the image to deliver a fixed-length vector for classification purposes.

The convolution operation generates *feature maps* that represent the presence of certain features in the image [95]. Because this operation exhibits the property of equivariance, a change in the location of the feature in the input will correspondingly change its location in the output. In images, the first layers of a CNN learn to detect edges. Sharing these parameters across the entire image allows the network to detect edges throughout the input while keeping the parameter count to a minimum. With shoeprints, these edges define high-level patterns like blocks, lines, and dots in addition to low-level features like scratches, tears, and cuts.

3.2.2 CNN architecture

We approach the problem of modelling wear patterns by implementing a CNN architecture in the style of an auto-encoder, inspired by the work in Tatarchenko et al. [98] and Vukotić et al. [99]. This architecture consists of three branches—an *encoder* $F_\theta(\cdot)$ that takes as input a shoeprint image X , a *delta branch* $G_\omega(\cdot)$ that encodes a representation of time from a parameter Δ_t , and a *decoder* $H_\beta(\cdot)$ that learns an upsampling function to predict the wear pattern. These branches take the form shown in (3.3).

$$\begin{aligned}
 F_\theta &= \sigma(X * \theta) \equiv f, \\
 G_\omega &= \sigma(\Delta_t \cdot \omega) \equiv g, \\
 H_\beta &= \sigma(f \uplus g *' \beta), \quad (3.3)
 \end{aligned}$$

where σ represents an activation function—ReLU or sigmoid, $*$ represents the convolution operation, \uplus the concatenation of two tensors, and $*'$ the transpose convolution. Bias terms are omitted for notational convenience.

The encoder is made up of 5 convolutional layers that act as feature extractors by performing discrete convolutions over the input image, with an increasing depth of feature maps. We double feature maps with each layer, going from 32 in the first layer, to 512 in the last convolutional layer. The delta branch consists of 2 fully connected layers; the output of this branch is reshaped and concatenated with the output of the last convolutional layer. This tensor is then fed into the 5 transpose convolutional layers of the decoder, which successively upsample the extracted feature maps and the output of the delta branch to produce an output of the same dimensions as the input.

Transpose convolutional layers are used here as a learnable upsampling function, as opposed to

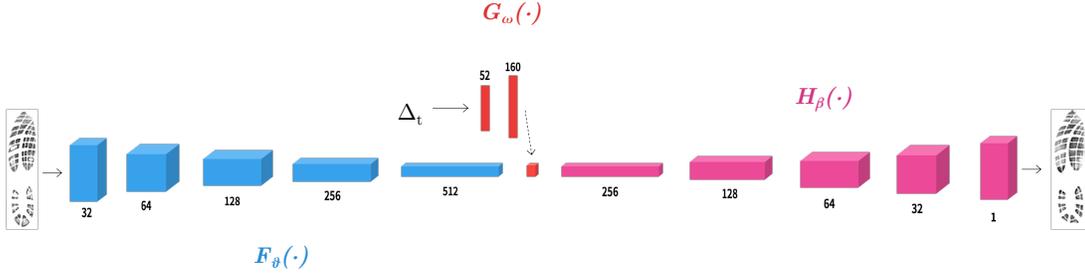


Figure 3.1: A visualisation of our proposed CNN architecture.

a fixed upsampling function (such as bilinear) in combination with 2D convolutions, as frequently seen in the literature. We discard pooling layers as traditionally seen in convolutional architectures since our denoising method—detailed in Chapter 2 removes redundant information in the image in the pre-processing stage.

The parameters of the network are updated by minimising the squared error loss (3.4):

$$L = \frac{1}{n} \sum_{i=1}^n \|H_\beta(F_\theta(X_i) \uparrow G_\omega(\Delta_t)) - Y_i\|_2^2 \quad (3.4)$$

where n is the number of training images presented to the network in one epoch, $X_i \in \mathbb{R}^2$ is the i th image, and Y_i the ground-truth image that corresponds to the input.

A visualisation of our architecture is shown in Figure 3.1. Using this architecture we have trained two models:

Moving forward—outsole wear prediction

Model A is designed to extrapolate wear patterns present in the shoeprint and form a prediction of what they might look like after a given period of time, denoted by Δ_t . The input image is presented at current relative time $t_o = 0$. We train this model to predict the appearance of the input, after the elapsed time Δ_t , where $\Delta_t \in [0, 52]$. Δ_t is incremented in steps of 2 to maintain consistency with the timeframe captured in our dataset.

Formally we train the model by feeding inputs as batches of $\{X, \Delta_t, Y\}$ tuples, where X represents the input image centered at a current relative time $t_o = 0$, Δ_t represents the desired temporal displacement, and Y represents the ground truth shoeprint image after the desired temporal displacement, $t_o + \Delta_t$.

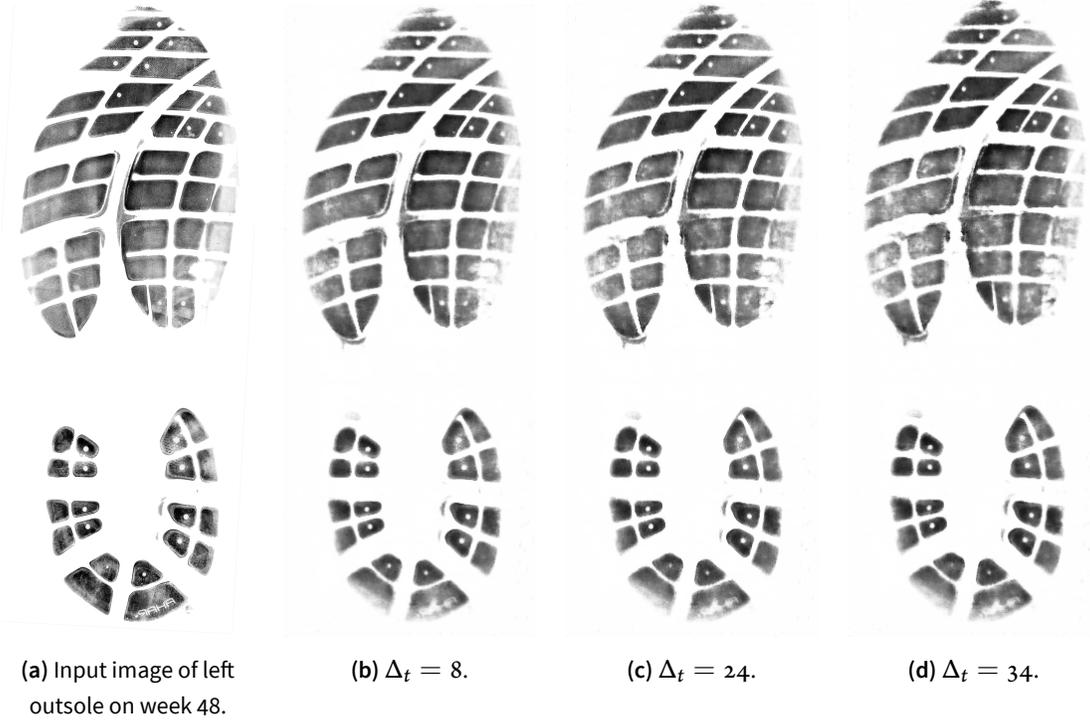


Figure 3.2: Predictions of Model A, given week 48 as input and a range of values for Δ_t .

Moving backward—outsole reconstruction

Model B is one that reconstructs the input shoeprint back to its state on any given week in a time-frame of one year. For this task, we use the same architecture as Model A. The difference here lies in how we design the Δ_t parameter. Here, Δ_t is represented as a logical vector $\in \mathbb{R}^{52 \times 1}$; wherein each element represents a week of the year, taking a value in $\{0, 1\}$, such that the desired week corresponding to the ground truth Y is represented as 1, and all other weeks represented as 0.

Once again we train the model by presenting $\{X, \Delta_t, Y\}$ tuples, and design the logical Δ_t vector in increments of 2 to correspond with the fortnightly nature of our captured dataset.

3.3 Experiments

3.3.1 Model training

Model training is performed by dividing the dataset of 52 images into an 80/20 training/test split. The first 42 images—both left and right outsoles—are used to train the model in conjunc-

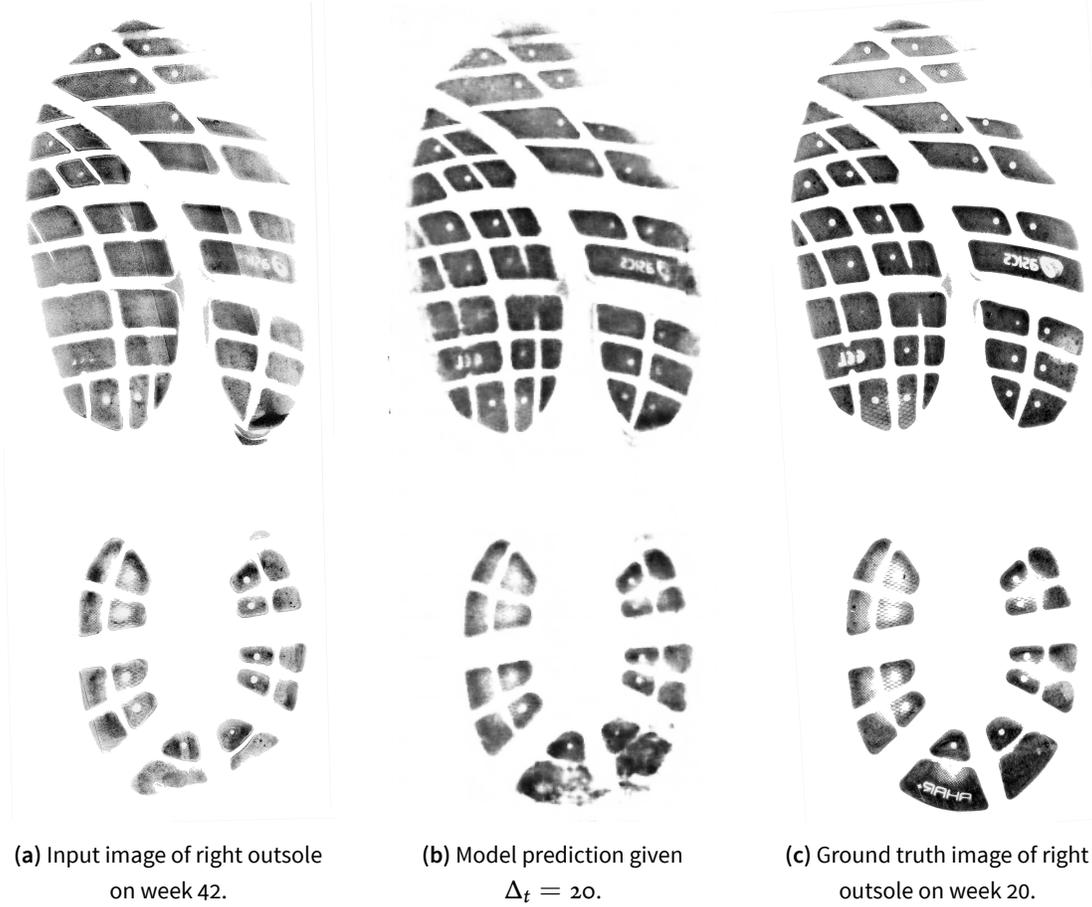


Figure 3.3: Outsole reconstruction predicted by Model B.

tion with a MSE loss and the Adam optimisation algorithm [100]. For training Model B, the split is reversed—i.e. we use the last 42 images for training, and test with the remaining images in our dataset that capture the start of our timeline. We use a learning rate of $1e - 5$ and train for 10,000 epochs. Activation functions throughout are the ReLU; except for the last layer which uses a sigmoid function to obtain outputs $\in [0, 1]$. We have found the same hyperparameters to be effective for both models.

Our dataset is composed of 52 grayscale images with a resolution of 13750×5500 . Fitting this dataset into memory during training required downsampling it to 640×256 . We train both models end-to-end from random initialisation. Alternative learning rates, initialisation schemes, optimisers, and loss functions were evaluated before settling on the above.

Figure 3.2 shows a sample of predictions from Model A. Figure 3.3 shows a reconstruction sample produced by Model B, along with the ground-truth image. Further information about model experiments, implementation, and source code are given in Appendix A.

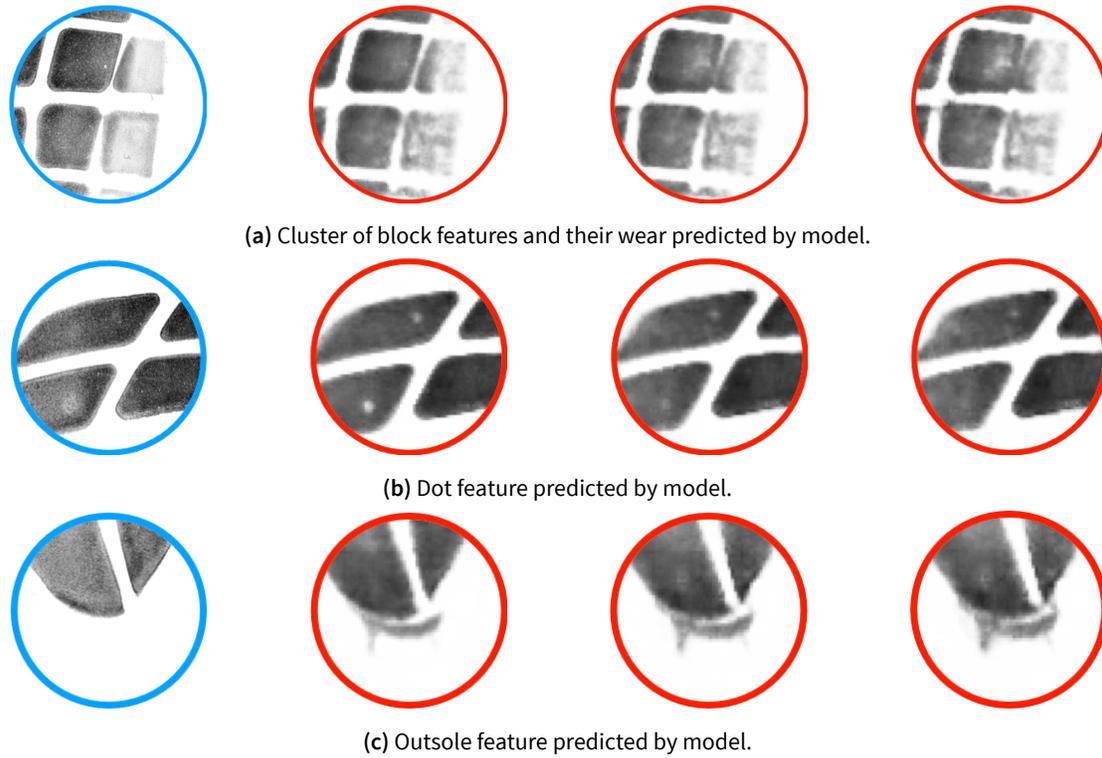


Figure 3.4: Highlights of relevant regions of predictions shown in Figure 3.2. Encircled in blue is the input to the model of the left outsole on week 48. Circled in red are the predicted wear patterns of the model, given Δ_t values of 8, 24, and 34, respectively from left to right.

3.3.2 Results

Our network successfully learns to model the high-level wear pattern embedded in the shoeprints. From observing the outputs, it is evident that the models have formed an internal representation sufficiently capable of predicting the wear pattern found in the dataset. Relevant regions of Figure 3.2 have been cropped and highlighted in Figure 3.4. Similarly, Figure 3.5 consists of crops of Figure 3.3. We have compared the predictions from the model against the ground-truth images from our dataset and noted the observations given below.

- In Figure 3.4a we see a cluster of four *block features* on the right edge of the outsole. In the

model’s predictions, we see them degrade and eventually merge in the final prediction, $\Delta_t = 34$. From the ground-truth image of week 52, we confirm that this change has indeed occurred; although clearly the model’s estimation of 20 weeks is far off from the reality of this eventual merger materialising in 4 weeks.

- Figure 3.4b shows two *dot features* visible in the first prediction of $\Delta_t = 8$ —note that these two features are not present in the input image of week 46. This feature is present throughout the outsole on many of the block features but have disappeared through wear-and-tear. It also happens to be visible in this exact region in all of the training images—weeks 0 through 42—but had eroded from the outsole by the time the input shoeprint was captured.

Interestingly, in the model’s latter two predictions— $\Delta_t = 24, 34$ —we see the feature degrade and eventually disappear, in line with the ground-truth; showing that the model has learned the wear development on this and similar features, despite consistently observing the dot features in this region throughout the training images.

- Figure 3.4c highlights a *ridge feature* seen in all the predictions, but not in the input. We verify through our dataset that this is in fact a feature of the outsole, seen in roughly half of the images in the training set, but is missing in the input image. Note how the model’s predictions show this ridge growing progressively larger in size, as the outsole erodes.
- In the outsole reconstructions of Model B, we see the successful reproduction of the *Asics* brand logo (Figure 3.5a), and the separation of block features that had merged through wear (Figure 3.5b). Also note the reconstruction of the feature that spells the word ‘GEL,’ imprinted by the manufacturer.
- The bottom region of the heel in the prediction seen in Figure 3.3b is blurry and poorly defined. This is due to the inconsistency of the appearance of this region in the dataset. During the data collection phase this region was either frequently occluded by fingerprints and debris, or ill-formed due to a lack of pressure between the outsole and the gel while collecting the impression. We surmise that this inconsistency in appearance is what has led the model to develop a fuzzy representation of this region.

From our evaluation of the results, we ascertain that our methodology is sufficient to capture the wear pattern from our dataset, and to accurately perform both outsole prediction and reconstruction.

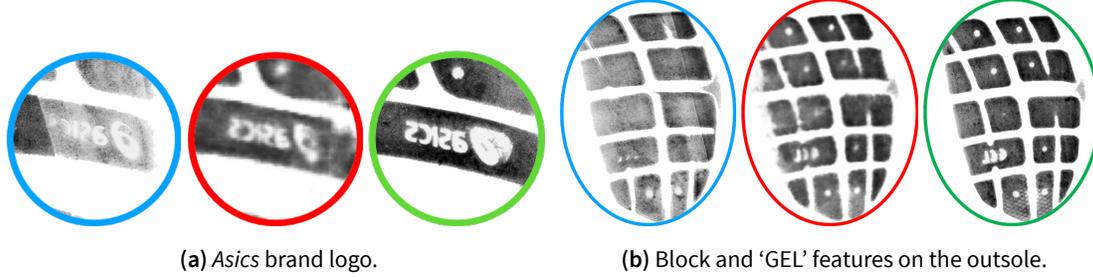


Figure 3.5: Highlights of relevant regions of outsoles shown in Figure 3.3. Encircled in blue is the input to the model of the right outsole on week 42. Circled in red is the predicted reconstruction of the model for week 20, and circled in green is the ground-truth image of week 20 from the dataset.

We also note that our network is handicapped by a lack of training data. In the era of deep learning, where models are routinely trained with millions of datapoints, we have sufficed with a meagre 52 images. Despite the size of the dataset, each pixel in the input image is a feature the model can learn from, and the 640×256 resolution of our training data is purely limited by processing power; allowing for a more robust model to be trained using higher resolution images. The generalisation ability of deep learning models can also benefit from an adequately sized dataset that fully captures the diversity of the problem domain.

Empirical evaluations are given in the next subsection.

3.3.3 Evaluations

For an objective evaluation of the performance of our models, we return to the metrics of SSIM (2.1) and PSNR (2.2), by comparing the predictions of the models against the ground-truth images from the validation dataset. Results are given in Tables 3.1 and 3.2.

Table 3.1: Mean and standard deviation of SSIM scores.

	Model A	Model B
Mean	0.8645	0.8596
STD	0.0381	0.0345

Going by the SSIM metric, our models are 86% accurate, while the PSNR score is 22dB on average.

* * *

In this chapter we presented a CNN architecture that can model the evolution of wear patterns on outsoles. We have shown that models trained with this architecture learn an accurate representation

Table 3.2: Mean and standard deviation of PSNR scores.

	Model A	Model B
Mean	22.0065	21.3681
STD	1.8930	1.8820

of the pattern of wear-and-tear found in our dataset of shoeprints, by applying it to predict the wear pattern on the outsole after a given temporal displacement, and by having it reconstruct the outsole back to its original state at a previous point in time. We addressed the drawbacks of the model and presented objective evaluations of its performance, which show the predictions of both models to be 86% accurate with average PSNR scores of 22dB.

This work adds to the scant literature on shoeprint wear patterns by presenting a computational model of outsole wear. The framework presented here can be applied to supplement the skills and expertise of the forensic examiner in their analysis of crime scene shoeprints, and to train the novice forensic scientist to hone their skills.

What is simple is always wrong. What is not is unusable.

Paul Valéry

4

Conclusion

Footwear outsoles present a potential method of identifying the individual(s) involved with a crime. By careful analysis of the geometric and accidental characteristics present on the outsole, the forensic footwear examiner may be able to link the outsole to the scene print. One challenge, however, is the accumulation of additional wear characteristics between the perpetration of a crime and the identification of suspect(s). In such situations, it falls on the forensic scientist to evaluate the suspect's outsole and determine if it matches the scene print while accounting for the formation of additional wear features. This involves the careful analysis of the outsole and requires intimate knowledge of the breadth of factors and variables that influence wear patterns.

We began by surveying the applications of computer science to the study of shoeprints and forensics. Identifying a knowledge gap around the formation of wear characteristics on footwear outsoles, we set out to model this process and better understand it. A unique dataset that captures the life of a pair of shoes were collected for this purpose and described in Chapter 2; along with a novel denoising method developed to account for accidental characteristics while removing unwanted artefacts. We proceeded to describe our CNN architecture in Chapter 3 and the two models trained using this architecture on our dataset.

While our denoising methodology was specifically designed to meet the constraints of our dataset, we believe it is capable of generalising to any high-resolution dataset of shoeprints where

the primary concern is to study the wear pattern. Having said that, we see some areas for improvement:

- **Parameters**—A total of 7 parameters were tuned to fit the methodology to our dataset. Reducing or eliminating the parameter count would make the methodology easier to use.
- **Interface**—Our denoising method is implemented in `MATLAB` and contains no graphical user interface (GUI) to aid the user. A GUI would lower the barrier to entry and allow the user to view and compare outputs from intermediate stages of the method.
- **Automation**—Given the size of our dataset, a semi-automated denoising method was sufficient; but when dealing with a larger set of images it may prove impractical to babysit the denoising process.

We have conclusively shown that CNNs can learn a sufficient internal representation of the wear pattern present on shoeprints, and use this representation to both predict and reconstruct the outsole's appearance. We see some possible avenues to expand on this work:

- **Longer timeline**—Collecting our dataset required a non-trivial time commitment and yet, the size of the dataset has proven to be the biggest hurdle to our task. A longer timeline would provide more data for the model to learn from, as well as additional data to evaluate the model's performance.
- **Diverse dataset**—Presently, our model is not capable of generalising beyond the *Asics*-brand shoeprint from our dataset. Expanding the training data to include different outsole patterns, worn by more individuals, in varied environments would enable a stronger model capable of generalising to unseen shoeprints. A robust model requires a robust dataset—one that adequately captures the diversity found in real life.
- **Advanced models**—Beyond prediction and reconstruction, experimenting with more advanced models and architectures could prove fruitful. An alternative architecture may be trained to predict different levels of wear given an input shoeprint, for example.

Modelling naturally occurring phenomena can present unique challenges; but can also provide deep insights into the process itself. The ideal model would deliver accurate predictions *and* the reasoning behind them—serving as an unbiased, objective point of reference to assist in the investigation. This thesis is presented as the first step towards machine understanding of the process that occurs between the outsole and the ground.



Model Experiments & Source Code

Machine learning models are only as good as the data they're trained on and it was apparent from the start that the size of our dataset represented a hurdle. One approach for dealing with a lack of data is a deep learning technique called *transfer learning* [101]. The idea here is to leverage pre-existing knowledge learned by a model in another domain of interest, and apply it to a related task—thereby reducing training time and the quantity of data needed. When applying this paradigm to a CNN, the weights of layers from a pre-trained model are used as a starting point for a new model and fine-tuned to the task. Transfer learning has been successfully applied to many popular CNN architectures and their pre-trained weights are provided online.

To apply this approach to our problem, we used the shoeprint datasets provided in Kortylewski et al. [59] and Richetelli et al. [63] to train a binary classifier—given an image, classify it as a shoeprint or not. To train this model we used the architecture of our encoder and attached a few fully connected layers to the end. The learned weights of this model were then transferred to a CNN architecture similar to the one detailed in Chapter 3—minus the delta branch—and further fine-tuned to wear prediction using our dataset.

What we found was that while transfer learning allowed the model to start forming coherent shoeprint images earlier in the training process, it did not have any noticeable effect on the ability of the network to model the wear pattern. This is most likely due to the wear prediction being

learned in the upsampling layers of the decoder. With no good reasons to use transfer learning, we discarded this approach in favor of end-to-end training of the network using our dataset exclusively.

Another strategy evaluated was to recursively feed the predicted output of the model back into it as an input, thereby getting iterative predictions of wear formation, with a fixed delta of 2. This approach very quickly devolved into noisy outputs, much like recursively taking a photocopy of a photocopy of a piece of paper would. The delta ‘trick’ we use not only expands on the number of unique training samples available, it also has the side effect of acting like a regulariser, preventing the model from over-fitting to the dataset.

Other architectural novelties considered were the addition of residuals, batch normalisation, and dropout layers. These techniques address over-fitting and/or optimised gradient flow in very deep neural networks, and did not prove to be of much use in our relatively shallow network.

* * *

Below, we provide the source code used to define and train our models. As noted in Chapter 3, one architecture is shared between our models, and is defined below in `model.py`. Our implementation is in Python, using the Keras [59] deep learning framework and the TensorFlow [102] backend. Training was monitored and visualised in TensorBoard. The dataset was pre-processed using scikit-image [103] by normalising it to the range $[0, 1]$, and resizing to 640×256 . The processed images were saved to disk in NumPy’s `.npy` file format, which greatly accelerated model training.

We defined custom data generators for each of our models by extending the base generator class exposed by Keras. These generators pick a random input file X , pick a random value for Δ_t , and load the corresponding ground-truth Y . This $\{X, \Delta_t, Y\}$ tuple is then fed to the optimiser during training as batches, the size of which is defined as a hyperparameter in `train.py`.

Source Code A.1: model.py

```
from keras.models import Model
from keras.layers import Input, Conv2D, Conv2DTranspose, Dense, Concatenate, Reshape

def DeltaEncoder(img_shape=(3438, 1375, 1), delta_shape=(1,), **kwargs):
    img = Input(shape=img_shape)

    # Encoder.
    x = Conv2D(32, (3, 3), strides=(2, 2),
              activation='relu', padding='same')(img)
    x = Conv2D(64, (3, 3), strides=(2, 2),
              activation='relu', padding='same')(x)
    x = Conv2D(128, (3, 3), strides=(2, 2),
              activation='relu', padding='same')(x)
    x = Conv2D(256, (3, 3), strides=(2, 2),
              activation='relu', padding='same')(x)
    x = Conv2D(512, (3, 3), strides=(2, 2),
              activation='relu', padding='same')(x)

    # Delta branch.
    delta = Input(shape=delta_shape)
    d = Dense(64, activation='relu', name='delta_fc1')(delta)
    d = Dense(160, activation='relu', name='delta_fc2')(d)
    d = Reshape((20, 8, 1))(d)

    x = Concatenate()([x, d])

    # Decoder.
    x = Conv2DTranspose(256, (3, 3), strides=(
        2, 2), activation='relu', padding='same')(x)
    x = Conv2DTranspose(128, (3, 3), strides=(
```

```
        2, 2), activation='relu', padding='same')(x)
x = Conv2DTranspose(64, (3, 3), strides=(
    2, 2), activation='relu', padding='same')(x)
x = Conv2DTranspose(32, (3, 3), strides=(
    2, 2), activation='relu', padding='same')(x)
x = Conv2DTranspose(1, (3, 3), strides=(
    2, 2), activation='sigmoid', padding='same')(x)

model = Model([img, delta], x, name='deltaencoder')

return model
```

Source Code A.2: generator.py

```
import keras
import random
import numpy as np

class DataGenerator(keras.utils.Sequence):
    'Generates data for the DeltaEncoder model'

    def __init__(self, path, list_IDs, batch_size=32, dim=(3438, 1375),
                n_channels=1, n_classes=1, shuffle=True):
        'Initialization'
        self.path = path
        self.list_IDs = list_IDs
        self.batch_size = batch_size
        self.dim = dim
        self.n_channels = n_channels
        self.n_classes = n_classes
        self.shuffle = shuffle
        self.on_epoch_end()

    def __len__(self):
        'Denotes the number of batches per epoch'
        return int(np.floor(len(self.list_IDs) / self.batch_size))

    def __getitem__(self, idx):
        'Generate one batch of data'
        # Generate indexes of the batch.
        indexes = self.indexes[idx * self.batch_size:(idx+1) * self.batch_size]

        # Find list of IDs.
        list_IDs_temp = [self.list_IDs[k] for k in indexes]
```

```

# Generate data.
X, deltaTMat, Y = self.__data_generation(list_IDs_temp)

return ([X, deltaTMat], Y)

def on_epoch_end(self):
    'Updates indexes after each epoch'
    self.indexes = np.arange(len(self.list_IDs))
    if self.shuffle == True:
        np.random.shuffle(self.indexes)

def __data_generation(self, list_IDs_temp):
    # X : (n_samples, *dim, n_channels)
    'Generates data containing batch_size samples'
    # Initialize.
    X = np.empty((self.batch_size, *self.dim, self.n_channels))
    deltaTMat = np.empty((self.batch_size, 1))
    Y = np.empty((self.batch_size, *self.dim, self.n_channels))

    # Generate data.
    for i, ID in enumerate(list_IDs_temp):
        # Store input.
        im_x = np.load(self.path + list_IDs_temp[i])
        im_x = im_x / 255.
        X[i, ] = im_x

        # Pick a delta value at random.
        while True:
            delta = random.choice(np.arange(int(list_IDs_temp[i][:2].lstrip(
                '0') or '00'), 53, 2)) - int(list_IDs_temp[i][:2].lstrip('0') or '00')
            deltaWeek = delta + \
                int(list_IDs_temp[i][:2].lstrip('0') or '00')
            if deltaWeek != 6:

```

break

Get name of delta file and construct delta matrix.

```
deltaFile = str(deltaWeek).rjust(2, '0') + \  
    list_IDs_temp[i][2:3] + '.npy'  
deltaTMat[i, ] = np.int(delta)
```

Store ground truth.

```
im_y = np.load(self.path + deltaFile)  
im_y = im_y / 255.  
Y[i, ] = im_y
```

return X, deltaTMat, Y

Source Code A.3: train.py

```
import os
import keras
import datetime
from model import DeltaEncoder
from generator import DataGenerator
from keras.callbacks import TensorBoard, ModelCheckpoint

# Hyperparameters.
model_name = 'deltaencoder-mark-iii-mse-adam'
path = '/data/data/decoder-ntp/'
params = {'dim': (640, 256),
          'batch_size': 6,
          'n_channels': 1,
          'shuffle': True}
height, width = params['dim']
learning_rate = 1e-4
lr_decay = 0.0
partition = {}

# Setup train/test split.
dataset = sorted(os.listdir(path))
partition['train'] = dataset[:46] # 80% for train.
partition['validation'] = dataset[46:50] # 20% for val.

# Setup data generators and model.
train_generator = DataGenerator(path, partition['train'], **params)
validation_generator = DataGenerator(
    path, partition['validation'], dim=(height, width))
model = DeltaEncoder(input_shape=(height, width, 1), delta_shape=(1,))

keras.utils.print_summary(model)
```

```

print('Training ' + model_name)

# Compile.
model.compile(loss='mse', optimizer=keras.optimizers.Adam(
    lr=learning_rate, decay=lr_decay))

# Checkpoints and Logs.
timestamp = '{0:%Y%m%d-%H%M%S}'.format(datetime.datetime.now())
filepath = 'checkpoints/deltaencoder/{time}-{name}.hdf5'.format(
    time=timestamp, name=model_name)
checkpoint = ModelCheckpoint(
    filepath, monitor='val_loss', verbose=1, save_best_only=True)
tensorboard = TensorBoard(
    log_dir='logs/{time}-{name}'.format(time=timestamp, name=model_name))
callbacks_list = [tensorboard, checkpoint]

# Fit.
model.fit_generator(
    generator=train_generator,
    steps_per_epoch=100 // params['batch_size'],
    epochs=3000,
    callbacks=callbacks_list,
    validation_data=validation_generator,
    validation_steps=50 // params['batch_size'],
    use_multiprocessing=True,
    workers=8)

print('Done.')

```

Bibliography

- [1] B. E. McKnight, “The washing away of wrongs: forensic medicine in thirteenth-century China,” *VRÜ Verfassung und Recht in Übersee*, vol. 17, no. 1, pp. 114–115, 1983.
- [2] M. Benecke, “DNA typing in forensic medicine and in criminal investigations: a current survey,” *Naturwissenschaften*, vol. 84, pp. 181–188, 1997.
- [3] J. R. Robertson, *Forensic examination of hair*, CRC Press, 2002.
- [4] J. S. Buckleton, J. Bright, and D. Taylor, *Forensic DNA evidence interpretation*, CRC press, 2016.
- [5] R. Kaur and S. G. Mazumdar, “Fingerprint based gender identification using frequency domain analysis,” *International Journal of Advances in Engineering and Technology*, vol. 3, no. 1, pp. 295–299, 2012.
- [6] W. Vernon, “The development and practice of forensic podiatry,” *Journal of Clinical Forensic Medicine*, vol. 13, no. 6-8, pp. 284–287, Aug. 2006.
- [7] M. Liukkonen, H. Majamaa, and J. Virtanen, “The role and duties of the shoeprint/toolmark examiner in forensic laboratories,” *Forensic Science International*, vol. 82, no. 1, pp. 99–108, 1996.
- [8] S. N. Srihari and Y. Tang, “Computational Methods for the Analysis of Footwear Impression Evidence,” in *Computational Intelligence in Digital Forensics: Forensic Investigation and Applications*, A. K. Muda, Y.-H. Choo, A. Abraham, and S. N Srihari, Eds., pp. 333–383. Springer International Publishing, Cham, 2014.
- [9] W. J. Bodziak, *Footwear impression evidence: detection, recovery and examination*, CRC Press, 1999.
- [10] T. J. Napier, “Scene linking using footwear mark databases,” *Science & Justice*, vol. 42, no. 1, pp. 39–43, Jan. 2002.

- [11] Z. Geradts and J. Keijzer, "The image-database REBEZO for shoeprints with developments on automatic classification of shoe outsole designs," *Forensic Science International*, vol. 82, no. 1, pp. 21–31, 1996.
- [12] V. L. Naples and J. S. Miller, "Making tracks: the forensic analysis of footprints and footwear impressions," *The Anatomical Record*, vol. 279, no. 1, pp. 9–15, 2004.
- [13] R. S. Stone, "Footwear examinations: mathematical probabilities of theoretical individual characteristics," *Journal of Forensic Identification*, vol. 56, no. 4, pp. 577, 2006.
- [14] N. Richetelli, M. Nobel, W. J. Bodziak, and J. A. Speir, "Quantitative assessment of similarity between randomly acquired characteristics on high quality exemplars and crime scene impressions via analysis of feature size and shape," *Forensic Science International*, vol. 270, pp. 211–222, 2017.
- [15] I. W. Evett, J. A. Lambert, and J. S. Buckleton, "A Bayesian approach to interpreting footwear marks in forensic casework," *Science & Justice*, vol. 38, no. 4, pp. 241–247, 1998.
- [16] C. Champod, R. Voisard, and A. Girod, "A statistical study of air bubbles on athletic shoesoles," *Forensic Science International*, vol. 109, no. 2, pp. 105–123, 2000.
- [17] D. W. Vernon, *The functional analysis of shoe wear patterns: theory and application.*, Ph.D. thesis, Sheffield Hallam University, 2000.
- [18] C. G. Aitken and F. Taroni, *Statistics and the evaluation of evidence for forensic scientists*, vol. 16, Wiley Online Library, 2004.
- [19] K. Franke and S. N. Srihari, "Computational forensics: Towards hybrid-intelligent crime investigation," in *Third International Symposium on Information Assurance and Security*. IEEE, 2007, pp. 383–386.
- [20] M. J. Saks and J. J. Koehler, "The Coming Paradigm Shift in Forensic Identification Science," *Science*, vol. 309, no. 5736, pp. 892–895, 2005.
- [21] G. Alexandre, "Computerized classification of the shoeprints of burglars' soles," *Forensic Science International*, vol. 82, no. 1, pp. 59–65, 1996.
- [22] S. Mikkonen, V. Suominen, and P. Heinonen, "Use of footwear impressions in crime scene investigations assisted by computerised footwear collection system," *Forensic Science International*, vol. 82, no. 1, pp. 67–79, 1996.

- [23] S. Mikkonen and T. Astikainen, “Databased Classification System for Shoe Sole Patterns — Identification of Partial Footwear Impression Found at a Scene of Crime,” *Journal of Forensic Sciences*, vol. 39, no. 5, pp. 1227–1236, Sept. 1994.
- [24] C. W. Monckton, ““Shoe-fit” - a computerised shoe print database,” *IET Conference Proceedings*, pp. 86–89, Jan. 1995.
- [25] W. Ashley, “What shoe was that? The use of computerised image database to assist in identification,” *Forensic Science International*, vol. 82, no. 1, pp. 7–20, 1996.
- [26] A. Alexander, A. Bouridane, and D. Crookes, “Automatic classification and recognition of shoeprints,” in *Seventh International Conference on Image Processing and Its Applications*, 1999, pp. 638–641.
- [27] A. Bouridane, A. Alexander, M. Nibouche, and D. Crookes, “Application of fractals to the detection and classification of shoeprints,” in *Proceedings of the International Conference on Image Processing*, 2000, pp. 474–477.
- [28] P. de Chazal, J. Flynn, and R. B. Reilly, “Automated processing of shoeprint images based on the Fourier transform for use in forensic science,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 3, pp. 341–350, Mar. 2005.
- [29] L. Zhang and N. Allinson, “Automatic shoeprint retrieval system for use in forensic investigations,” in *UK Workshop On Computational Intelligence*, 2005, pp. 137–142.
- [30] H. Su, D. Crookes, and A. Bouridane, “Thresholding of noisy shoeprint images based on pixel context,” *Pattern Recognition Letters*, vol. 28, no. 2, pp. 301–307, 2007.
- [31] M. Pavlou and N. M. Allinson, “Automatic extraction and classification of footwear patterns,” in *International Conference on Intelligent Data Engineering and Automated Learning*. Springer, 2006, pp. 721–728.
- [32] M. Pavlou and N. M. Allinson, “Automated encoding of footwear patterns for fast indexing,” *Image and Vision Computing*, vol. 27, no. 4, pp. 402–409, 2009.
- [33] L. Ghouti, A. Bouridane, and D. Crookes, “Classification of shoeprint images using directional filter banks,” in *International Conference on Visual Information Engineering*, 2006, pp. 167–173.

- [34] H. Su, D. Crookes, and A. Bouridane, "Shoeprint Image Retrieval by Topological and Pattern Spectra," in *International Machine Vision and Image Processing Conference*, Sept. 2007, pp. 15–22.
- [35] D. Crookes, A. Bouridane, H. Su, and M. Gueham, "Following the Footsteps of Others: Techniques for Automatic Shoeprint Classification," in *Second NASA/ESA Conference on Adaptive Hardware and Systems*, Aug. 2007, pp. 67–74.
- [36] M. Gueham, A. Bouridane, and D. Crookes, "Automatic Recognition of Partial Shoeprints Using a Correlation Filter Classifier," in *International Machine Vision and Image Processing Conference*, Sept. 2008, pp. 37–42.
- [37] M. Gueham, A. Bouridane, D. Crookes, and O. Nibouche, "Automatic Recognition of Shoeprints using Fourier-Mellin Transform," in *NASA/ESA Conference on Adaptive Hardware and Systems*, June 2008, pp. 487–491.
- [38] G. Al Garni and M. Hamiane, "A novel technique for automatic shoeprint image retrieval," *Forensic Science International*, vol. 181, no. 1–3, pp. 10–14, 2008.
- [39] R. Xiao and P. Shi, "Computerized matching of shoeprints based on sole pattern," in *International Workshop on Computational Forensics*. Springer, 2008, pp. 96–104.
- [40] M. Q. Jing, W.-J. Ho, and L. H. Chen, "A novel method for shoeprints recognition and classification," in *International Conference on Machine Learning and Cybernetics*, 2009, pp. 2846–2851.
- [41] O. Nibouche, A. Bouridane, M. Gueham, and M. Laadjel, "Rotation Invariant Matching of Partial Shoeprints," in *13th International Machine Vision and Image Processing Conference*, Sept. 2009, pp. 94–98.
- [42] F. Dardi, F. Cervelli, and S. Carrato, "A texture based shoe retrieval system for shoe marks of real crime scenes," in *International Conference on Image Analysis and Processing*. Springer, 2009, pp. 384–393.
- [43] F. Cervelli, F. Dardi, and S. Carrato, "Comparison of footwear retrieval systems for synthetic and real shoe marks," in *2009 Proceedings of 6th International Symposium on Image and Signal Processing and Analysis*. IEEE, 2009, pp. 684–689.

- [44] F. Dardi, F. Cervelli, and S. Carrato, "A combined approach for footwear retrieval of crime scene shoe marks," in *3rd International Conference on Imaging for Crime Detection and Prevention*, Dec. 2009, pp. 1–6.
- [45] R. Wang, W. Hong, and N. Yang, "The research on footprint recognition method based on wavelet and fuzzy neural network," in *2009 Ninth International Conference on Hybrid Intelligent Systems*. IEEE, 2009, pp. 428–432.
- [46] P. M. Patil and J. V. Kulkarni, "Rotation and intensity invariant shoeprint matching using Gabor transform with application to forensic science," *Pattern Recognition*, vol. 42, no. 7, pp. 1308–1317, 2009.
- [47] W. Pei, Y. Y. Zhu, Y. N. Na, and X. G. He, "Multiscale Gabor Wavelet for Shoeprint Image Retrieval," in *2nd International Congress on Image and Signal Processing*, Oct. 2009, pp. 1–5.
- [48] Y. Tang, S. N. Srihari, H. Kasiviswanathan, and J. J. Corso, "Footwear Print Retrieval System for Real Crime Scene Marks," in *Computational Forensics*, H. Sako, K. Y. Franke, and S. Saitoh, Eds., Berlin, Heidelberg, 2011, pp. 88–100, Springer Berlin Heidelberg.
- [49] Z. Li, C. Wei, Y. Li, and T. Sun, "Research of shoeprint image stream retrieval algorithm with scale-invariance feature transform," in *International Conference on Multimedia Technology*, July 2011, pp. 5488–5491.
- [50] M. Hasegawa and S. Tabbone, "A Local Adaptation of the Histogram Radon Transform Descriptor: An Application to a Shoe Print Dataset," in *Structural, Syntactic, and Statistical Pattern Recognition*, G. Gimel'farb, E. Hancock, A. Imiya, A. Kuijper, M. Kudo, S. Omachi, T. Windeatt, and K. Yamada, Eds., Berlin, Heidelberg, 2012, pp. 675–683, Springer Berlin Heidelberg.
- [51] C.-H. Wei, Y. Li, and C.-Y. Gwo, "The Use of Scale-Invariance Feature Transform Approach to Recognize and Retrieve Incomplete Shoeprints," *Journal of Forensic Sciences*, vol. 58, no. 3, pp. 625–630, 2013.
- [52] X. Kong, C. Yang, and F. Zheng, "A Novel Method for Shoeprint Recognition in Crime Scenes," in *Biometric Recognition*, Z. Sun, S. Shan, H. Sang, J. Zhou, Y. Wang, and W. Yuan, Eds., Cham, 2014, pp. 498–505, Springer International Publishing.

- [53] X. Li, M. Wu, and Z. Shi, "The retrieval of shoeprint images based on the integral histogram of the gabor transform domain," in *International Conference on Machine Learning and Cybernetics*, 2014.
- [54] X. Min and G. Qi, "Automatic shoeprint classification based on DFT," in *12th International Conference on Signal Processing*, Oct. 2014, pp. 821–826.
- [55] S. Almaadeed, A. Bouridane, D. Crookes, and O. Nibouche, "Partial shoeprint retrieval using multiple point-of-interest detectors and SIFT descriptors," *Integrated Computer-Aided Engineering*, vol. 22, no. 1, pp. 41–58, 2015.
- [56] Y. Li and H. X. Wang, "A Secondary Positioning Algorithm of the Shoe Print," in *International Conference on Network and Information Systems for Computers*, Jan. 2015, pp. 560–563.
- [57] C.-Y. Gwo and C.-H. Wei, "Shoeprint retrieval: Core point alignment for pattern comparison," *Science & Justice*, vol. 56, no. 5, pp. 341–350, 2016.
- [58] S. Alizadeh and C. Kose, "Automatic retrieval of shoeprint images using blocked sparse representation," *Forensic Science International*, vol. 277, pp. 103–114, 2017.
- [59] A. Kortylewski, T. Albrecht, and T. Vetter, "Unsupervised footwear impression analysis and retrieval from crime scene data," *International Conference on Intelligent Information Processing*, pp. 644–658, 2015.
- [60] X. Wang, C. Zhang, Y. Wu, and Y. Shu, "A manifold ranking based method using hybrid features for crime scene shoeprint retrieval," *Multimedia Tools and Applications*, vol. 76, no. 20, pp. 21629–21649, Oct. 2017.
- [61] X. Wang, H. Sun, Q. Yu, and C. Zhang, "Automatic Shoeprint Retrieval Algorithm for Real Crime Scenes," in *Asian Conference on Computer Vision*, D. Cremers, I. Reid, H. Saito, and M.-H. Yang, Eds., Cham, 2015, pp. 399–413, Springer International Publishing.
- [62] T. Luostarinen and A. Lehmussola, "Measuring the accuracy of automatic shoeprint recognition methods," *Journal of Forensic Sciences*, vol. 59, no. 6, pp. 1627–1634, 2014.
- [63] N. Richetelli, M. C. Lee, C. A. Lasky, M. E. Gump, and J. A. Speir, "Classification of footwear outsole patterns using Fourier transform and local interest points," *Forensic Science International*, vol. 275, pp. 102–109, 2017.

- [64] T. L. Fruchtenicht, W. P. Herzig, and R. D. Blackledge, "The discrimination of two-dimensional military boot impressions based on wear patterns," *Science & Justice*, vol. 42, no. 2, pp. 97–104, Apr. 2002.
- [65] J. Skerrett, C. Neumann, and I. Mateos-Garcia, "A Bayesian approach for interpreting shoe-mark evidence in forensic casework: Accounting for wear features," *Forensic Science International*, vol. 210, no. 1, pp. 26–30, 2011.
- [66] N. K. Damary, M. Mandel, S. Wiesner, Y. Yekutieli, Y. Shor, and C. Spiegelman, "Dependence among randomly acquired characteristics on shoeprints and their features," *Forensic Science International*, vol. 283, pp. 173–179, 2018.
- [67] Y. Yekutieli, Y. Shor, S. Wiesner, and T. Tsach, "Expert assisting computerized system for evaluating the degree of certainty in 2d shoeprints," Tech. Rep. TP-3211, National Institute of Justice, 2012.
- [68] M. J. Cassidy, *Footwear Identification*, Public Relations Branch of the Royal Canadian Mounted Police Ottawa, 1980.
- [69] J. M. Wyatt, K. Duncan, and M. A. Trimpe, "Aging of shoes and its effect on shoeprint impressions," *Journal of Forensic Identification*, vol. 55, no. 2, pp. 181, 2005.
- [70] T. W. Adair, J. Lemay, A. McDonald, R. Shaw, and R. Tewes, "The Mount Bierstadt study: An experiment in unique damage formation in footwear," *Journal of Forensic Identification*, vol. 57, no. 2, pp. 199, 2007.
- [71] T. N. Moorthy and S. Chelliah, "Study on Stability of Individualizing Characteristics on Footwear Impression as Aging of Footwear for Forensic Investigation," *College Sadhana - Journal for Bloomers of Research*, vol. 6, no. 1, Aug. 2013.
- [72] N. D. K. Petraco, C. Gambino, T. A. Kubic, D. Olivio, and N. Petraco, "Statistical Discrimination of Footwear: A Method for the Comparison of Accidentals on Shoe Outsoles Inspired by Facial Recognition Techniques," *Journal of Forensic Sciences*, vol. 55, no. 1, pp. 34–41, 2010.
- [73] H. D. Sheets, S. Gross, G. Langenburg, P. J. Bush, and M. A. Bush, "Shape measurement tools in footwear analysis: A statistical investigation of accidental characteristics over time," *Forensic Science International*, vol. 232, no. 1, pp. 84–91, 2013.

- [74] R. Poplin, A. V. Varadarajan, K. Blumer, Y. Liu, M. V. McConnell, G. S. Corrado, L. Peng, and D. R. Webster, "Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning," *Nature Biomedical Engineering*, vol. 2, no. 3, pp. 158–164, 2018.
- [75] A. Silburt, M. Ali-Dib, C. Zhu, A. Jackson, D. Valencia, Y. Kissin, D. Tamayo, and K. Menou, "Lunar Crater Identification via Deep Learning," *Icarus*, vol. 317, pp. 27–38, 2018.
- [76] T. Perol, M. Gharbi, and M. Denolle, "Convolutional neural network for earthquake detection and location," *Science Advances*, vol. 4, no. 2, 2018.
- [77] W. Sun, D. Taniar, and T. Torabi, "Image Mining: A Case for Clustering Shoe prints," in *Database Technologies: Concepts, Methodologies, Tools, and Applications*, pp. 1552–1567. IGI Global, 2009.
- [78] V. Ramakrishnan and S. Srihari, "Extraction of shoe-print patterns from impression evidence using Conditional Random Fields," in *19th International Conference on Pattern Recognition*, Dec. 2008, pp. 1–4.
- [79] A. Kortylewski and T. Vetter, "Probabilistic Compositional Active Basis Models for Robust Pattern Recognition.," in *British Machine Vision Conference*, 2016.
- [80] B. Kong, D. Ramanan, and C. Fowlkes, "Cross-Domain Forensic Shoeprint Matching," in *British Machine Vision Conference*, London, Sept. 2017.
- [81] Y. Zhang, H. Fu, E. Dellandréa, and L. Chen, "Adapting Convolutional Neural Networks on the Shoeprint Retrieval for Forensic Use," in *Chinese Conference on Biometric Recognition*, J. Zhou, Y. Wang, Z. Sun, Y. Xu, L. Shen, J. Feng, S. Shan, Y. Qiao, Z. Guo, and S. Yu, Eds., Shenzhen, China, 2017, pp. 520–527, Springer International Publishing.
- [82] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [83] A. Buades, B. Coll, and J. Morel, "A non-local algorithm for image denoising," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 60–65.
- [84] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

- [85] W. Yang, C. Hui, Z. Chen, J. Xue, and Q. Liao, "FV-GAN: Finger Vein Representation Using Generative Adversarial Networks," *IEEE Transactions on Information Forensics and Security*, 2019.
- [86] A. Czajka, K. W. Bowyer, M. Krumdick, and R. G. VidalMata, "Recognition of Image-Orientation-Based Iris Spoofing," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 9, pp. 2184–2196, Sept. 2017.
- [87] A. Genovese, V. Piuri, K. N. Plataniotis, and F. Scotti, "PalmNet: Gabor-PCA Convolutional Networks for Touchless Palmprint Recognition," *IEEE Transactions on Information Forensics and Security*, 2019.
- [88] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. Efros, "Context Encoders: Feature Learning by Inpainting," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [89] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5505–5514.
- [90] C. Dong, C. C. Loy, K. He, and X. Tang, "Image Super-Resolution Using Deep Convolutional Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 295–307, Feb. 2016.
- [91] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, July 2017.
- [92] C. Dong, Y. Deng, C. Change Loy, and X. Tang, "Compression artifacts reduction by a deep convolutional network," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 576–584.
- [93] M. Mathieu, C. Couprie, and Y. LeCun, "Deep multi-scale video prediction beyond mean square error," in *International Conference on Learning Representations*, 2015.
- [94] C. Finn, I. Goodfellow, and S. Levine, "Unsupervised Learning for Physical Interaction through Video Prediction," in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds., pp. 64–72. Curran Associates, Inc., 2016.

- [95] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.
- [96] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, pp. 533 EP –, 1986.
- [97] Y. LeCun and Y. Bengio, “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, pp. 1995, 1995.
- [98] M. Tatarchenko, A. Dosovitskiy, and T. Brox, “Multi-view 3d models from single images with a convolutional network,” in *European Conference on Computer Vision*. Springer, 2016, pp. 322–337.
- [99] V. Vukotić, S.-L. Pinteá, C. Raymond, G. Gravier, and J. Van Gemert, “One-Step Time-Dependent Future Video Frame Prediction with a Convolutional Encoder-Decoder Neural Network,” in *International Conference on Image Analysis and Processing*. Springer, 2017, pp. 140–151.
- [100] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations*, 2015.
- [101] S. J. Pan and Q. Yang, “A Survey on Transfer Learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, pp. 1345–1359, 2010.
- [102] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: A System for Large-Scale Machine Learning,” in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, Berkeley, CA, USA, 2016, pp. 265–283, USENIX Association.
- [103] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, “Scikit-Image: Image Processing in Python,” *PeerJ*, vol. 2, pp. e453, June 2014.



Declaration

Name of candidate: Soonu Xavier Francis

This Thesis/Dissertation/Research Project entitled: Modelling Wear Patterns on Footwear Outsoles

is submitted in partial fulfillment for the requirements for the Unitec degree of Master of Computing

Principal Supervisor: Dr Hamid Sharifzadeh

Associate Supervisor: Angus Newton

Associate Supervisor: Dr Nilufar Baghaei

CANDIDATE'S DECLARATION

I confirm that:

- This Thesis/Dissertation/Research Project represents my own work;
- The contribution of supervisors and others to this work was consistent with the Unitec Regulations and Policies.
- Research for this work has been conducted in accordance with the Unitec Research Ethics Committee Policy and Procedures, and has fulfilled any requirements set for this project by the Unitec Research Ethics Committee.

Research Ethics Committee Approval Number: 2016-1022

Candidate Signature:

A handwritten signature in black ink, appearing to read 'Soonu Xavier Francis', written over a light grey signature line.

Date: 05 / 06 / 2019

Student ID: 1487947

Full name of author: Soonu Xavier Francis

ORCID number (Optional): 0000-0001-9793-9310

Full title of thesis/dissertation/research project ('the work'): Modelling Wear Patterns on Footwear Outsoles

Practice Pathway: School of Computing and Information Technology

Degree: Master of Computing

Year of presentation: 2019

Principal Supervisor: Dr Hamid Sharifzadeh

Associate Supervisor: Angus Newton

Associate Supervisor: Dr Nilufar Baghaei

Permission to make open access

I agree to a digital copy of my final thesis/work being uploaded to the Unitec institutional repository and being made viewable worldwide.

Copyright Rights:

Unless otherwise stated this work is protected by copyright with all rights reserved.
I provide this copy in the expectation that due acknowledgement of its use is made.

AND

Copyright Compliance:

I confirm that I either used no substantial portions of third party copyright material, including charts, diagrams, graphs, photographs or maps in my thesis/work or I have obtained permission for such material to be made accessible worldwide via the Internet.

Signature of author:

Date: 05 / 06 / 2019

