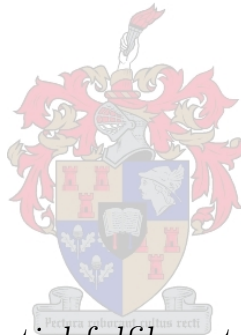


Modifying and generalising the Radon transform for improved curve-sensitive feature extraction

by

Carlien Fick



*Thesis presented in partial fulfilment of the requirements for
the degree of Master of Applied Mathematics in the Faculty
of Science at Stellenbosch University*

Supervisor: Dr J. Coetzer

Co-supervisor: Dr J.P. Swanepoel

December 2017

The financial assistance of the CSIR and ARMSCOR towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the CSIR and ARMSCOR.

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Name: Carlien Fick

Date: December 2017

Copyright © 2017 Stellenbosch University
All rights reserved.

Abstract

In this thesis a *novel* and *generic* feature extraction protocol that is based on the well-known *standard discrete Radon transform* (SDRT) is presented. The SDRT is traditionally associated with computerised tomography and involves the calculation of projection profiles of an image from a finite set of angles. Although the SDRT has already been successfully employed for the purpose of feature extraction, it is limited to the detection of *straight lines*.

The proposed feature extraction protocol is based on modifications to the SDRT that facilitate the detection of not only straight lines, but also *curved* lines (with various curvatures), as well as *textural* information. This is made possible by first constructing a novel appropriately normalised *multiresolution polar transform* (MPT) of the image in question. The origin of said MPT may be adjusted according to the type of features targeted. The SDRT, or the novel *modified discrete Radon transform* (MDRT) conceptualised in this thesis, is subsequently applied to the MPT.

The extraction of textural information based on different textural *periodicities* is facilitated by considering different projection angles associated with the MDRT, while the extraction of textural information based on different textural *orientations* is facilitated by specifying different origins for the MPT. The extraction of information pertaining to curved lines is made possible by specifying origins for the MPT that are located at different distances from the edge of the image in question – the SDRT is subsequently applied to a given MPT from a specific angle of 90° .

An existing system that only employs SDRT-based features constitutes a *benchmark*. Two novel texture-based systems, that target different textural *periodicities* and *orientations* respectively, are developed. A novel system, that constitutes a *generalisation* of the SDRT-based benchmark, and is geared towards the detection of different curved lines, is also developed.

The proficiency of the proposed systems is gauged by considering a data set that contains authentic handwritten signature images and *skilled* forgeries associated with 51 writers. All of the proposed systems outperform the SDRT-based benchmark. The improvement in proficiency associated with each individual texture-based system is *statistically significant*. The proficiency of the proposed systems also compares favourably with that of existing state-of-the-art systems within the context of offline signature verification.

Uittreksel

In hierdie tesis word ‘n *nuwe* en *generiese* kenmerk-onttrekkingsprotokol, wat op die bekende *gestandaardiseerde diskrete Radon-transformasie* (SDRT) gebaseer is, voorgelê. Die SDRT word tradisioneel met rekenaarmatige tomografie geassosieer, en behels die berekening van projeksie-profiel van ‘n beeld vanuit ‘n eindige versameling hoeke. Alhoewel die SDRT reeds vir kenmerk-onttrekking aangewend is, is dit beperk tot die opsporing van *reguit lyne*.

Die voorgestelde kenmerk-onttrekkingsprotokol is op aanpassings van die SDRT gebaseer, en fasiliteer die opsporing van benewens reguit lyne, ook *krom lyne* (met verskeie krommings), asook *tekstuur*-inligting. Dit word bewerkstellig deur eers ‘n nuwe korrek-genormaliseerde *multiresolusie-pooltransformasie* (MPT) op die betrokke beeld toe te pas. Die oorsprong van so ‘n MPT kan, na gelang van die tipe kenmerke wat geteiken word, aangepas word. Die SDRT, of die nuwe *aangepaste diskrete Radon-transformasie* (MDRT) soos gekonseptualiseer in hierdie tesis, word vervolgens op die MPT toegepas.

Die onttrekking van tekstuur-inligting op grond van verskillende tekstuur-*periodisiteit* word gefasiliteer deur verskillende projeksie-hoeke geassosieer met die MDRT te beskou, terwyl die onttrekking van tekstuur-inligting op grond van verskillende tekstuur-*oriëntasies* moontlik gemaak word deur verskillende oorspronge vir die MPT te spesifiseer. Die onttrekking van inligting rakende krom lyne word gefasiliteer deur oorspronge op verskillende afstande vanaf die rand van die betrokke beeld vir die MPT te spesifiseer – die SDRT word vervolgens op ‘n gegewe MPT vanaf ‘n spesifieke hoek van 90° toegepas.

‘n *Maatstaf* word gestel deur ‘n bestaande stelsel wat slegs SDRT-gebaseerde kenmerke gebruik. Twee nuwe tekstuur-gebaseerde stelsels, wat onderskeidelik verskillende tekstuur-*periodisiteit* en -*oriëntasies* teiken, word ontwikkel. ‘n Nuwe stelsel, wat gebaseer is op ‘n *veralgemening* van die SDRT-gebaseerde maatstaf, en gerig is op die opsporing van verskillende krom lyne, word ook ontwikkel.

Die vaardigheid van die voorgestelde stelsels word afgeskat deur ‘n datastel te beskou wat statiese handtekeninge en *hoë-kwaliteit* vervalsings, geassosieer met 51 skrywers, bevat. Al die voorgestelde stelsels vaar beter as die SDRT-gebaseerde maatstaf. Die vaardigheidsverbetering geassosieer met elke individuele tekstuur-gebaseerde stelsel is *statisties beduidend*. Die vaardigheid van die voorgestelde stelsels vergelyk ook goed met dié van bestaande stand-van-die-kuns-stelsels binne die konteks van statiese handtekeningverifikasie.

Acknowledgements

I would like to express my sincere gratitude to the following people and/or institutions for enabling me to successfully complete this study:

- My supervisor, Dr Hanno Coetzer, for his invaluable insight, guidance and support throughout my postgraduate studies. I am grateful and honoured to have been his student.
- My co-supervisor, Dr Jacques Swanepoel, for his input and guidance throughout this study.
- The CSIR and ARMSCOR, for their financial assistance.
- My family and friends, for their unconditional love and support.

Contents

Declaration	i
Abstract	ii
Uittreksel	iii
Acknowledgements	iv
Contents	v
List of Figures	vii
List of Tables	xiv
List of Acronyms	xv
Nomenclature	xvi
1 Introduction	1
1.1 Background and motivation	1
1.2 Key concepts	2
1.3 Scope and objectives	6
1.4 System design	7
1.5 Abbreviated results	10
1.6 Contributions	11
1.7 Thesis outline	13
2 Literature study	15
2.1 Introduction	15
2.2 SDRT for feature extraction	15
2.3 Texture-based recognition	19
2.4 Generic curve detection	21
2.5 Concluding remarks	22
3 Feature extraction	23
3.1 Introduction	23

3.2	The standard discrete Radon transform (SDRT)	24
3.3	A multiresolution polar transform (MPT)	25
3.4	Modification of the SDRT for texture detection	32
3.5	Generalisation of the SDRT for curved line detection	43
3.6	Feature normalisation	45
3.7	Concluding remarks	47
4	Feature matching, threshold selection and classifier ensembles	48
4.1	Introduction	48
4.2	Dynamic Time Warping (DTW)	49
4.3	Score normalisation	55
4.4	Threshold selection	55
4.5	Classifier ensemble construction	56
4.6	Classifier combination	57
4.7	Concluding remarks	57
5	Experiments	59
5.1	Introduction	59
5.2	Data	59
5.3	System parameters	60
5.4	Protocol	61
5.5	Results	61
5.6	Concluding remarks	69
6	Conclusion and future work	70
6.1	Conclusion	70
6.2	Future work	71
	Bibliography	79

List of Figures

1.1	Conceptualisation of the EER, which determines the selection of an appropriate global threshold based on the <i>optimisation</i> set for a specific trial.	5
1.2	A conceptualisation of the systems proposed in this thesis.	8
1.3	The four classifier ensembles constructed in this thesis.	10
3.1	Overview of the proposed feature extraction protocol.	24
3.2	Geometric interpretation of the SDRT (Toft, 1996).	25
3.3	Conceptualisation of the standard polar transform. The arc length Δs increases when the radius r increases, which leads to an insufficient sampling rate for large values of r . The largest radius considered is denoted by r_{\max}	26
3.4	Illustration of the deficiencies associated with the standard polar transform. (a) The original image. (b) The pixels that are considered (visited) when the standard polar transform is employed. (c) The reconstructed image using information obtained from the standard polar transform. Numerous inaccuracies are observed for large values of r . The centre of the image coincides with $r = 0$	27
3.5	Illustration of the benefits associated with the MPT. (a) The original image. (b) The reconstructed image using information obtained from the MPT. The slight inaccuracies for large values of r are the result of rounding errors. The centre of the image coincides with $r = 0$	28
3.6	Conceptualisation of the MPT which illustrates how more emphasis is placed on pixels closer to the origin ($r = 0$). Normalisation is therefore required.	29

- 3.7 A practical illustration of the rationale behind the normalisation of the MPT. **(a)** The original image. **(b)** The unnormalised MPT of the image in (a) for which the origin ($r = 0$) is located at the top right corner of (a) as indicated by the red dot. The pixel values closer to the origin ($r = 0$) are clearly emphasised through duplication. **(c)** The normalised version of the MPT depicted in (b). The values of r and ϕ are specified as indicated in Figure 3.8 (b). This normalised version is displayed through the use of a colour map in order to effectively convey the difference in pixel values. This protocol therefore ensures that the undesired phenomenon of emphasising certain pixels at the cost of others is avoided by assigning smaller values (depicted by darker colours) to those (duplicated) pixels that are located closer to the origin ($r = 0$). 30
- 3.8 **(a)** Conceptualisation of an *unnormalised standard* polar transform that is comprised of 360 different angles and $K + 1$ different radii. The interval between consecutive angles is given by $\Delta\phi = 1^\circ$, while the interval between consecutive radii is given by $\Delta r = 1$ pixel. Each entry in the transform (denoted by a black square) is effectively assigned a weight of one. **(b)** Conceptualisation of a *normalised MPT* that is comprised of T different angles and $K + 1$ different radii, with T defined as in Equation (3.4). The interval between consecutive angles $\Delta\phi$ is dependent on the radius and defined in Equation (3.1), while the interval between consecutive radii is given by $\Delta r = 1$ pixel. The collection of entries in the transform that are boxed in red and highlighted in yellow coincide with a single pixel in the original image and is assigned a combined (summed) weight of approximately one. Therefore, when $r = 0$, each individual entry (denoted by a black square) is assigned a weight of approximately $1/T$. However, when $r = r_{\max}$, each individual entry is assigned a value of one. 31
- 3.9 Geometric interpretation of the SDRT when applied to a normalised MPT. 32
- 3.10 Conceptualisation of spirals that coincide with distinct beams associated with the SDRT when applied to an MPT, where **(a)** $\theta = 22.5^\circ$, **(b)** $\theta = 45^\circ$ and **(c)** $\theta = 67.5^\circ$. The spirals conceptualised here are obtained specifically when $r = 0$ is chosen to be at the centre of the original image. Note that the spirals for $\theta = 22.5^\circ$ do not consider (visit) all radii in the interval $r \in [0, r_{\max}]$. The information extracted by these spirals is therefore insufficient. . . . 33
- 3.11 Conceptualisation of the SDRT being applied to a normalised MPT. The relevant targeted curves in the original image $I(x, y)$ are specified for different values of the projection angle θ 34

- 3.12 Conceptualisation of a selected number of beams (associated with the SDRT) that extract information from a normalised MPT. Each beam coincides with a segment of a specific spiral in Cartesian space. Only the green beam extracts information across all radii (that is $r \in [0, r_{\max}]$) and all polar angles (that is $\phi \in [0^\circ, 360^\circ]$). The other beams extract relatively little information. 35
- 3.13 The proposed protocol for constructing composite images. The original image and the eight specified origins for the MPT (indicated by red dots) are depicted in **(a)**. The specified origins and the corresponding composite images are depicted in **(b)** to **(i)**. For each of the aforementioned scenarios, the original image is reflected through the purple axis to obtain the reflected image depicted in purple. Similarly, the original image is reflected through the blue axis to obtain the reflected image depicted in blue, while the original image is reflected through the selected origin (denoted by the red dot) to obtain the reflected image depicted in pink. 36
- 3.14 Examples of how composite images are constructed when the origin of the MPT ($r = 0$) is chosen to be at **(a)** the bottom right corner of the original image, **(b)** the top right corner of the original image, and **(c)** the point midway between the two bottom corners of the original image. In each case, the composite image (boxed in blue) consists of the original image (boxed in green) and three reflections thereof. The relevant origin ($r = 0$) is denoted by a red dot. 37
- 3.15 **(a)** A spiral that coincides with a distinct beam associated with the SDRT that is applied to an MPT for the projection angle $\theta = 22.5^\circ$. The composite image is constructed in the same way as the one depicted in Figure 3.14 (a). The origin therefore coincides with the bottom right corner of the original image. **(b)** The same scenario as the one depicted in (a). In this case, eight spirals that coincide with eight different parallel beams are shown. **(c)** The same scenario as the one depicted in (a). In this case, those segments of the spiral that are located in the reflected images are superimposed onto the original image. This indicates that (in reality) pixel values within the original image are summed along a path that resembles (in shape) a set of shark fins. **(d)** The same scenario as the one depicted in (c), but analysed using the protocol in (b). The eight different parallel beams of the SDRT are now associated with eight different paths, each one resembling a set of shark fins. 38
- 3.16 Geometric interpretation of the MDRT. For a projection angle of $\theta = 22.5^\circ$, the concatenation of a single duplicate of the original normalised MPT is sufficient to ensure that the entire range of radii is covered. 40

- 3.17 **(a)** A spiral that coincides with a distinct beam associated with the MDRT for the projection angle $\theta = 22.5^\circ$. The composite image is constructed in the same way as the ones depicted in Figure 3.14 (a). The origin therefore coincides with the bottom right corner of the original image. **(b)** The same scenario as the one depicted in (a). In this case, eight spirals that coincide with eight different parallel beams are shown. **(c)** The same scenario as the one depicted in (a). In this case, those segments of the spiral that are located in the reflected images are superimposed onto the original image. This indicates that (in reality) pixel values within the original image are summed along a path that resembles (in shape) a set of shark fins. **(d)** The same scenario as the one depicted in (c), but analysed using the protocol in (b). The eight different parallel beams of the MDRT are now associated with eight different paths, each one resembling a set of shark fins. 41
- 3.18 The concentric curves that form part of a shark fin shaped path in Cartesian space. The curves represented in green and orange constitute two distinct sets of concentric paths. For this specific case, the origin of the MPT ($r = 0$) is located at the bottom right corner of the original image and the projection angle associated with the MDRT is given by $\theta = 22.5^\circ$ 42
- 3.19 Overview of the proposed feature extraction protocol that is based on the detection of textural information. 43
- 3.20 Conceptualisation of curved lines that may be detected when the origin (that is where $r = 0$) of the MPT is located on the line that passes through the centre and the bottom right corner of the image. The distance d between the origin of the MPT and the bottom right corner is specified as follows: **(a)** $d = 0$ pixels, **(b)** $d = 70$ pixels, **(c)** $d = 140$ pixels, **(d)** $d = 210$ pixels, **(e)** $d = 280$ pixels, **(f)** $d = 350$ pixels and **(g)** $d \rightarrow \infty$ 44
- 3.21 Conceptualisation of how curved lines with identical curvatures may be defined as convex or concave based on their orientation. The curved lines in **(a)** are concave, while those in **(b)** are convex with respect to the bottom right corner. 45
- 3.22 Conceptualisation of the concentric ellipses associated with different values of d . For a specific value of d , the origins of the respective MPTs are located at 16 evenly distributed positions on said ellipse. These positions are indicated by red dots for the specific ellipse associated with $d = 350$ pixels. 46
- 3.23 Overview of the proposed feature extraction protocol that is based on the detection of curved lines. 46

- 4.1 Conceptualisation of the four distinct groups of feature sets extracted by the systems proposed in this thesis. **(a)** A feature set that contains information extracted via the standard discrete Radon transform (SDRT). **(b)** A group of feature sets that are extracted by applying the modified discrete Radon transform (MDRT) to a normalised multiresolution polar transform (MPT). Each feature set is associated with a specific textural periodicity. **(c)** A group of feature sets that are extracted by applying the MDRT to a normalised MPT. Each feature set is associated with a specific textural orientation. **(d)** A group of feature sets that are extracted by applying the SDRT to a normalised MPT by only considering a projection angle of $\theta = 90^\circ$. Each of the aforementioned feature sets is geared towards the detection of specific curved lines. 49
- 4.2 SDRT-based feature vectors associated with a projection angle of $\theta = 22.5^\circ$. **(a)** An example of a genuine reference sample I_ω . **(b)** An example of a positive questioned sample I_q^+ . **(c)** An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. **(d)** The feature vector associated with (a). **(e)** The feature vector associated with (b) is depicted in green, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 36.68. **(f)** The feature vector associated with (c) is depicted in red, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 63.09. 50
- 4.3 MDRT-based feature vectors associated with a projection angle of $\theta = 33.75^\circ$. These feature vectors are geared towards the detection of textural information based on specific *periodicities*. The origin of the MPT is located at the bottom right corner of the original image. **(a)** An example of a genuine reference sample I_ω . **(b)** An example of a positive questioned sample I_q^+ . **(c)** An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. **(d)** The composite image associated with (a). **(e)** The composite image associated with (b). **(f)** The composite image associated with (c). **(g)** The feature vector associated with (d). **(h)** The feature vector associated with (e) is depicted in green, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 6.58. **(i)** The feature vector associated with (f) is depicted in red, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 15.11. 51

- 4.4 MDRT-based feature vectors associated with a projection angle of $\theta = 11.25^\circ$. These feature vectors are geared towards the detection of textural information based on specific *orientations*. The origin of the MPT is located at the bottom left corner of the original image. **(a)** An example of a genuine reference sample I_ω . **(b)** An example of a positive questioned sample I_q^+ . **(c)** An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. **(d)** The composite image associated with (a). **(e)** The composite image associated with (b). **(f)** The composite image associated with (c). **(g)** The feature vector associated with (d). **(h)** The feature vector associated with (e) is depicted in green, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 16.45. **(i)** The feature vector associated with (f) is depicted in red, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 19.53. 52
- 4.5 Feature vectors extracted by applying the SDRT to an MPT from a projection angle of $\theta = 90^\circ$. These feature vectors are geared towards the detection of curved lines. The origin of said MPT is located on the line that passes through the centre of the original image and its bottom right corner. The distance between the origin of the MPT and said corner of the image is given by $d = 210$ pixels. **(a)** An example of a genuine reference sample I_ω . **(b)** An example of a positive questioned sample I_q^+ . **(c)** An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. **(d)** The feature vector associated with (a). **(e)** The feature vector associated with (b) is depicted in green, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 10.26. **(f)** The feature vector associated with (c) is depicted in red, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 20.29. 53
- 4.6 Conceptualisation of the difference between **(a)** the Euclidean distance and **(b)** a DTW-based distance (Swanepoel, 2015). 54

- 5.1 Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{II} = \{C_2, C_3, C_4, C_5, C_6, C_7, C_8\}$. The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar. 62
- 5.2 Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{III} = \{C_9, C_{10}, C_{11}, C_{12}, C_{13}, C_{14}, C_{15}, C_{16}\}$. The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar. 63
- 5.3 Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{IV} = \{C_1, C_{17}, C_{18}, C_{19}, C_{20}, C_{21}, C_{22}\}$. Recall that classifier C_1 is included within ensemble E_{IV} . The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar. 64
- 6.1 A conceptualisation of how rotational variations may be compensated for within the context of the SDRT-based benchmark system. The questioned feature set is shifted a specified number of columns towards the left or right with respect to the reference feature set. A wrap-around procedure is facilitated by the periodicity of the projection angle in the interval $\theta \in [0^\circ, 360^\circ)$ 73
- 6.2 Examples of images associated with various authentication scenarios in which the systems proposed in this thesis may be useful within the context of future research. 77

List of Tables

1.1	Results for classifier combination through majority voting. The results that constitute a statistically significant improvement in proficiency over the SDRT-based benchmark system (E_I) are depicted in boldface.	11
1.2	Results for classifier combination through score averaging.	12
5.1	Result for ensemble E_I	65
5.2	Results for ensemble E_{II}	65
5.3	Results for ensemble E_{III}	65
5.4	Results for ensemble E_{IV}	65
5.5	The mean AER for systems evaluated on Dolfig's data set. The best results for the systems proposed in this study are presented in boldface.	66
5.6	Statistics for a t -test, where sample set A contains the AERs associated with ensemble E_I and sample set B contains the AERs associated with ensemble E_{II}	68
5.7	Statistics for a t -test, where sample set A contains the AERs associated with ensemble E_I and sample set B contains the AERs associated with ensemble E_{III}	69
6.1	Oracle results for ensembles E_{II} , E_{III} and E_{IV}	73

List of Acronyms

AER	Average Error Rate
DTW	Dynamic Time Warping
EER	Equal Error Rate
FAR	False Acceptance Rate
FRR	False Rejection Rate
HT	Hough Transform
MDRT	Modified Discrete Radon Transform
MPT	Multiresolution Polar Transform
SDRT	Standard Discrete Radon Transform

Nomenclature

Feature extraction

$I(x, y)$	Original image in Cartesian space
θ	Projection angle
N_θ	Number of projection angles
R_θ	Projection profile associated with projection angle θ
x'	Axis perpendicular to projection beams in Cartesian space
ϕ	Polar angle
r	Polar radius
$\Delta\phi$	Interval between consecutive polar angles
Δr	Interval between consecutive polar radii
Δs	Arc length between consecutive samples
r_{\max}	Maximum polar radius for a given image
N_ϕ	Number of polar angles
K	Number of rows in MPT
T	Number of columns in MPT

$M(r, \phi)$	Unnormalised MPT
$M_{\text{norm}}(r, \phi)$	Normalised MPT
$\eta(r)$	Weight employed for normalising r th row of MPT
N_{MPT}	Number of specified origins for MPT
ϕ'	Axis perpendicular to projection beams in polar space
d	Distance between origin of MPT and corner of $I(x, y)$

Template matching and thresholding

X_i	Feature set associated with i th base classifier
V_i	i th feature vector of feature set X_i
V_{dim}	Feature vector dimension
ω	Claimed class
I_ω	Genuine reference sample of class ω
I_q	Questioned sample
I_q^+	Positive questioned sample
I_q^-	Negative questioned sample
N	Number of genuine reference samples
\mathbf{I}_ω	Set of N reference samples for class ω
\mathbf{X}_ω	Group of N feature sets associated with N reference samples of class ω
X_q	Feature set associated with a questioned sample
$D(X_q, \mathbf{X}_\omega)$	Average raw DTW-based distance
$D_{\text{norm}}(X_q, \mathbf{X}_\omega)$	Normalised DTW-based distance
$S_{\text{norm}}(X_q, \mathbf{X}_\omega)$	Normalised score

Ensemble construction and classifier combination

μ_{AER}	Mean average error rate
τ	Global threshold
C_i	i th base classifier
E_i	i th classifier ensemble
$C^{(\text{s})}$	Combined classifier through score averaging
$C^{(\text{m})}$	Combined classifier through majority voting

Chapter 1

Introduction

1.1 Background and motivation

Feature extraction through textural pattern and curved line detection has been implemented successfully and independently on various occasions. The proficiency of generic object recognition systems may therefore be improved through enhanced feature detection based on the aforementioned approach. Research areas that may particularly benefit from these enhancements include (but are not limited to) vehicle-related object detection, health-related object detection and biometric authentication.

Vehicle-related object detection. A recent increase in security concerns involving vehicles has led to an increased demand for effective vehicle-related object detection strategies. Vehicle-related object detection includes scenarios in which vehicles are detected, as well as scenarios in which the vehicle is used as a tool for detection. Whenever a vehicle that exhibits suspicious behaviour is detected, the appropriate authorities may be notified. Driver safety may also be enhanced through the detection of road signs and road boundaries. The survey by Sivaraman and Trivedi (2013) highlights a number of research projects that focus on the detection of vehicles and analyse driver behaviour. Zhang and Zheng (2010) employ texture-based features for the purpose of vehicle detection. Yu and Jain (1997) on the other hand utilise Hough transforms for the purpose of detecting straight lines that coincide with lane boundaries. It should therefore be clear that the improved detection of textural patterns and curved lines within an image will probably enhance the proficiency of vehicle-related object detection algorithms.

Health-related object detection. Medical image analysis based on the detection of textural patterns and curved lines may provide a means of detecting medical anomalies at an early stage, which could result in lives being saved. In particular, line and texture detection strategies may aid in establishing whether components within a medical image should be deemed normal or abnormal. In addition to this, the aforementioned strategies may be utilised

to render certain medical practices more efficient or aid medical staff in their work. Examples of medical systems that rely on line detection include those proposed by Nguyen *et al.* (2013) and Wu *et al.* (2013). Algorithms that involve texture detection within the context of health sciences are proposed by Castellano *et al.* (2004) and Nanni *et al.* (2010). Improved texture and line detection within medical images may therefore benefit the health industry.

Biometric authentication. Numerous existing biometric authentication systems rely on either curved line or texture detection. Wu and Ye (2009) employ the standard discrete Radon transform (SDRT) for the purpose of detecting straight lines in finger vein patterns within the context of biometric authentication. Coetzer *et al.* (2004), Jayadevan *et al.* (2009) and Panton and Coetzer (2010) utilise the SDRT for the purpose of detecting straight lines in static handwritten signature images, thereby facilitating successful offline signature verification. Wei *et al.* (2008) and Nikam and Agarwal (2008) successfully employ textural information for the purpose of iris and fingerprint detection, respectively. The value of curved line and texture detection is therefore evident within the context of biometric authentication.

The increased demand for improved feature extraction algorithms through the detection of curved lines and textural patterns is clear. The aim of this thesis is therefore to develop novel generic systems for this purpose.

1.2 Key concepts

1.2.1 Texture detection

Objects within static images may be described by either their external shape or their internal properties. In many scenarios the texture of an object constitutes a characteristic internal property that is well suited for distinguishing it from other objects. The texture of an object may be described by quantifying its coarseness, periodicity, orientation, uniformity, etc. Textural information may furthermore be extracted from images in either a local or a global manner. Local textural information pertains to fine detail within an image, typically associated with a small subimage (retina). Global textural information, on the other hand, is typically extracted from an image in its entirety, as is proposed in this thesis. The SDRT has been used on previous occasions for the purpose of texture detection based on textural orientation (Hai-peng *et al.*, 2005), as well as for texture analysis based on textural periodicity (Jafari-Khouzani and Soltanian-Zadeh, 2005). In this thesis the SDRT is employed for the purpose of extracting textural information based on both textural orientation and textural periodicity.

1.2.2 Curved line detection

The feature extraction protocol proposed in this thesis is based on the SDRT and modifications thereof. The SDRT of an input image is obtained by calculating projection profiles of said image from a number of equally spaced projection angles in the interval $\theta \in [0^\circ, 180^\circ)$. A prominent peak within a specific projection profile is indicative of a prominent straight line within the input image. The SDRT is therefore geared towards the detection of straight lines within static images. The SDRT is discussed in more detail in Section 3.2.

Although many systems have been proposed for the purpose of curved line detection, the Hough transform (HT) and adaptations thereof remain the most well-known and widely used. The linear HT is similar to the SDRT, which is employed in this thesis, in the sense that both transforms are geared towards the detection of straight lines within an image.

The linear HT differs from the SDRT in the sense that it targets straight lines in *parameterised* form, $y = ax + b$. A predefined number of accumulator cells in the ab (parameter) space is initially set to zero. For each allotted value of the gradient a , each set of pixel coordinates (x_i, y_i) in Cartesian space is iteratively considered, in which case the y -intercept b is estimated. During each iteration, the value of an accumulator cell (a, b) is incremented. The eventual value of a certain accumulator cell (a, b) is therefore indicative of the number of one-valued pixels in the image that approximately lie on the line $y = ax + b$, assuming that said image is binary.

It is assumed that zero-valued pixels constitute the background, while one-valued pixels form part of objects within a binary image. It is therefore important to note that the HT is only defined for binary images, while the SDRT may be applied to either binary or grey-scale images. The linear HT is also not invertible as is the case for the SDRT.

Since $a \in (-\infty, \infty)$ and $b \in (-\infty, \infty)$, the normal representation of a straight line $\rho = x \cos \gamma + y \sin \gamma$ is often used within the context of the linear HT in order to ensure a finite parameter space. Said parameter space is therefore defined for $\gamma \in [-90^\circ, 90^\circ]$ and $\rho \in [-H, H]$, where H represents the maximum distance between opposite corners of the image. The reader is referred to Gonzales and Woods (2010) for a more detailed description of the aforementioned representation of straight lines.

Although the HT may also detect more complex parameterised curves (typically conic sections) by considering higher dimensional parameter spaces, such strategies are both inefficient and limited to the parameterised curves in question.

Due to the above-mentioned limitations, a novel strategy for detecting curved lines, which involves appropriate modifications to the SDRT, is developed in this thesis (Section 3.5).

1.2.3 Biometric authentication

The utilisation of biometric data is becoming increasingly popular within the context of identity authentication. This is due to the fact that it is deemed to be more secure than knowledge-based or possession-based data which is typically easier to gain access to and may be easily replicated.

Biometric data is categorised into physiological and behavioural data. The former includes fingerprints, retinal patterns, irises, hand-veins, etc. On the other hand, identity authentication based on behavioural biometrics includes voice recognition, handwritten signature verification, gait recognition, keystroke recognition, etc. Behavioural biometrics are employed, in the form of offline handwritten signatures, to gauge the proficiency of the systems proposed in this thesis. The reader is referred to Wang and Geng (2009) for an in-depth discussion on behavioural biometrics.

Handwritten signatures may be categorised into offline (static) signatures and online signatures. Offline signatures are typically extracted from legal documents, bank cheques or credit card slips, while online signatures are typically acquired dynamically through a digitising device like a tablet. The systems proposed in this thesis are implemented on an ideal (noiseless) *offline* signature database (see Section 5.2) for the purpose of detecting skilled forgeries. A skilled forgery is produced when the forger has access to one or more authentic samples of the victim's signature and/or witnessed the signing event.

1.2.4 Performance metrics

Performance metrics typically employed for verification systems are utilised to gauge the proficiency of the systems proposed in this thesis. Although similar performance metrics may be employed within the context of recognition systems, it is important to note that recognition and verification systems are conceptually different. A recognition system aims to determine to which class (among a number of classes) a questioned sample belongs to. A verification system, on the other hand, simply aims to determine whether or not a questioned sample belongs to a claimed class.

Within the context of offline signature verification, the systems developed in this thesis match a questioned sample (signature) to a model of the claimed writer. The resulting distance is subsequently normalised (on a writer-specific basis) and converted into a normalised score. A global threshold τ is applied to all scores (regardless of the claimed writer) so that a false acceptance rate (FAR) and a false rejection rate (FRR) are obtained. These performance metrics are estimated as follows,

$$\text{FAR} = \frac{\text{Number of forgeries accepted}}{\text{Total number of forgeries}}, \quad (1.1)$$

$$\text{FRR} = \frac{\text{Number of authentic samples rejected}}{\text{Total number of authentic samples}}. \quad (1.2)$$

During experimentation a number of independent trials are conducted. For each trial, the data set is partitioned into an optimisation set and an evaluation set as explained in Section 5.4. These sets are comprised of *different* writers. The optimisation writers are first used to obtain the global threshold τ that results in an equal error rate (EER) as explained in Section 4.4. The EER is achieved when the FAR equals the FRR (see Figure 1.1). The aforementioned threshold is subsequently applied to questioned samples in the evaluation set, after which an FAR and FRR are reported for the specific evaluation set. Said FAR and FRR are not expected to differ substantially for a specific classifier. The average of these two rates is referred to as the average error rate (AER) which quantifies the proficiency of said classifier based on the trial in question. A sufficient number of trials are conducted in order to ensure statistically sig-

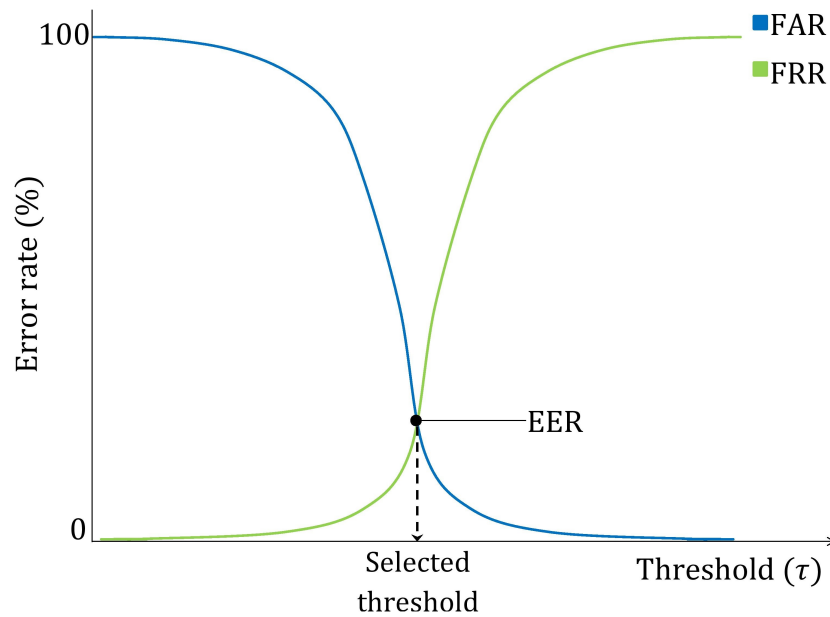


Figure 1.1: Conceptualisation of the EER, which determines the selection of an appropriate global threshold based on the *optimisation* set for a specific trial.

nificant results, as explained in Section 5.4. System performance is quantified by the mean AER across all trials. Said performance metric is denoted by μ_{AER} .

1.2.5 Local and global features

Features extracted from a questioned image may be categorised into local and global features. A local feature quantifies a localised attribute of an object like an individual pen stroke within a handwritten signature or an individual segment within an object boundary. On the other hand, global features describe an image, or an object within an image, in its entirety. Various existing

systems, including the system proposed by Panton and Coetzer (2010) employ a combination of both local and global features. The systems proposed in this thesis however focus exclusively on the extraction of *global* features.

1.2.6 Classifier ensembles and combination

A classifier ensemble constitutes an assembly of base classifiers (often associated with different feature sets) that consider similar information. In order to improve upon the proficiency of the individual base classifiers within an ensemble, said base classifiers may be fused into a combined classifier. The aforementioned fusion may be achieved in three ways. The first approach involves score-level fusion, where the scores associated with the individual base classifiers are fused into a single score. The systems proposed in this thesis consider amongst other things score-level fusion through *score averaging* (see Section 4.6). The second approach to combining the individual base classifiers within an ensemble involves decision-level fusion. The systems proposed in this thesis consider amongst other things decision-level fusion through *majority voting* (see Section 4.6). Feature level fusion presents a third approach, but is not considered in this thesis. The aforementioned fusion strategy typically combines various features before a score is calculated.

1.3 Scope and objectives

The standard discrete Radon transform (SDRT) has been successfully implemented for the purpose of object recognition on several previous occasions (Coetzer *et al.*, 2004; Wang *et al.*, 2007; Hjouj and Kammler, 2008; Jayadevan *et al.*, 2009; Panton and Coetzer, 2010; Swanepoel, 2015). It is however limited in the sense that it is only capable of detecting *straight* lines within images, while a typical object within an image also contains curved lines.

The scope of this work is limited to an investigation into *novel* alternative feature extraction strategies for the purpose of object recognition. Said strategies involve modifications to the SDRT for the purpose of detecting curved lines and textural information within an image.

Feature matching is achieved through a dynamic time warping (DTW) algorithm. It is assumed that a limited number of reference samples are available for training purposes. An investigation into the training of appropriate statistical models, like probabilistic graphical models, as an alternative to template matching, is therefore not conducted in this thesis since it is deemed to be outside the scope of this work.

The novel feature extraction techniques conceptualised in this thesis are *generic* in the sense that it is reasonable to presume that they will be of value within any authentication scenario where objects (or their boundaries) are characterised by distinct prominent curved lines. The proficiency of the

systems proposed in this thesis is however evaluated by gauging their ability to detect skilled forgeries within the context of offline signature verification. The data set considered contains ideal (noiseless) handwritten signatures with a fixed pen-stroke width. An investigation into appropriate pre-processing techniques that may involve noise removal and pen-stroke width normalisation is therefore not conducted in this thesis since it is deemed to be outside the scope of this work. In this thesis we therefore aim to:

- (1) Adapt the SDRT in such a way that not only straight lines, but also curved lines with a wide range of curvatures may be detected;
- (2) Adapt the SDRT in such a way that the intrinsic textural information (based either on periodicity or orientation) associated with an object within an image may be detected;
- (3) Demonstrate that the combined classifiers that utilise the features extracted in (1) and (2) outperform an existing benchmark classifier that only utilises SDRT-based features;
- (4) Demonstrate that a specific base classifier obtained in (1) and (2) may be more adept at detecting objects belonging to a *specific* class (writer), which opens up the possibility of successful dynamic classifier selection and/or dynamic weighted score fusion in future research.

1.4 System design

In this section a concise overview of the systems proposed in this thesis is provided. Said systems are conceptualised in Figure 1.2.

1.4.1 Data

The experimental data constitutes handwritten signatures that were originally captured online by Hans Dolfing for his PhD thesis (Dolfing, 1998). These dynamic signatures were subsequently converted by Coetzer *et al.* (2004) into ideal (noiseless) static signature images with a fixed pen-stroke width. The aforementioned data set contains authentic signature samples and *skilled* forgeries from 51 different writers and is discussed in detail in Section 5.2. Five authentic reference samples constitute a *template* for a specific class (writer).

1.4.2 Feature extraction

A feature set that is based on the SDRT is first extracted from each input image according to the protocol proposed by Coetzer *et al.* (2004). In this way a *benchmark system* that is geared towards the detection of straight lines is

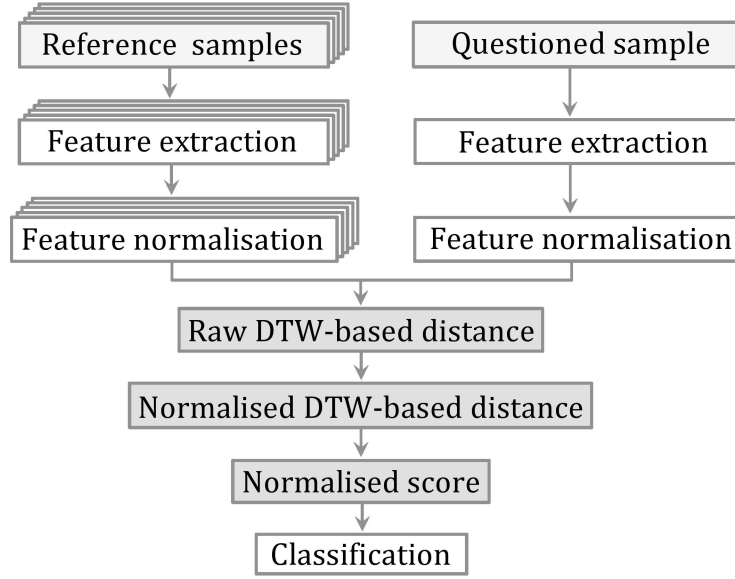


Figure 1.2: A conceptualisation of the systems proposed in this thesis.

obtained. The *first* “group” of feature sets therefore constitutes a *single* feature set based on the SDRT.

In order to facilitate the feature extraction protocols proposed in this thesis, appropriately normalised multiresolution polar transforms (MPT) are first obtained (Section 3.3). The origins of the MPTs are specified according to the type of features that are targeted for extraction.

Texture detection. In order to extract textural information, an MPT is applied to a composite image that contains the original input image, as well as appropriately appended reflections thereof. When the SDRT is applied to said MPT, each projection beam (associated with the SDRT) coincides with a spiral in Cartesian space. A spiral within the aforementioned composite image coincides with a shark-fin shaped pattern in the original (input) image which is well suited for extracting textural information.

To ensure that spirals are not limited to a single revolution, a modified discrete Radon transform (MDRT) is proposed. Said MDRT, constitutes the application of the SDRT to a periodical continuation of the normalised MPT (see Section 3.4.3). Various textural orientations may be considered by specifying different origins for the MPT. This concept is explained in Section 3.4.4. The *second* group of feature sets therefore contain textural information based on specific textural *periodicities*. Furthermore, the *third* group of feature sets contain textural information based on specific textural *orientations*.

Curved line detection. The SDRT is finally applied to a normalised MPT from a projection angle of 90° . This facilitates the detection of concentric circles in Cartesian space. The range of curvatures associated with the aforementioned concentric circles may be adjusted. This is achieved through specifying that the origin of the MPT is located at a certain distance from the

edge of the input image as explained in Section 3.5. The *fourth* and final group of extracted feature sets therefore contain information pertaining to concentric curved lines with a specific range of curvatures.

1.4.3 Feature normalisation

Since the extracted feature vectors may differ substantially simply due to variations in scale or translation, appropriate feature vector normalisation is required. Therefore, once the feature vectors are extracted they are subsequently rendered scale and translation invariant. The feature vector normalisation protocol and the reason for not enforcing rotation invariance are discussed in Section 3.6.

1.4.4 Feature matching

In order to quantify the difference between the corresponding normalised feature vectors associated with a specific feature set and extracted from two different samples, a DTW-based matching algorithm is employed. For a specific feature set, the *average* of the DTW-based distances between the individual feature vectors associated with a questioned sample and the *corresponding* feature vectors associated with a number of authentic reference samples is first calculated (see Section 4.2). The above-mentioned raw distance is subsequently *normalised* on a class-specific (writer-specific) basis through a strategy that is similar to the well-known *z*-score as explained in Section 4.3. Said normalised distance is subsequently converted into a normalised score. Feature-specific score normalisation is not conducted since the feature vectors associated with a specific *group* of feature sets are already appropriately normalised through the strategy touched upon in Section 1.4.3.

1.4.5 Classification

A number of independent trials are conducted. For each trial a *global* threshold across all classes (writers) is determined by considering an optimisation set, after which said threshold is applied to an evaluation set. The aforementioned protocol was briefly outlined in Section 1.2.4. A questioned sample is accepted as authentic by a specific classifier when its associated score is larger than or equal to the global threshold. It is otherwise rejected as fraudulent. The experimental protocol is detailed in Section 5.4.

1.4.6 Ensemble construction and classifier combination

Four classifiers ensembles are constructed by assembling base classifiers (associated with specific feature sets) that consider similar information. Each of

these ensembles correspond to a specific *group* of feature sets as outlined in Section 1.4.2.

The first ensemble E_I contains the benchmark SDRT-based classifier. The second and third ensembles, E_{II} and E_{III} , contain the classifiers that detect textural information based on textural periodicity and orientation, respectively. The fourth and final ensemble E_{IV} contains the classifiers that extract information pertaining to concentric curved lines. The construction of the aforementioned classifier ensembles is described in detail in Section 4.5.

The base classifiers within a specific ensemble may be combined through score averaging or majority voting as discussed in Section 4.6. The classifier ensembles constructed in this thesis are conceptualised in Figure 1.3.

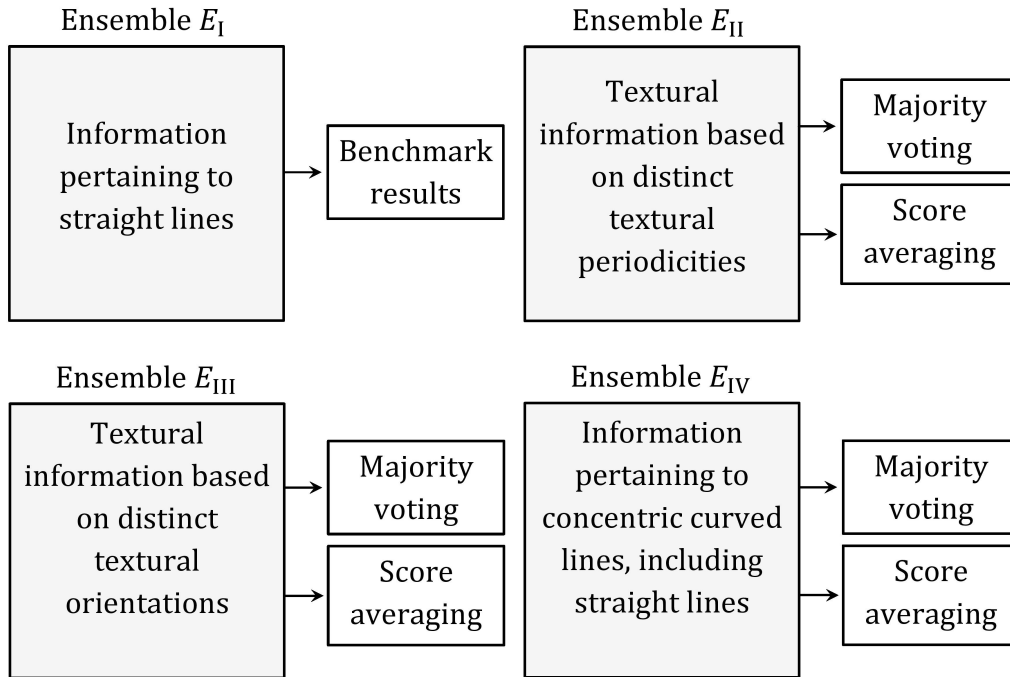


Figure 1.3: The four classifier ensembles constructed in this thesis.

1.5 Abbreviated results

The proficiency of the generic systems developed in this thesis is estimated by considering a data set that contains authentic handwritten signatures and skilled forgeries associated with 51 writers.

Recall that ensemble E_I constitutes the single SDRT-based benchmark classifier. Ensemble E_{II} contains a number of base classifiers that extract textural information based on different textural periodicities, while ensemble E_{III} is

comprised of base classifiers that extract textural information based on different textural orientations. Finally, ensemble E_{IV} contains base classifiers that extract information pertaining to curved lines with different curvatures.

Majority voting. When the respective base classifiers within ensembles E_{II} , E_{III} and E_{IV} are combined through majority voting, each of the combined classifiers are more proficient than the SDRT-based benchmark system (E_I). The relevant results are presented in Table 1.1. Note that the combined classifiers associated with ensembles E_{II} , E_{III} and E_{IV} show relative improvements in proficiency over the SDRT-based benchmark system (E_I) of 35.86%, 28.29% and 11.67% respectively. It is demonstrated in Section 5.5.4 that said improvement in proficiency is statistically significant within the context of ensembles E_{II} and E_{III} .

Table 1.1: Results for classifier combination through majority voting. The results that constitute a statistically significant improvement in proficiency over the SDRT-based benchmark system (E_I) are depicted in boldface.

Ensemble	$\mu_{AER}(\%)$
E_I	9.51
E_{II}	6.10
E_{III}	6.82
E_{IV}	8.40

Score averaging. When the respective base classifiers within ensembles E_{II} , E_{III} and E_{IV} are combined through score averaging, all of the combined classifiers, except one, are more proficient than the SDRT-based benchmark system (E_I). The relevant results are presented in Table 1.2. Note that the combined classifiers associated with ensembles E_{II} and E_{III} show relative improvements in proficiency over the benchmark SDRT-based system (E_I) of 13.56% and 11.78% respectively.

Note that, within the context of ensembles E_{II} , E_{III} and E_{IV} , the combined classifiers that employ majority voting are consistently more proficient than their counterparts that employ score averaging.

The proficiency of the systems developed in this thesis is placed into perspective by comparing it to a number of existing systems in Section 5.5.3.

1.6 Contributions

Recall that the main objective of this study is to develop novel feature extraction strategies that are based on, and improve upon, existing feature extraction techniques. Within this context, the contributions of this thesis are as follows:

Table 1.2: Results for classifier combination through score averaging.

Ensemble	$\mu_{\text{AER}}(\%)$
E_{I}	9.51
E_{II}	8.22
E_{III}	8.39
E_{IV}	10.14

The facilitation of novel and improved feature extraction through the construction of appropriately normalised and appropriately centred multiresolution polar transforms (MPTs). Traditionally the standard discrete Radon transform (SDRT) has been employed for the purpose of feature extraction within the context of detecting prominent straight lines in an image. In these scenarios the SDRT is applied *directly* to the image in question. In this thesis it is proposed that an appropriately normalised and appropriately centred MPT of the original image is first constructed. The SDRT or a modified discrete Radon transform (MDRT) is subsequently applied to the MPT. This facilitates the detection of various curved lines, as well as the extraction of textural information associated with spirals in the original image. In this thesis it is clearly demonstrated that the above-mentioned protocol facilitates the extraction of richer features than those associated with the traditional implementation of the SDRT. It is also demonstrated that classifiers based on these novel features outperform a traditional SDRT-based classifier. This superiority is often statistically significant. To the best of our knowledge, the proposed protocol for appropriately normalising and centring an MPT for the purpose of extracting features based on the detection of curved lines and textural information is novel.

A novel protocol that facilitates the detection of textural information. In this thesis a novel strategy for detecting textural information within an input image based on either textural periodicity or textural orientation is proposed. This is made possible by applying the MDRT to appropriately normalised and appropriately centred MPTs. Although the proficiency of the proposed texture detection protocol is clearly demonstrated by considering a data set containing binary (signature) images, the proposed protocol should be equally applicable to grey-scale images.

A novel protocol that facilitates the detection of curved lines. In this thesis a novel strategy for detecting curved lines within an input image is proposed. This is made possible by applying the SDRT from an angle of 90° to appropriately normalised and appropriately centred MPTs. This provides

a valuable alternative to the computationally complex Hough transform (HT) that detects curved lines (mostly conic sections) within a high dimensional parameter space. The HT is also restricted in the sense that it is only able to detect curved lines within binary images, while the proposed protocol is not.

Proficient off-line signature verification. It is clearly demonstrated that the proficiency of the systems developed in thesis compares favourably to that of existing state-of-the-art systems within the context of writer-dependent off-line signature verification. The proposed feature extraction protocol is however generic, since said protocol is expected to be proficient within the context of other authentication scenarios as well, as long as the objects in question (or their boundaries) are characterised by distinct prominent curved lines.

Avenues for future research. It is demonstrated that an investigation into more sophisticated classifier fusion strategies as part of future research is likely to result in a further improvement in system proficiency. The proposed feature extraction protocol may also form part of a system that trains appropriate generative or discriminative statistical models for classification purposes.

1.7 Thesis outline

The remainder of this thesis is structured as follows:

Chapter 2: Literature study. The systems proposed in this thesis are placed into perspective by comparing the aforementioned systems to existing systems that apply similar research principles. Existing research on the utilisation of the SDRT for feature extraction is discussed. Existing research on texture-based object recognition and generic curve detection is also investigated.

Chapter 3: Feature extraction. Three feature extraction protocols, associated with four distinct groups of feature sets, are discussed. The first protocol pertains to the benchmark SDRT-based system. The second protocol involves the application of an MDRT to a normalised MPT. The third protocol involves the application of the SDRT from a specific projection angle of 90° to a normalised MPT. An appropriate feature vector normalisation strategy is finally presented.

Chapter 4: Feature matching, threshold selection and classifier ensembles. A feature matching protocol that is based on a DTW-based distance measure is described. Score normalisation and threshold selection strategies are also presented. The construction of classifier ensembles and the combination of the base classifiers within an ensemble are finally explained.

Chapter 5: Experiments. The offline signature data that is considered for experimental purposes is first discussed, after which the system parameters are specified. The experimental protocol is subsequently explained and the experimental results are finally presented. The proficiency of the proposed systems is placed into perspective through a comparison with relevant existing systems. The protocol for statistical significance testing is also presented.

Chapter 6: Conclusion and future work. The traits associated with the systems proposed in this thesis are revisited and assessed in hindsight. As a continuation of this study, possible future research that may improve upon the proficiency of the systems proposed in this study is presented and analysed.

Chapter 2

Literature study

2.1 Introduction

It is infeasible to directly compare the specifications and proficiency of the novel systems proposed in this thesis (in their entirety) to those of existing systems. This is due to the fact that the proposed systems employ a unique combination of feature extraction and experimental evaluation protocols. However, the relevance of this research may be placed into perspective by investigating previous work that was motivated by similar objectives and/or employed similar protocols. We therefore opt to investigate existing research as follows: In Section 2.2 we discuss existing research that employs the standard discrete Radon transform (SDRT) for the purpose of feature extraction. In Section 2.3 we investigate existing research pertaining to texture-based object recognition. Finally, in Section 2.4 we investigate existing research pertaining to the detection of curved lines in generic images.

2.2 SDRT for feature extraction

The SDRT constitutes an integration transform that may be used for the purpose of feature extraction for various applications. Recall that the SDRT calculates projection profiles of the original image from a predefined set of equally distributed angles. When the projection profiles are packed into the columns of a matrix, said matrix is often referred to as a sinogram. High-intensity values within a sinogram are indicative of the presence of straight lines in the original image. In addition to feature extraction, the SDRT may also be used to simulate X-rays within the context of CAT scans and is invertible through strategies like filtered back projection. In this thesis the SDRT is applied to a multiresolution polar transform (MPT) both directly, for the purpose of extracting curved lines from generic images, and also as part of a more complex protocol for the purpose of texture-based recognition. We therefore opt to investigate existing research that is based on the utilisation of

the SDRT for the purpose of feature extraction within the context of generic applications (see Section 2.2.1) and for the specific purpose of offline signature verification (see Section 2.2.2).

2.2.1 Generic applications

Multimedia fingerprinting. Seo *et al.* (2004) exploit the scale, translation and rotational properties of the SDRT for the purpose of feature extraction within the context of multimedia fingerprinting, specifically for images. The authors note that cryptographic hash functions are sensitive to the manipulation of images. In order to guarantee a more robust system, they propose that multimedia fingerprinting is employed instead. The authors investigate various manipulations of 1000 images and report false rejection rates between 0.00% and 4.90%.

Facial recognition. Vankayalapati and Kyamakya (2009) employ the SDRT and the wavelet transform for the purpose of facial recognition. These two feature extraction techniques are first employed in isolation and then combined. Finally, each of these feature extraction techniques is combined with a technique based on linear discriminant analysis (LDA). The results for the aforementioned systems are then compared to the respective recognition rates obtained for systems based exclusively on LDA and principal component analysis (PCA). The authors report that the combination of the SDRT and wavelet transform outperforms both LDA and PCA. The reported recognition rates range between 97% and 74% (for 40 to 400 front facing samples) in the case of LDA and between 88% and 49% (for 40 to 400 front facing samples) in the case of PCA. When implemented on the same data, the combination of the SDRT and wavelet transform achieves recognition rates that range between 100% and 91%.

Iris recognition. Kumar *et al.* (2012) use the SDRT to extract local features for automated iris recognition. The authors compare the proficiency of the SDRT for feature extraction to that of a one-dimensional log-Gabor filter and a monogenic log-Gabor filter. It is reported that the SDRT outperforms the aforementioned techniques. Their best results for the SDRT-based approach are reported for the CASIA V4 database (that is a recognition rate of 90.43%) and their worst results are reported for the FRGC database (that is a recognition rate of 33.20%).

2.2.2 Offline signature verification

Signature verification has been researched extensively due to its significance within the context of biometric identification. In order to gain a broad perspective on the various protocols and data sets employed by recent research

in offline signature verification, the reader is referred to the following surveys: Hou *et al.* (2004), Impedovo and Pirlo (2008), Al-Omari *et al.* (2011), Yadav *et al.* (2013), Khan and Dhole (2014), and Hafemann *et al.* (2015).

Since static signature images are employed for gauging the proficiency of the generic systems proposed in this thesis, we elaborate on recently proposed protocols for detecting *skilled* forgeries within the context of offline signature verification. These protocols include the utilisation of dynamic time warping (DTW) algorithms, the construction of hidden Markov models (HMMs) and the use of support vector machines (SVMs). A skilled forgery is produced when the forger has access to one or more authentic samples of the targeted writer's signature and sufficient time to practice imitating it.

It is important to note that no single standard international database currently exists and that various different databases have been employed for gauging the proficiency of existing systems.

The signatures within the data set considered for our study were originally captured online by Hans Dolfing for his PhD thesis (Dolfing, 1998). The pen-tip coordinates of these dynamic signatures were subsequently used to construct static signature images that contain no noise and have a fixed stroke width of five pixels. Said data set is therefore *ideal* and is subsequently referred to as Dolfing's data set. Since this data set is not frequently employed for the purpose of offline signature verification, relatively few studies exist to which our results may be directly compared. We do however opt to detail the results of existing research that employs similar modelling techniques and similar experimental protocols to those employed in this study. This may provide a general indication of what may be expected in terms of system proficiency.

Global features modelled by DTW-based templates and ring-structured HMMs. Offline signature verification may be achieved, as proposed in Coetzer *et al.* (2004) and Coetzer (2005), through first calculating the SDRT of each original signature image. A DTW-based distance measure may then be used to quantify the dissimilarity between corresponding projection profiles. Alternatively, the projection profiles associated with authentic signatures belonging to a specific writer may be used to train an appropriate HMM. The aforementioned HMM has a ring topology and one state skip and is constructed in such a way that the probability of entering the HMM at a specific state is equal to the probability of entering the HMM at any other state. The rationale behind this topology is to ensure rotational invariance. The so-called Stellenbosch data set and Dolfing's data set are employed for experimentation purposes. Dolfing's data set, that is the same data set employed in our study, consists of 60 skilled forgeries, 15 authentic test samples and 15 authentic training samples per writer, from a total of 51 different writers. Coetzer *et al.* (2004) report a best equal error rate (EER) of 17.70% for the Stellenbosch data set and an EER of 12.20% for Dolfing's data set when skilled forgeries are targeted by the HMM-based systems. In this thesis, the

DTW-based system proposed by Coetzer *et al.* (2004) (for which the proficiency was not reported when evaluated on Dolfing’s data set) is effectively reproduced and subsequently significantly improved upon by also extracting features that are able to detect generic curves and textural information. The protocol for extracting the aforementioned features is detailed in Sections 3.4 and 3.5. In Section 5.5 it is shown that the systems proposed in this thesis are significantly more proficient than the HMM-based system proposed by Coetzer *et al.* (2004) when all of the aforementioned systems are evaluated on Dolfing’s data set. This is achieved despite the fact that the systems proposed in this thesis utilise fewer training samples (5 instead of 15) and fewer projection angles (8 instead of 128) than the HMM-based system proposed by Coetzer *et al.* (2004).

An alternative DTW-based algorithm for matching horizontal and vertical projection profiles. Jayadevan *et al.* (2009) employ the SDRT in order to specifically obtain the vertical and horizontal projection profiles of an original signature image. The SDRT is therefore only calculated for projection angles of 0° and 90° . They also proceed to combine said projection profiles into a single feature vector. A set of three feature vectors is therefore obtained for a specific signature image. A DTW algorithm is subsequently employed for feature matching. However, instead of using DTW as proposed in our research, Jayadevan *et al.* (2009) employ a DTW algorithm to determine a so-called difference cost, but also to calculate the size of the closed area(s) formed between the diagonal of the DTW-grid and the so-called matching path (as determined by the DTW algorithm). The product of the aforementioned difference cost and size of the closed area(s) subsequently determines the dissimilarity score. The system is evaluated on the GPDS signature database and the authors report a total error rate of 44.02% (that is an average error rate of 22.01%) when five genuine signature samples per writer are employed for training purposes. This happens to be the same number of samples that are employed for training purposes by the systems proposed in this thesis. However, since the GPDS database contains non-ideal (noisy) offline signatures, while Dolfing’s database contains ideal (noiseless) signature images, it is not unexpected that our systems achieve superior results.

Local features modelled by ensembles of HMM-based classifiers. Panton and Coetzer (2010) propose a feature extraction protocol that is also based on the SDRT, but instead of exclusively focusing on global features, they propose that a number of local *and* global features are appropriately combined. This is in contrast to the systems proposed in this thesis that focus exclusively on the extraction of *global* features. The local features are extracted by calculating the SDRT of local retinas within the signature in question. Each different retina is modelled by a different HMM. When their system is implemented on Dolfing’s data set, the authors report an EER of

8.89%. This constitutes a significant improvement in the proficiency reported for the benchmark system proposed by Coetzer *et al.* (2004) and also indicates that the utilisation of classifier ensembles within the context of the SDRT warrants further investigation. Such an investigation is conducted in this thesis.

A writer-independent approach. Swanepoel (2015) proposes a writer-independent approach to online and offline signature verification. Within the context of offline signature verification, the SDRT is employed for the purpose of extracting feature vectors from a questioned signature. Said feature vectors are subsequently matched (through DTW) to the corresponding feature vectors extracted from a reference signature that is known to belong to the claimed writer. In this way, a dissimilarity vector is obtained of which the dimension is equal to the number of projection angles associated with the SDRT. Examples of positive and negative dissimilarity vectors may be obtained from a set of so-called guinea pig writers, which may subsequently be used to train two different models in dissimilarity space, that is a model that represents *all* positive signatures and a model that represents *all* negative signatures. These models are therefore *writer-independent*. The authors employ quadratic discriminant analysis (QDA) and SVMs to model the aforementioned two classes. The aforementioned systems are evaluated on Dolfing’s data set and EERs of 6.06% and 5.55% are reported for the SVM-based and QDA-based systems respectively when the parameters are chosen in a similar way as proposed in this thesis. As is the case for the systems proposed in this thesis, five training signatures are utilised per writer. A conventional *writer-dependent* approach to signature modelling is however proposed in this thesis. The contribution of this thesis is furthermore focussed on the development of alternative feature extraction strategies (that are geared towards curved line detection and the extraction of textural information) and *not* signature modelling per se. Therefore, a direct comparison between the results reported in this thesis and those reported by Swanepoel (2015) does provide valuable perspective, but it is not indicative of the contribution of this thesis.

2.3 Texture-based recognition

In this thesis a novel approach to texture detection is proposed, which involves the application of a modified SDRT to a normalised MPT. This approach enables one to detect straight lines in polar space that coincide with spirals in Cartesian space (Fick *et al.*, 2016). Said spirals are well suited for extracting periodic (textural) information from an image as explained in detail in Section 3.4.

In this section existing research on the extraction of textural information from images is discussed. These conventional approaches to texture detection employ amongst other things, Gabor filters, log-polar transforms, the SDRT

and the wavelet transform.

A log-polar transform and a discrete wavelet packet transform. Pun and Lee (2003) propose that a log-polar transform of a questioned image which contains textural information is obtained, after which a discrete wavelet packet transform is applied to the result. They calculate the log-polar transform in two steps. The first step involves the calculation of the polar transform of the original image, while a logarithmic function is applied to said polar transform during the second step. Although the log-polar transform is related to the normalised MPT as proposed in this thesis, the discrete wavelet packet transform is conceptually different from the SDRT and its modifications.

A combination of the SDRT and wavelet transforms. Jafari-Khouzani and Soltanian-Zadeh (2005) employ the SDRT in combination with wavelet transforms for the purpose of extracting rotation-invariant information from images that contain textural patterns. Although this approach differs from the strategy employed in this thesis, it does present an interesting perspective on how the use of the SDRT may contribute towards successful texture classification. In order to convert rotational variations into translational variations, the SDRT is applied to the original image. The subsequent implementation of an efficient translation-invariant wavelet transform for the purpose of texture-based feature extraction successfully eliminates rotational variations. They also investigate the optimal number of projections required for accurate and efficient computation of the SDRT. The authors consider a database that contains two separate data sets. These data sets contain 25 and 24 texture images respectively. Maximum correct classification percentages of 97.90% and 97.40% are reported for the respective data sets.

The Gabor filter. Kekre and Bharadi (2010) propose that the Gabor filter (which is traditionally associated with texture recognition) is employed for the purpose of signature verification, albeit online signature verification. The Gabor filter is defined as a band-pass filter that extracts textural information pertaining to both orientation and frequency. A Gaussian envelope is typically employed for any Gabor filter. Said envelope is adjusted according to a sinusoidal function of a particular frequency that is orientated along a particular axis. The authors acquired their own data set that contains 250 signatures from 25 individual writers. The Euclidean distance is used for the purpose of feature matching and therefore differs from the DTW-based approach proposed in this thesis. The authors report an EER of 10.00% when no temporal information is taken into account. The research conducted by Kekre and Bharadi (2010) indicates that the extraction of textural information within the context of offline signature verification warrants further investigation. Such an investigation is conducted in this thesis.

2.4 Generic curve detection

Recall that the linear Hough transform (HT) detects straight lines in parameterised form as explained in Section 1.2.2. As mentioned in Section 1.2.2 the detection of other generic curves within the context of the HT is made possible by considering higher dimensional parameter spaces. The aforementioned ability however comes at the cost of a significant increase in computational complexity. A three-dimensional parameter space $c_1c_2c_3$ is for example required for detecting circles of the form $(x - c_1)^2 + (y - c_2)^2 = c_3^2$.

For the purpose of this thesis, appropriately normalised MPTs (that constitute grey-scale images) are however defined for the purpose of detecting generic curved lines, as will be explained in detail in Section 3.5. In order to achieve this, the SDRT is applied to the MPT in question. In this way, the detection of concentric circles with arbitrary origins is facilitated.

A number of noteworthy existing systems that use the conventional HT or adaptations thereof for the purpose of generic curved line detection is subsequently discussed. Said adaptations have been developed in order to exploit the valuable characteristics of the HT without the notorious computational complexity associated with the *conventional* HT. The review of the HT published by Illingworth and Kittler (1998) provides an in-depth exploration of its historical use, whereas the survey by Mukhopadhyay and Chaudhuri (2015) provides perspective on the latest advancements in research regarding the HT.

A generalised HT for arbitrary shape detection. Ballard (1981) introduced the concept of a generalised HT for shape detection. This approach was novel at the time in the sense that it is capable of detecting shapes regardless of whether the object boundary is connected or not. In order to generalise the HT, the author proposes the use of directional information contained within the object boundary, which leads to an increase in both efficiency and proficiency. For an arbitrary shape, the utilisation of a so-called *R*-table is proposed for representing the associated generalised HT.

Ellipse detection through the separation of HT parameter spaces. Nair and Saunders (1996) propose an algorithm for ellipse detection, by considering five parameters within the context of the HT. This is achieved by separating a five-dimensional parameter space into various sub-spaces, after which the parameters are calculated independently (and not simultaneously). This facilitates a significant decrease in computational complexity. The authors also propose a variation of the HT that aims to detect an ellipse in accumulation (parameter) space by searching for clusters instead of local max-

ima. This strategy facilitates a further decrease in computational complexity.

A segment-based HT for the detection of primitive curves associated with segments of an object boundary. Li *et al.* (2007) introduced a segment-based HT that is specifically geared towards the detection of primitive curves associated with segments of object boundaries in binary images. The edge pixels associated with the object boundary are first isolated, after which individual segments of said boundary are independently considered by the HT in order to decrease computational complexity. The proposed strategy is however only geared towards the detection of primitive curves such as straight lines and circular arcs. For a database containing more than 30 000 images, a recognition rate of more than 95% is reported. The aforementioned recognition rate is achieved despite a significant decrease in computational complexity.

2.5 Concluding remarks

In this chapter the systems proposed in this thesis were placed into perspective by outlining the architecture and proficiency of a number of existing systems that are in some ways related to our systems. The architecture of the systems proposed in this thesis is detailed in Chapters 3 and 4, while system proficiency is investigated and reported in Chapter 5.

Chapter 3

Feature extraction

3.1 Introduction

In order to effectively compare two images, we propose that so-called *global* feature vectors are extracted for the purpose of capturing the essence of the respective images in their *entirety*. These feature vectors are therefore constructed in such a way that they extract as much relevant information pertaining to the image in question as possible. This maximises the probability of accurate verification at a later stage.

The proposed feature extraction protocol is based on the well-known *standard* discrete Radon transform (SDRT). Said transform is geared towards the detection of straight lines in the original image (see Section 3.2). The SDRT has proved to be robust (with respect to noise, as well as variations in scale, position and orientation) (Wang *et al.*, 2007; Hjouj and Kammler, 2008). Furthermore, it was shown to be proficient in detecting skilled forgeries (within the context of signature verification) on numerous occasions (Coetzer *et al.*, 2004; Jayadevan *et al.*, 2009; Panton and Coetzer, 2010; Swanepoel, 2015).

In this thesis we aim to further exploit the above-mentioned attributes. The SDRT is adapted in such a way that not only straight lines are detected within the original image, but also curved lines (with various curvatures) and certain textural properties of the image in question.

The detection of the above-mentioned features is made possible by first constructing so-called *multiresolution* polar transforms (MPTs) of the original image (Section 3.3). A modified version of the SDRT, namely the *modified* discrete Radon transform (MDRT) may then be applied to an MPT in order to obtain textural features (Section 3.4). It is shown in Section 5.5 that classifier ensembles based on the above-mentioned features are either complementary to or outperform a classifier that is *solely* based on the SDRT.

When the SDRT is applied to different MPTs, each with an origin located at a different distance from the edge of the original image, different concentric curved lines may be detected (Section 3.5). When the origin of the MPT is

at a distance (from the edge of the image in question) that strives to infinity, the SDRT is effectively obtained. These features therefore constitute a *generalisation* of the SDRT that is able to detect curved lines with different curvatures. It is shown in Section 5.5 that classifier ensembles based on such a generalisation of the SDRT are either complementary to or outperform a classifier that is *solely* based on the SDRT. The proposed feature extraction protocol is conceptualised in Figure 3.1.

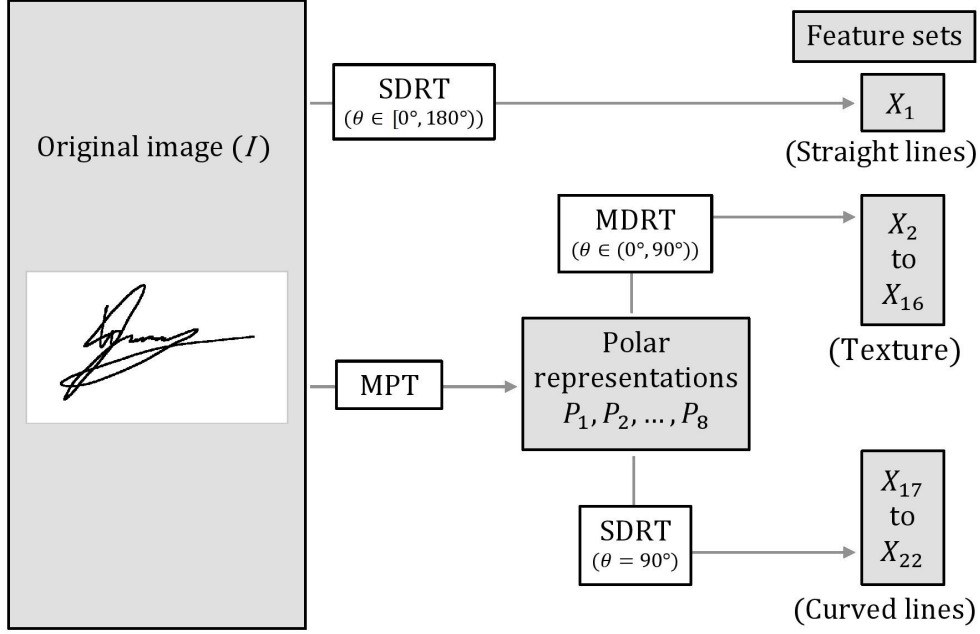


Figure 3.1: Overview of the proposed feature extraction protocol.

3.2 The standard discrete Radon transform (SDRT)

The SDRT is geared towards the detection of straight lines in an image. In order to achieve the aforementioned objective, projection profiles (R_θ) of the image in question, $I(x, y)$, are calculated from specified projection angles θ (see Figure 3.2).

An axis x' is defined that is orientated at a projection angle of θ with respect to the horizontal x -axis. A projection profile ($R_\theta(x')$) associated with said projection angle θ is then obtained by accumulating pixel values along parallel beams perpendicular to the x' -axis.

Eight evenly distributed projection angles are specified within half a revolution, that is $\theta \in [0^\circ, 180^\circ)$, since projection profiles obtained by considering the remaining angles, that is $\theta \in [180^\circ, 360^\circ)$, constitute reflections of the projection profiles already obtained.

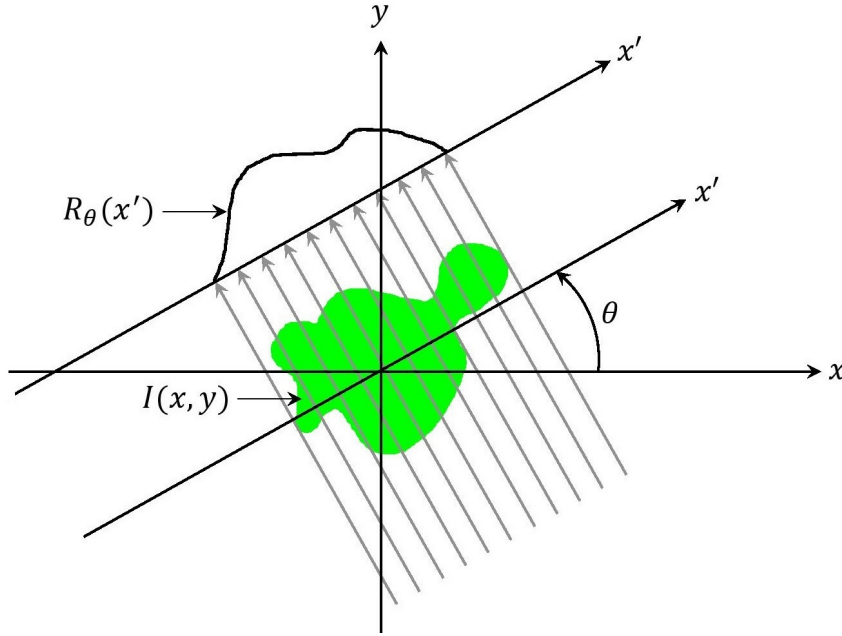


Figure 3.2: Geometric interpretation of the SDRT (Toft, 1996).

A specific feature vector constitutes a projection profile of the image obtained from one of the eight specified projection angles θ . The individual feature vectors are often stored as columns within a matrix, that is commonly referred to as a *sinogram*. When a sinogram is displayed as a grey-scale image, pixels with large intensity values coincide with prominent lines in the original image.

Feature set X_1 (see Figure 3.1) consists of appropriately normalised versions of eight different projection profiles (feature vectors) of the image in question. Feature normalisation is discussed in Section 3.6.

Due to its robustness and proficiency, the SDRT (within the context of feature set X_1) has become a benchmark for feature extraction, especially regarding offline signature verification (Coetzer *et al.*, 2004; Panton and Coetzer, 2010; Jayadevan *et al.*, 2009; Swanepoel, 2015). The SDRT is therefore adapted/modified in such a way that it is geared towards the detection of curved lines and textural patterns within an original image. We elaborate on this in the remaining sections of this chapter.

3.3 A multiresolution polar transform (MPT)

3.3.1 Definition of the MPT

In order to obtain information that pertains to curved lines and texture within a given image, a polar transform may be employed. The *standard* polar transform $P(r, \phi) = I(r \cos \phi, r \sin \phi)$ constitutes a transformed version in polar

coordinates of the original Cartesian image $I(x, y)$. It is obtained by considering a set of radii (different values of r) and angles (different values of ϕ) with respect to a specified origin ($r = 0$). The resolution of the polar transform, and therefore the number of pixels considered (visited) in the original Cartesian image is determined by a *sampling rate*.

The set of values for r may be specified in such a way that the interval between two consecutive values of r is given by $\Delta r = 1$ pixel. However, when the set of values for ϕ is specified in such a way that the interval between consecutive values of ϕ is *fixed* at $\Delta\phi = 1^\circ$, the arc length between consecutive samples on the perimeter of a circle with radius r is given by $\Delta s = r\Delta\phi$. An increase in r therefore results in a *linear* increase in the arc length Δs . This implies that, although $\Delta r = 1$ pixel, not all of the pixel values in the original image are considered (visited) since the arc length Δs determines the sampling rate for a specific value of r . In fact, for a given radius r , the sampling rate is inversely proportional to the arc length, and the sampling rate therefore becomes increasingly insufficient for larger values of r (see Figure 3.3). A substantial number of pixel values in the original image is therefore *not*

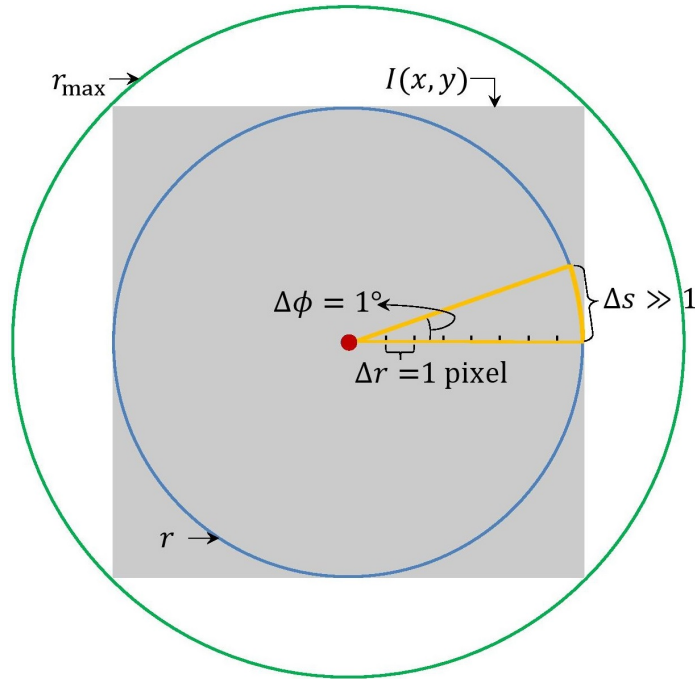


Figure 3.3: Conceptualisation of the standard polar transform. The arc length Δs increases when the radius r increases, which leads to an insufficient sampling rate for large values of r . The largest radius considered is denoted by r_{\max} .

considered (visited) for large values of r . This deficiency associated with the standard polar transform is conceptualised in Figure 3.4.

Since the goal is to extract as much information as possible from the original image, the standard polar transformation is therefore inadequate and has to be improved upon.

Note that unless otherwise specified, zero-valued pixels are rendered white, while one-valued pixels are rendered black. Within the context of static signature images, pen-strokes are therefore represented by black, one-valued pixels.

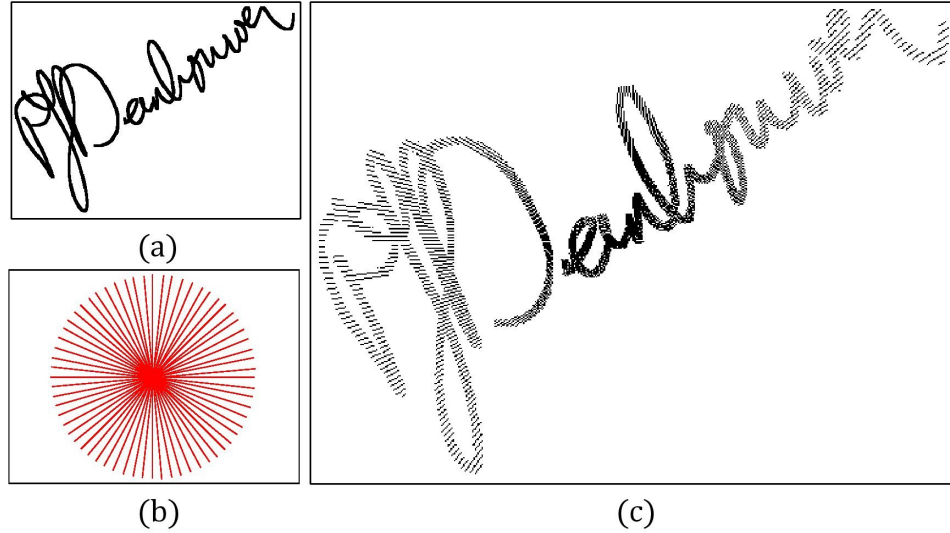


Figure 3.4: Illustration of the deficiencies associated with the standard polar transform. (a) The original image. (b) The pixels that are considered (visited) when the standard polar transform is employed. (c) The reconstructed image using information obtained from the standard polar transform. Numerous inaccuracies are observed for large values of r . The centre of the image coincides with $r = 0$.

A more refined approach is therefore proposed, namely the MPT. This transform involves the calculation of the polar transform in such a way that for each *different* specified value of r , a *different* set of values for ϕ is determined. More precisely, for larger values of r , a larger number of values for ϕ is specified. In order to ensure that all pixels in the original Cartesian image $I(x, y)$ are considered (visited), the constraint $\Delta s = r\Delta\phi \leq 1$ is imposed for all values of r , which implies that $\Delta\phi \leq 1/r$. In order to achieve this, and also avoid division by zero when $r = 0$, we propose that the interval between successive values of ϕ for a specific value of r is obtained as follows,

$$\Delta\phi(r) = \frac{1}{r+1} \text{ radians.} \quad (3.1)$$

Furthermore, the number of values (samples) for ϕ which has to be considered for a specific value of the radius r , that is $N_\phi(r)$, may be deduced by noting

the following,

$$\sum \Delta s = 2\pi r \implies \sum r \Delta \phi = 2\pi r \implies \sum \Delta \phi = 2\pi. \quad (3.2)$$

Therefore, when $\Delta \phi(r)$ is defined as in Equation (3.1), the following is obtained,

$$\sum_{\phi=\phi_0}^{N_\phi(r)-1} \frac{1}{r+1} = \frac{1}{r+1} \sum_{\phi=\phi_0}^{N_\phi(r)-1} 1 = N_\phi(r) \left[\frac{1}{r+1} \right] = 2\pi. \quad (3.3)$$

From Equation (3.3) we conclude that $N_\phi(r) = \lceil 2\pi(r+1) \rceil$, where $\lceil \cdot \rceil$ denotes the ceiling operator. It is therefore clear that when $N_\phi(r)$ distinct, equally spaced values of ϕ are specified for a specific value of r , the MPT ensures that $\Delta r = 1$ pixel *and* $\Delta s \leq 1$ pixel. All the pixels in Cartesian space are therefore considered, barring rounding errors (see Figure 3.5).

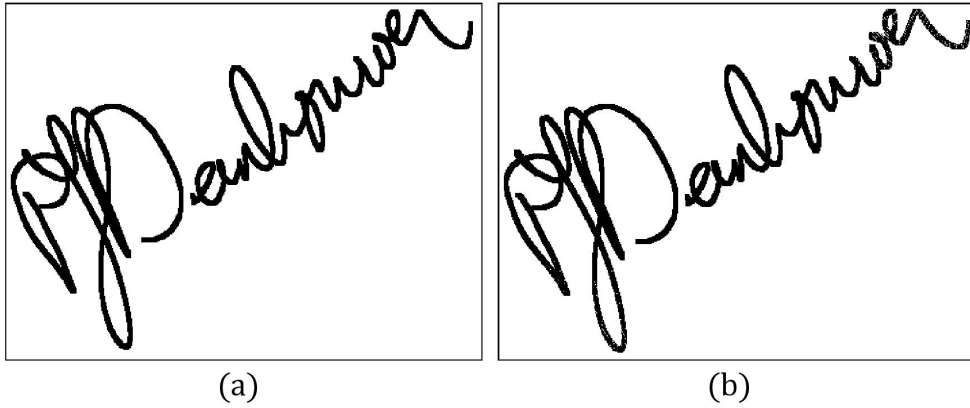


Figure 3.5: Illustration of the benefits associated with the MPT. (a) The original image. (b) The reconstructed image using information obtained from the MPT. The slight inaccuracies for large values of r are the result of rounding errors. The centre of the image coincides with $r = 0$.

3.3.2 Normalisation of the MPT

The largest value of r that is considered during the calculation of the MPT is denoted by r_{\max} and specified as the ceiling of half of the diagonal distance of the original image (see Figure 3.3). Therefore, the number of values for ϕ that coincide with $r = r_{\max}$ constitutes the largest number of values for ϕ (across all the specified values of r) that is considered during the calculation of the MPT. As a result, we opt to *fix* the second dimension, that is the number of columns T of the MPT, at $\lceil 2\pi(r_{\max} + 1) \rceil$ and specify that

$$T = N_\phi(r) = N_\phi(r_{\max}) = \lceil 2\pi(r_{\max} + 1) \rceil \quad (3.4)$$

for *all* values of r . Consequently, in order to store the MPT as a matrix, it is necessary to *duplicate* information when $r < r_{\max}$. This duplication of information may however place an undesired emphasis on pixels that are located close to the origin ($r = 0$). This emphasis is conceptualised in Figure 3.6. It is clear that the centre pixel (where the origin is located) is considered (visited) on at least four different occasions, while pixels that are located at a radius r are each considered only once. When $r = r_{\max}$ this problem becomes more apparent.

The MPT is therefore *normalised* in such a way that the sum of a subset of pixels in the MPT, that coincides with a *single* pixel in Cartesian space, equals one. In this way, a given pixel value in the original image is not emphasised at the cost of another.

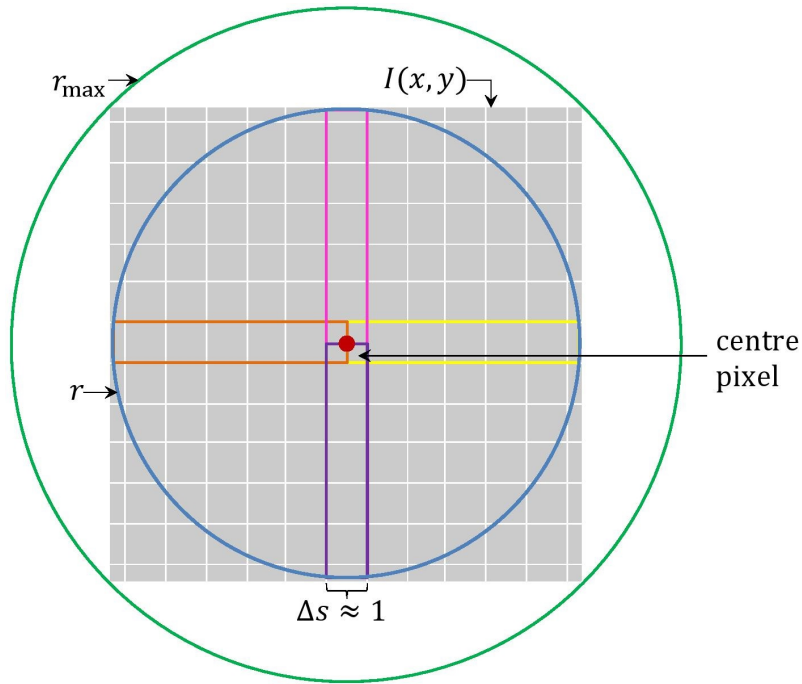


Figure 3.6: Conceptualisation of the MPT which illustrates how more emphasis is placed on pixels closer to the origin ($r = 0$). Normalisation is therefore required.

When the *unnormalised* version of the MPT is denoted by $M(r, \phi)$ and the *normalised* version by $M_{\text{norm}}(r, \phi)$, the normalisation process is conducted as follows,

$$M_{\text{norm}}(r, \phi) = \eta(r)M(r, \phi). \quad (3.5)$$

The weights $\eta(r)$ are defined as follows,

$$\eta(r) = \frac{\lceil 2\pi(r+1) \rceil}{f(r)\lceil 2\pi(r_{\max}+1) \rceil}, \quad (3.6)$$

where

$$f(r) = \frac{(1 - 2\pi)r}{r_{\max}} + 2\pi. \quad (3.7)$$

This ensures that a total weight of one is effectively assigned to all the pixels in the MPT that coincide with a single pixel in Cartesian space, and that a certain pixel is not emphasised at the cost of another.

The strategy employed for normalising the MPT is demonstrated through a practical example in Figure 3.7, while the key differences between the standard polar transform and the normalised MPT are illustrated in Figure 3.8. Note that effectively no normalisation is conducted when $r = r_{\max}$.

This carefully constructed MPT is considered in the subsequent feature extraction steps (Sections 3.4 to 3.6).

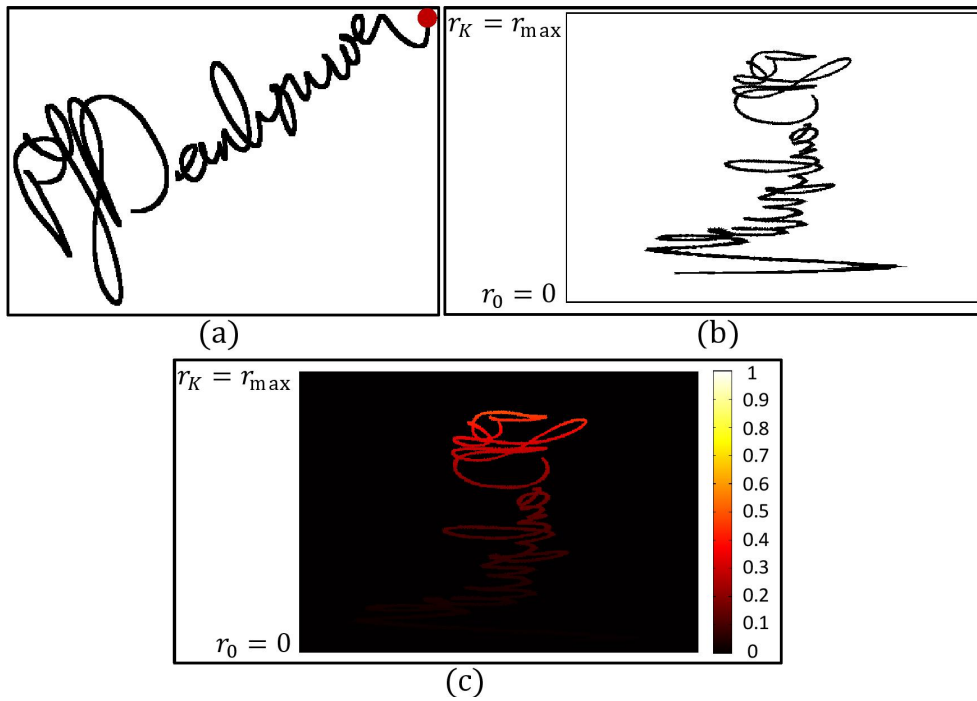


Figure 3.7: A practical illustration of the rationale behind the normalisation of the MPT. (a) The original image. (b) The unnormalised MPT of the image in (a) for which the origin ($r = 0$) is located at the top right corner of (a) as indicated by the red dot. The pixel values closer to the origin ($r = 0$) are clearly emphasised through duplication. (c) The normalised version of the MPT depicted in (b). The values of r and ϕ are specified as indicated in Figure 3.8 (b). This normalised version is displayed through the use of a colour map in order to effectively convey the difference in pixel values. This protocol therefore ensures that the undesired phenomenon of emphasising certain pixels at the cost of others is avoided by assigning smaller values (depicted by darker colours) to those (duplicated) pixels that are located closer to the origin ($r = 0$).

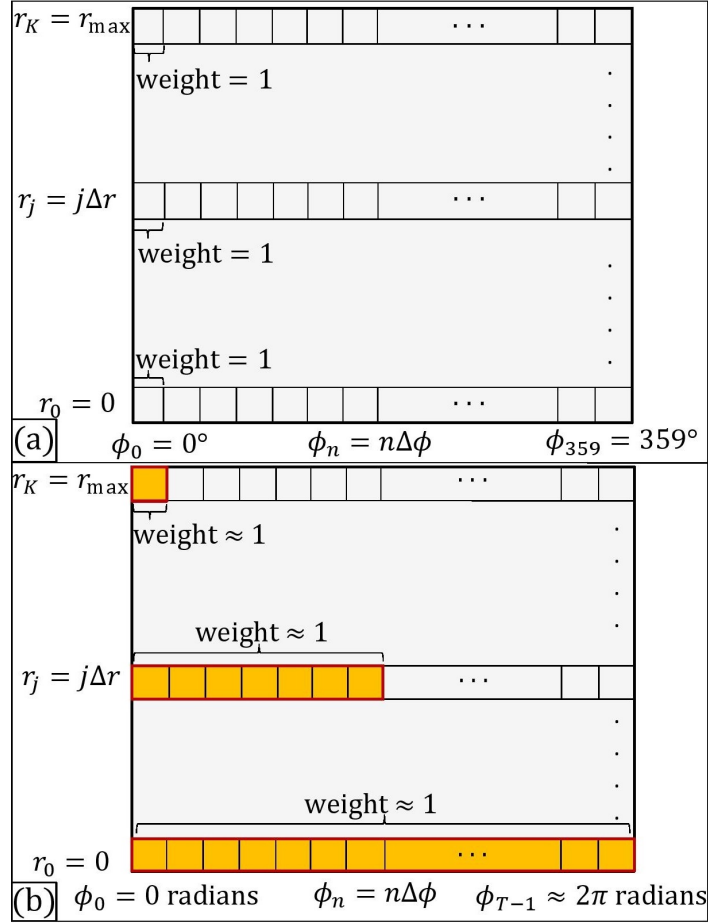


Figure 3.8: (a) Conceptualisation of an *unnormalised standard* polar transform that is comprised of 360 different angles and $K + 1$ different radii. The interval between consecutive angles is given by $\Delta\phi = 1^\circ$, while the interval between consecutive radii is given by $\Delta r = 1$ pixel. Each entry in the transform (denoted by a black square) is effectively assigned a weight of one. (b) Conceptualisation of a *normalised MPT* that is comprised of T different angles and $K + 1$ different radii, with T defined as in Equation (3.4). The interval between consecutive angles $\Delta\phi$ is dependent on the radius and defined in Equation (3.1), while the interval between consecutive radii is given by $\Delta r = 1$ pixel. The collection of entries in the transform that are boxed in red and highlighted in yellow coincide with a single pixel in the original image and is assigned a combined (summed) weight of approximately one. Therefore, when $r = 0$, each individual entry (denoted by a black square) is assigned a weight of approximately $1/T$. However, when $r = r_{\max}$, each individual entry is assigned a value of one.

3.4 Modification of the SDRT for texture detection

3.4.1 Detection of spirals via the MPT

When the SDRT is applied to an MPT from a projection angle of θ as depicted in Figure 3.9, where $\theta \in (0^\circ, 90^\circ)$, a single *beam* associated with the SDRT coincides with a single *spiral* in the original image space (Cartesian space). This is due to the fact that along such a beam both r and ϕ vary. Furthermore, when different values of the projection angle θ are considered, different *types* of spirals are obtained. In particular, when the value of the projection angle θ increases from 0° to 90° in the interval $(0^\circ, 90^\circ)$, the associated spirals inevitably become “sparser” (see Figure 3.10). Note that the spirals in Figure 3.10 (a), associated with the scenario where $\theta = 22.5^\circ$, do not cover the entire range of radii in the interval $r \in [0, r_{\max}]$ and therefore extract a sub-optimal amount of information. This issue will be addressed in Section 3.4.3.

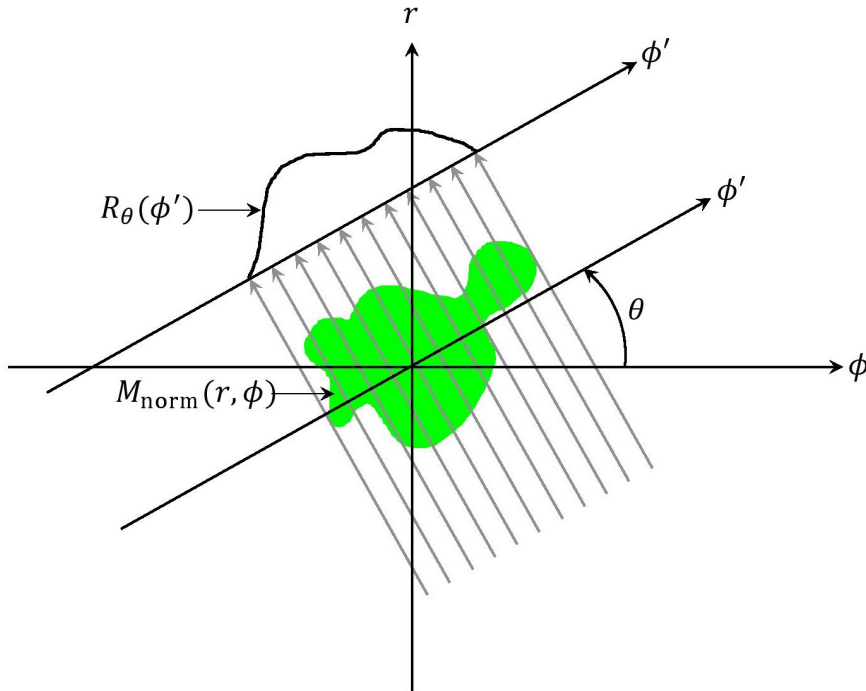


Figure 3.9: Geometric interpretation of the SDRT when applied to a normalised MPT.

Note that only angles in the interval $\theta \in (0^\circ, 90^\circ)$ need to be considered (see Figure 3.11), since angles in the interval $\theta \in (90^\circ, 180^\circ)$ coincide with *similar* spirals as those for $\theta \in (0^\circ, 90^\circ)$, except for the fact that the spirals associated with $\theta \in (0^\circ, 90^\circ)$ rotate *clockwise* as r increases, while those for $\theta \in (90^\circ, 180^\circ)$

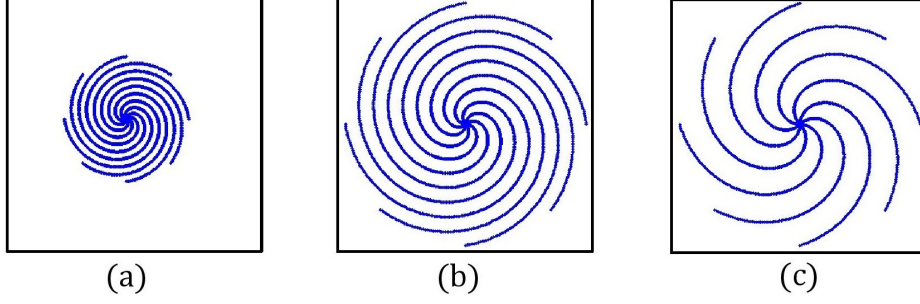


Figure 3.10: Conceptualisation of spirals that coincide with distinct beams associated with the SDRT when applied to an MPT, where **(a)** $\theta = 22.5^\circ$, **(b)** $\theta = 45^\circ$ and **(c)** $\theta = 67.5^\circ$. The spirals conceptualised here are obtained specifically when $r = 0$ is chosen to be at the centre of the original image. Note that the spirals for $\theta = 22.5^\circ$ do not consider (visit) all radii in the interval $r \in [0, r_{\max}]$. The information extracted by these spirals is therefore insufficient.

rotate *counterclockwise*. No additional information is therefore extracted by also considering the aforementioned spirals since the summed value acquired via the SDRT is independent of the direction (clockwise/counterclockwise) of summation within the context of the protocol explained in Section 3.4.2.

Similarly, since projections calculated from angles in the interval $\theta \in [180^\circ, 360^\circ)$ constitute *reflections* of those projections calculated from angles in the interval $\theta \in [0^\circ, 180^\circ)$, the former projections are also considered redundant.

For the specific scenarios where $\theta = 0^\circ$ and $\theta = 90^\circ$, the beams of the SDRT do not coincide with spirals at all, but rather with radiating straight lines and concentric circles respectively, therefore rendering these angles unsuitable for the purpose of texture detection in the current context (see Figure 3.11).

It is important to note that the conceptualisation in Figure 3.10 is only applicable to the specific case where the origin of the MPT (that is where $r = 0$) coincides with the centre of the original image. This brings us to the point of identifying two major *deficiencies* associated with the concept of simply applying the SDRT to a normalised MPT for the purpose of feature extraction:

1. Firstly, in order to facilitate the construction of various base classifiers within a classifier ensemble, various origins (where $r = 0$) may be chosen in calculating the MPT. The centre of the original image may not be a good choice for such an origin since pertinent information is subsequently associated with small radii. In fact, we rather opt to select the aforementioned origins either at points along the *edge* of the original image or at points located *outside* of the original image. Large portions of the resulting spirals will therefore inevitably lie *outside* the original image, rendering these portions of the spirals unable to extract any sig-

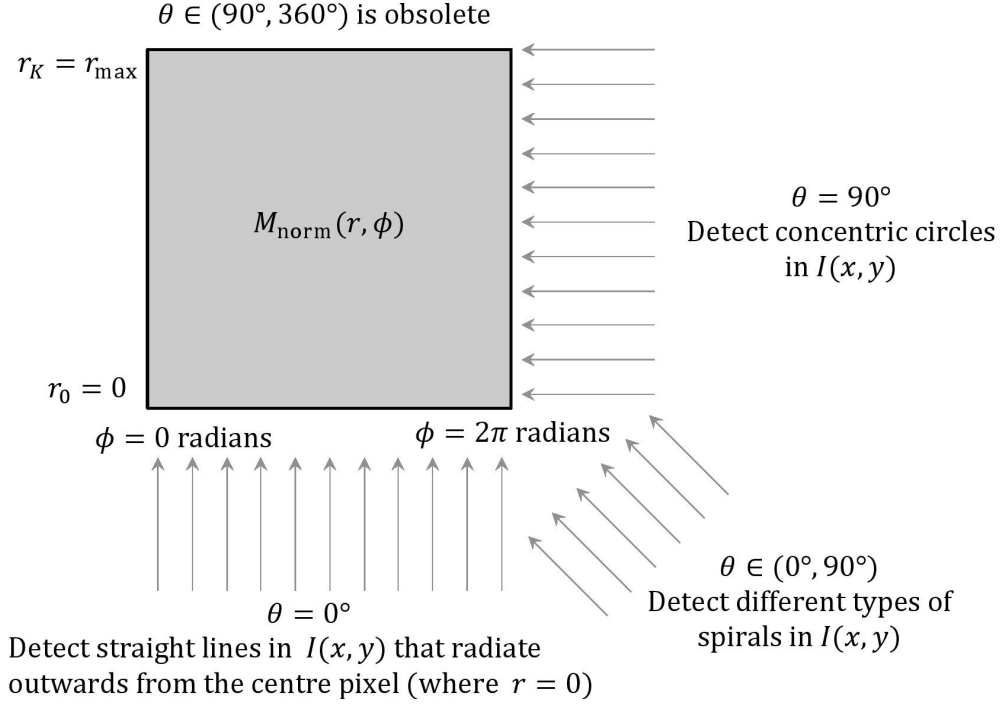


Figure 3.11: Conceptualisation of the SDRT being applied to a normalised MPT. The relevant targeted curves in the original image $I(x, y)$ are specified for different values of the projection angle θ .

nificant information. This deficiency may be rectified by constructing a composite image that is comprised of the original image, as well as appended reflections thereof. This protocol is discussed in more detail in Section 3.4.2.

2. Secondly, it is clear from Figures 3.10 and 3.12 that when the SDRT is applied to a normalised MPT at projection angles in the interval $\theta \in (0^\circ, 90^\circ)$, a specific beam never covers *both* the entire range of radii (that is $r \in [0, r_{\max}]$) *and* the entire range of polar angles (that is $\phi \in [0^\circ, 360^\circ)$). The only notable exception is the centre beam associated with the projection angle $\theta = 45^\circ$. This beam is conceptualised by the green line in Figure 3.12. Note that very little information is extracted by the blue, pink and brown lines which are associated with projection angles of $\theta = 22.5^\circ$, $\theta = 45^\circ$ and $\theta = 67.5^\circ$, respectively. In order to ensure that every beam covers at least all the radii in the interval $r \in [0, r_{\max}]$, we propose a so-called *modified* discrete Radon transform (MDRT) that considers a number of duplicates of the MPT in the ϕ -dimension. This is discussed in more detail in Section 3.4.3.

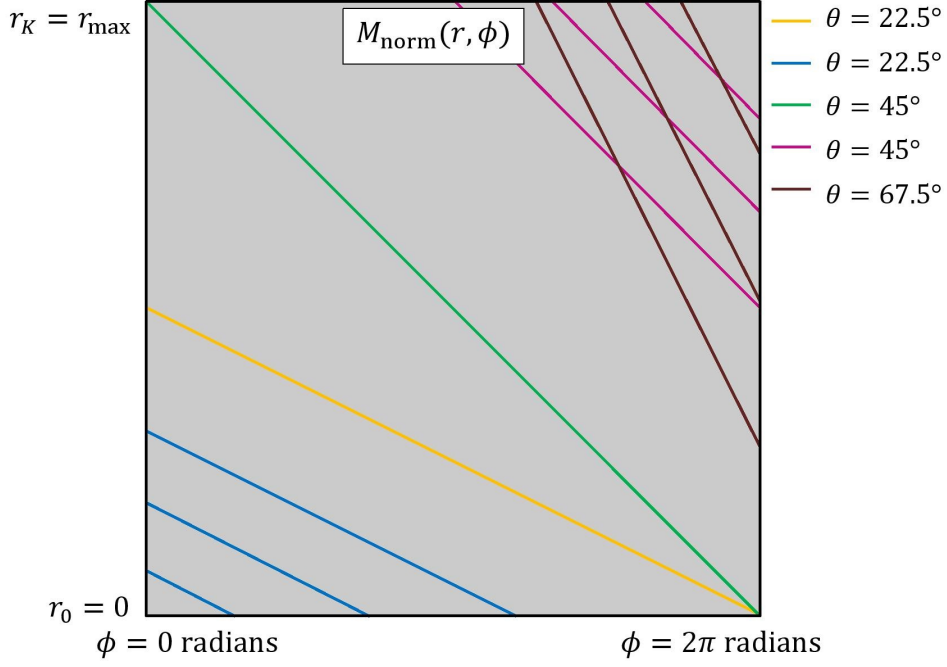


Figure 3.12: Conceptualisation of a selected number of beams (associated with the SDRT) that extract information from a normalised MPT. Each beam coincides with a segment of a specific spiral in Cartesian space. Only the green beam extracts information across all radii (that is $r \in [0, r_{\max}]$) and all polar angles (that is $\phi \in [0^\circ, 360^\circ]$). The other beams extract relatively little information.

3.4.2 Reflections of the original image

In order to facilitate the construction of various base classifiers within a classifier ensemble, we opt to specify different locations for the origin (that is where $r = 0$) of the MPT as conceptualised in Section 3.3. Said locations also coincide with the origins of different types of spirals in Cartesian space as explained in the previous section. It is furthermore proposed that eight different origins are specified, where each origin is either located along the edge of the original image or outside said image. Large portions of the corresponding spirals are consequently located outside the image in question. This is undesirable since we aim to extract as much information as possible from the image by effectively adding up pixel values along these spirals.

In order to facilitate the extraction of as much information as possible from an image, we opt to construct a composite image that not only contains the original image but also several reflections thereof. The manner in which these reflections are appended to the original image depends on where the origin of the MPT ($r = 0$) is located. The eight specified origins are depicted in Figure 3.13 (a). The respective protocols for constructing suitable composite images by appending reflected versions of the original image to said image, for each of the specified origins, are conceptualised in Figure 3.13 (b) to (i). In

each of the aforementioned scenarios, the reflected images are appended to the original image in such a way that the centre of the composite image coincides with the specified origin as pertaining to the original image. The centre of each composite image also coincides with the respective origins as pertaining to the reflections of the original image.

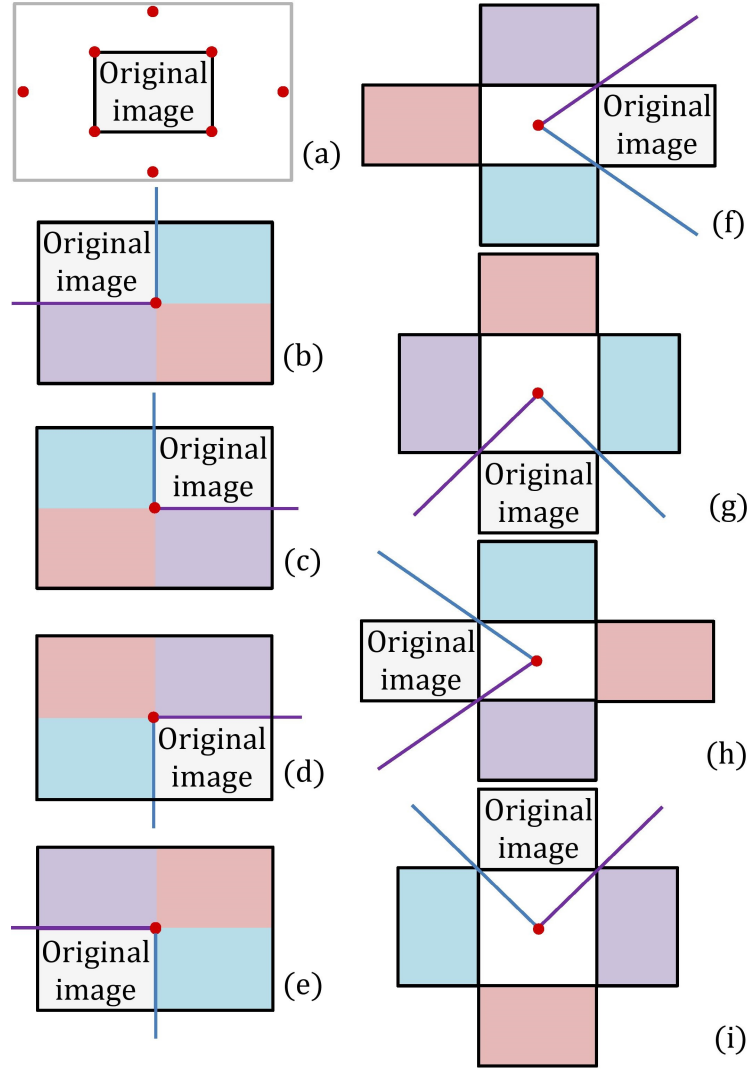


Figure 3.13: The proposed protocol for constructing composite images. The original image and the eight specified origins for the MPT (indicated by red dots) are depicted in (a). The specified origins and the corresponding composite images are depicted in (b) to (i). For each of the aforementioned scenarios, the original image is reflected through the purple axis to obtain the reflected image depicted in purple. Similarly, the original image is reflected through the blue axis to obtain the reflected image depicted in blue, while the original image is reflected through the selected origin (denoted by the red dot) to obtain the reflected image depicted in pink.

Three examples of how the above-mentioned composite images are constructed are presented in Figure 3.14, specifically when the origin of the MPT ($r = 0$) is located at three different points, either along the edge of the original image or outside of the original image.

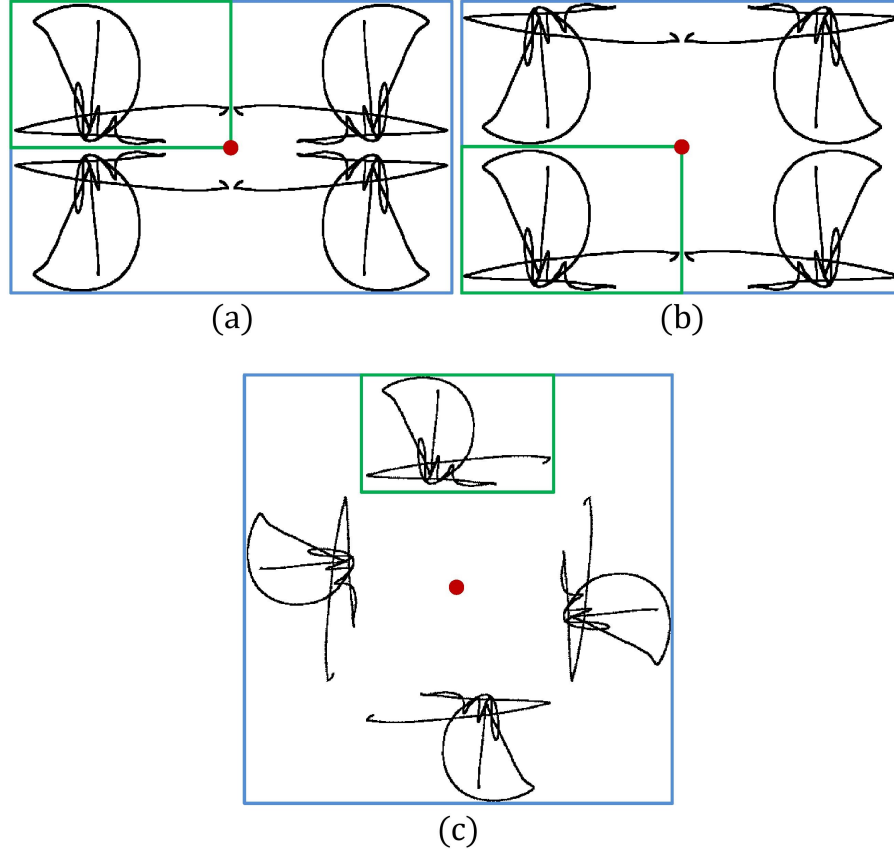


Figure 3.14: Examples of how composite images are constructed when the origin of the MPT ($r = 0$) is chosen to be at (a) the bottom right corner of the original image, (b) the top right corner of the original image, and (c) the point midway between the two bottom corners of the original image. In each case, the composite image (boxed in blue) consists of the original image (boxed in green) and three reflections thereof. The relevant origin ($r = 0$) is denoted by a red dot.

The MPT may now be applied to any of these newly constructed composite images, with the origin of the MPT ($r = 0$) located at the *centre* of the *composite* image in question. The SDRT may then be applied to such an MPT, thereby effectively extracting information from the original image via segments of spirals that are completely contained within the composite image (see Figure 3.15).

Although this strategy guarantees that segments of spirals will not be located outside the image, it does not guarantee that these spirals will cover the entire set of specified radii ($r \in [0, r_{\max}]$). This is especially evident for the

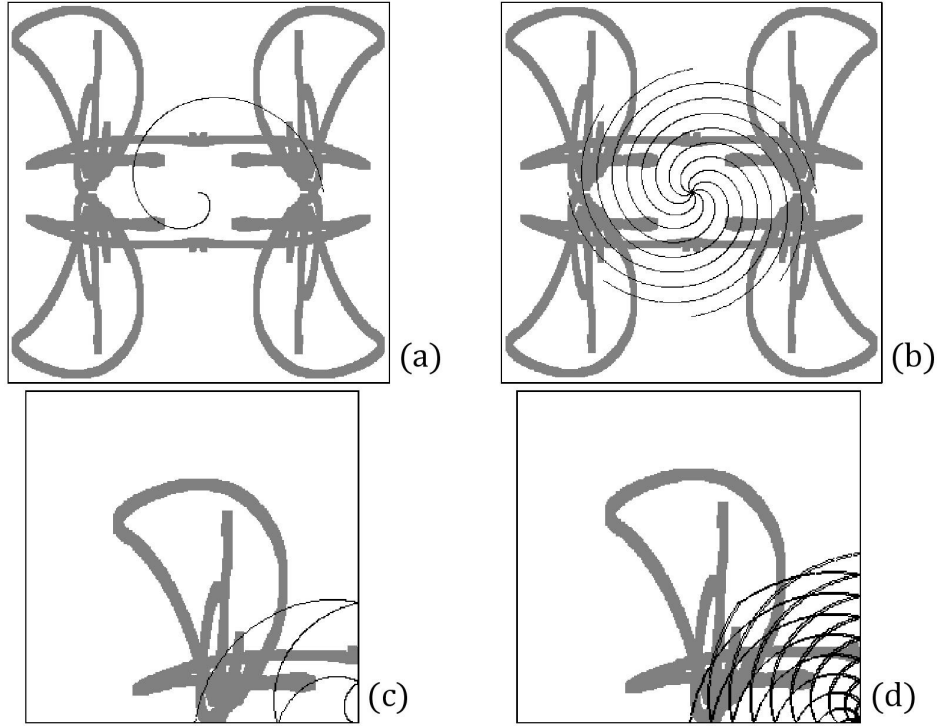


Figure 3.15: (a) A spiral that coincides with a distinct beam associated with the SDRT that is applied to an MPT for the projection angle $\theta = 22.5^\circ$. The composite image is constructed in the same way as the one depicted in Figure 3.14 (a). The origin therefore coincides with the bottom right corner of the original image. (b) The same scenario as the one depicted in (a). In this case, eight spirals that coincide with eight different parallel beams are shown. (c) The same scenario as the one depicted in (a). In this case, those segments of the spiral that are located in the reflected images are superimposed onto the original image. This indicates that (in reality) pixel values within the original image are summed along a path that resembles (in shape) a set of shark fins. (d) The same scenario as the one depicted in (c), but analysed using the protocol in (b). The eight different parallel beams of the SDRT are now associated with eight different paths, each one resembling a set of shark fins.

less sparse (more tightly packed) spirals associated with projections calculated from relatively small projection angles θ (see Figure 3.10 (a) and Figure 3.15). It is therefore proposed that the SDRT be *modified* in order to improve the feature extraction protocol via the MPT by ensuring that at least all the radii in the interval $r \in [0, r_{\max}]$ are covered (visited). This strategy is explained in greater detail in the following section.

3.4.3 The modified discrete Radon transform (MDRT)

When the SDRT is applied to an MPT for relatively small values of the projection angle θ , information is extracted from a composite image (like the ones

depicted in Figures 3.13 and 3.14) by adding up pixel values along relatively dense spirals. This phenomenon can be easily explained by considering the yellow beam in Figure 3.12 for which $\theta = 22.5^\circ$. Along said beam, the radius r increases less rapidly than the polar angle ϕ , which results in the fact that, although all angles in the interval $\phi \in [0^\circ, 360^\circ)$ are covered, only a relatively small subset of the radii in the interval $r \in [0, r_{\max}]$ is visited. This problem is evident in Figure 3.10 (a) and Figure 3.15. The associated projection angle for both of these examples is also $\theta = 22.5^\circ$. The essence of this problem is that, for any particular spiral, a maximum of one revolution is possible within the context of the SDRT.

In order to allow a spiral to complete more than one revolution and therefore cover the entire range of radii $r \in [0, r_{\max}]$, a normalised MPT has to be *duplicated* at least once. This leads to the concept of a modified discrete Radon transform (MDRT) which is now explained in greater detail.

Since the MPT is periodic in the ϕ -direction, one or more duplicates thereof may be concatenated to the original MPT-matrix as illustrated in Figure 3.16. When the SDRT of this concatenated matrix is obtained, where said matrix contains a sufficient number of duplicates, it should be clear from Figure 3.16 that a single beam is now able to cover the entire range of radii $r \in [0, r_{\max}]$. This therefore facilitates the possibility that the corresponding spiral (in Cartesian space) completes more than one revolution, thereby covering the entire composite image. Per definition, the MDRT therefore constitutes the SDRT of the periodic continuation of a normalised MPT.

For the example shown in Figure 3.16, the projection angle is given by $\theta = 22.5^\circ$. In this example, only half of the total range of radii is covered by the first revolution which only involves the original normalised MPT, but the presence of the concatenated duplicate of said MPT ensures that the entire range of radii is covered.

Similarly, for projection angles other than $\theta = 22.5^\circ$, the concatenation of an appropriate number of duplicates to the original MPT will ensure that the entire range of radii is covered. Figure 3.17 illustrates how the MDRT facilitates the coverage of the entire range of radii. The improvement is clear when Figure 3.17 is compared to Figure 3.15.

For the purpose of this thesis, projection angles smaller than 11.25° will not be considered. A fixed number of three duplicates will therefore be appended to the original normalised MPT.

Whenever we refer to the MDRT of an image in the remainder of this thesis, it is implied that the MDRT is applied to the composite version of said image as explained in Figure 3.13.

3.4.4 Texture detection

In this section we explain how features that are indicative of textural information within an image may be inferred from its MDRT. For this purpose,

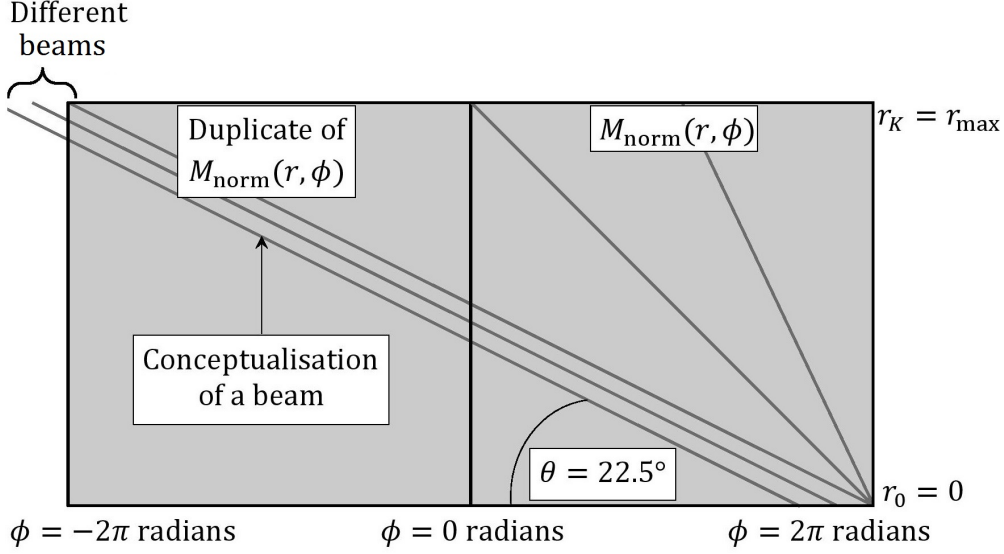


Figure 3.16: Geometric interpretation of the MDRT. For a projection angle of $\theta = 22.5^\circ$, the concatenation of a single duplicate of the original normalised MPT is sufficient to ensure that the entire range of radii is covered.

we refer to Figure 3.17 in which case the projection angle associated with the MDRT is given by $\theta = 22.5^\circ$. We concluded in the previous section that a specific beam associated with the MDRT coincides with a curve in Cartesian space that resembles a set of shark fins (see Figure 3.17 (c)). Through applying the MDRT, the pixels in the original image are therefore summed along this curve in order to obtain a specific feature value within a feature vector. The other features within said feature vector are associated with beams that are parallel to the beam in question.

Note that the shark fins in Figure 3.17 (c) is comprised of two distinct sets of concentric curves as conceptualised in orange and green in Figure 3.18.

The distance between consecutive concentric curves within a given set may be interpreted as a “period” which is inversely proportional to a “frequency”. It should be clear that for larger values of the projection angle θ , the associated spirals are more sparse (see Figure 3.10), which coincides with a smaller frequency and therefore a coarser texture.

Different parallel beams that are associated with the same projection angle are therefore geared towards the detection of textures with a similar coarseness while allowing for variations in phase. In order to compensate for these variations, corresponding feature vectors that are extracted from different images are matched through dynamic time warping (DTW) as detailed in Section 4.2.

Textural information may now be extracted from an image in two different ways:

1. Specify different projection angles θ for the MDRT. These feature vectors

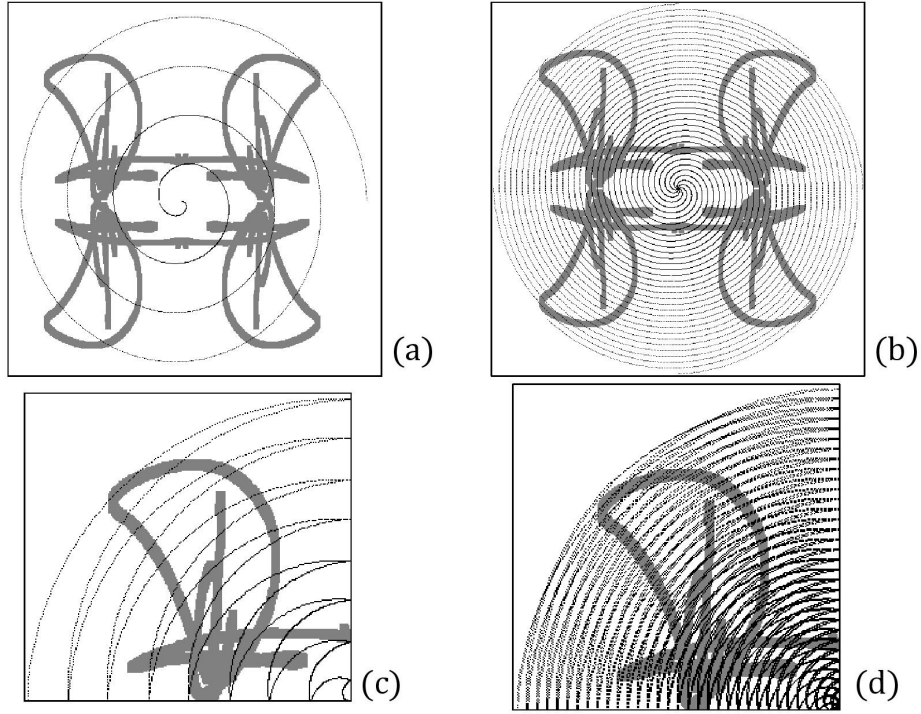


Figure 3.17: (a) A spiral that coincides with a distinct beam associated with the MDRT for the projection angle $\theta = 22.5^\circ$. The composite image is constructed in the same way as the ones depicted in Figure 3.14 (a). The origin therefore coincides with the bottom right corner of the original image. (b) The same scenario as the one depicted in (a). In this case, eight spirals that coincide with eight different parallel beams are shown. (c) The same scenario as the one depicted in (a). In this case, those segments of the spiral that are located in the reflected images are superimposed onto the original image. This indicates that (in reality) pixel values within the original image are summed along a path that resembles (in shape) a set of shark fins. (d) The same scenario as the one depicted in (c), but analysed using the protocol in (b). The eight different parallel beams of the MDRT are now associated with eight different paths, each one resembling a set of shark fins.

are therefore based on the *periodicity* or *coarseness* of the texture.and

2. Specify different origins for the MPT (that is where $r = 0$), either along the edge of the original image or outside said image. These feature vectors are therefore based on the *orientation* of the textural pattern.

Figure 3.19 provides an overview of the proposed texture extraction protocol.

The first approach towards texture detection involves the extraction of *seven* different *feature sets* X_i , $i = 2, \dots, 8$ based on seven different values of the projection angle θ , therefore representing seven distinct periodicities. The *eight* distinct *feature vectors* that form part of a specific feature set are associated with eight different locations of the origin (that is where $r = 0$) of

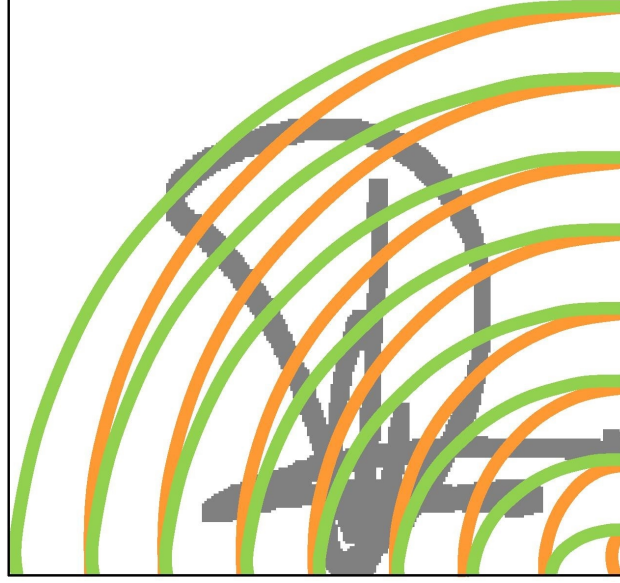


Figure 3.18: The concentric curves that form part of a shark fin shaped path in Cartesian space. The curves represented in green and orange constitute two distinct sets of concentric paths. For this specific case, the origin of the MPT ($r = 0$) is located at the bottom right corner of the original image and the projection angle associated with the MDRT is given by $\theta = 22.5^\circ$.

the MPT. Said origins are located either along the edge of the original image or outside said image. The distance between two images, based on one of these feature sets, is defined as the average DTW-based distance between the corresponding feature vectors. The seven feature sets X_i , $i = 2, \dots, 8$ will be associated with seven different base classifiers in a classifier ensemble. We elaborate on the relevant DTW-based algorithm for feature matching, as well as the protocol for ensemble construction in Sections 4.2 and 4.5 respectively.

The second approach towards texture detection involves the extraction of *eight* different *feature sets* X_i , $i = 9, \dots, 16$ based on eight different origins (that is where $r = 0$) of the MPT, where each origin is either located along the edge of the original image or outside said image, therefore representing eight different orientations. The *seven* distinct *feature vectors* that form part of a specific feature set are associated with seven different values of the projection angle θ , therefore representing seven distinct periodicities. The distance between two images, based on one of these feature sets, is defined as the average DTW-based distance between the corresponding feature vectors. The eight different feature sets X_i , $i = 9, \dots, 16$ will be associated with eight different base classifiers in a classifier ensemble.

As previously mentioned, the beams associated with the SDRT when applied to an MPT for projection angles of $\theta = 0^\circ$ and $\theta = 90^\circ$ do not coincide with spirals in Cartesian space and are therefore not employed for the purpose

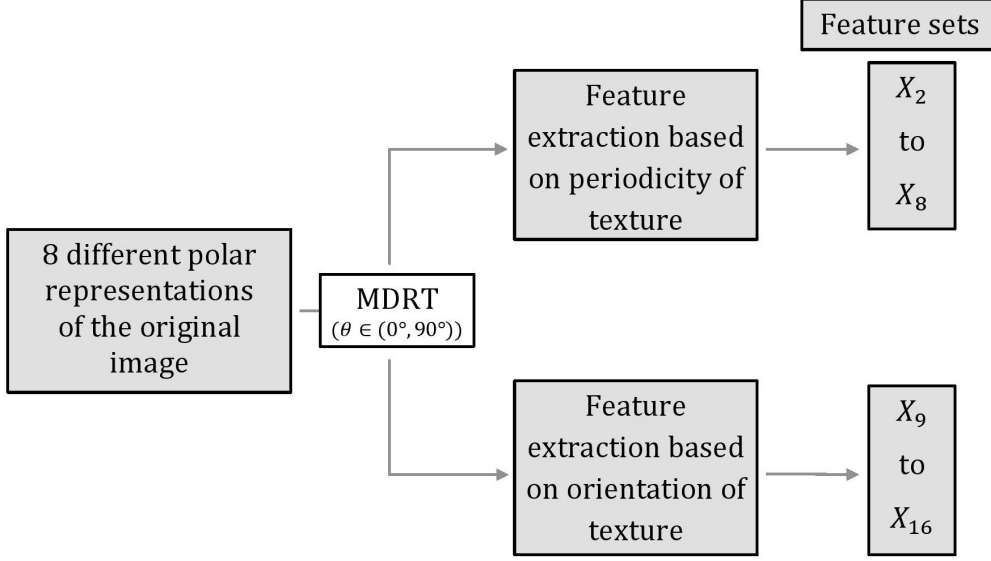


Figure 3.19: Overview of the proposed feature extraction protocol that is based on the detection of textural information.

of texture detection. In particular, we note that for $\theta = 0^\circ$ the associated beams coincide with straight lines in Cartesian space that radiate outwards from the chosen origin (where $r = 0$). Furthermore, for $\theta = 90^\circ$, the associated beams coincide with concentric circles in Cartesian space which are centred at the chosen origin. The scenario where $\theta = 90^\circ$ is however valuable within the context of curved line detection as we explain in the subsequent section.

3.5 Generalisation of the SDRT for curved line detection

For the specific scenario within the context of the SDRT in which the projection angle is given by $\theta = 90^\circ$, the beams associated with the SDRT when applied to an MPT coincide with *concentric circles* in the original image space (Cartesian space). This is attributed to the fact that each horizontal line (associated with a projection angle of $\theta = 90^\circ$) within an MPT corresponds to a single radius r and all polar angles ϕ . In this scenario the associated “spirals” have a maximum curvature resulting in concentric circles. This concept is illustrated in Figure 3.20 (a) for the specific scenario where the origin of the MPT is located at the bottom right corner of the original image. Only a fraction (25%) of the concentric circles are therefore located within the image.

As per usual, the interval between consecutive radii within the context of the MPT is given by $\Delta r = 1$ pixel. The origin of the MPT (where $r = 0$) may now be progressively shifted further away from the edge of the image while

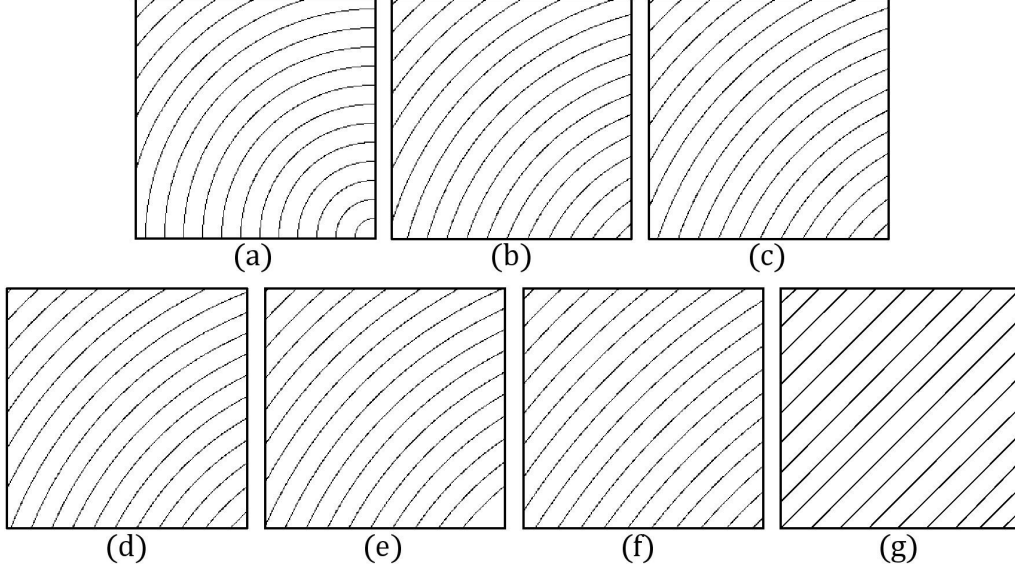


Figure 3.20: Conceptualisation of curved lines that may be detected when the origin (that is where $r = 0$) of the MPT is located on the line that passes through the centre and the bottom right corner of the image. The distance d between the origin of the MPT and the bottom right corner is specified as follows: **(a)** $d = 0$ pixels, **(b)** $d = 70$ pixels, **(c)** $d = 140$ pixels, **(d)** $d = 210$ pixels, **(e)** $d = 280$ pixels, **(f)** $d = 350$ pixels and **(g)** $d \rightarrow \infty$.

still being located on the line that passes through the centre of the image and the bottom right corner. When the constraint that $\Delta r = 1$ pixel is imposed throughout the process, the concentric curves depicted in Figure 3.20 (b) to (g) are obtained. When d is defined as the distance between said origin and a corner of the image, the curvature of the concentric curves progressively decreases as d increases. Straight lines are effectively obtained when $d \rightarrow \infty$, in which case the SDRT is effectively calculated. This strategy may therefore be perceived as a generalisation of the SDRT.

Note that the construction of a composite image that contains reflections of the original image is not required for a projection angle of $\theta = 90^\circ$ since a circle is completed within a single revolution. Furthermore, the duplication of the MPT is not required either, since that would result in a repetition of information. The application of the SDRT is therefore sufficient for extracting the features discussed in this section and the MDRT as defined in Section 3.4.3 is not required within the current context.

We opt to extract features in this way by first specifying that the origins (that is where $r = 0$) of the MPTs are located at 16 different evenly distributed locations on an ellipse that encloses the original image. The origins of these MPTs are therefore specified in a *different* way than was the case for the MPTs associated with the extraction of textural information. By adopting this protocol curves with 16 different orientations may be detected, of which eight are

deemed concave (see Figure 3.21 (a)) and eight convex (see Figure 3.21 (b)). The detection of curves with different curvatures is made possible by specifying different values for d as conceptualised in Figure 3.20.

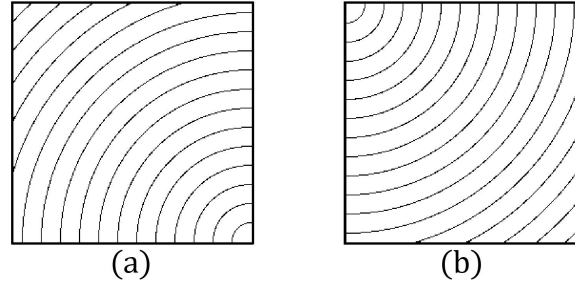


Figure 3.21: Conceptualisation of how curved lines with identical curvatures may be defined as convex or concave based on their orientation. The curved lines in (a) are concave, while those in (b) are convex with respect to the bottom right corner.

The benchmark SDRT is associated with the special case where $d \rightarrow \infty$. Feature set X_1 therefore consists of *eight* different *feature vectors*, where each feature vector constitutes a projection of the original image from a certain projection angle θ (see Figure 3.1). Within the context of the benchmark SDRT, only eight feature vectors are required since concavity or convexity are not attributed to straight lines. When six other values of d are specified, *six other feature sets*, X_i , $i = 17, \dots, 22$, may be constructed. The *sixteen* different *feature vectors* that constitute a specific feature set are associated with sixteen different locations of the origin (that is where $r = 0$) of the MPT. Said origins are located at evenly distributed intervals on an ellipse that is concentric to the innermost ellipse which passes through the four corners of the original image (see Figure 3.22). The innermost ellipse is associated with $d = 0$ pixels and feature set X_{17} . The subsequent ellipses are associated with $d = 70$ pixels, $d = 140$ pixels, $d = 210$ pixels and $d = 280$ pixels respectively (and therefore feature sets X_{18} , X_{19} , X_{20} and X_{21}), while the outermost ellipse is associated with $d = 350$ pixels and feature set X_{22} .

Figure 3.23 provides an overview of the proposed feature extraction protocol that is based on the detection of curved lines within an image. The distance between two images, based on one of these feature sets, is defined as the average DTW-based distance between the corresponding feature vectors. The above-mentioned six feature sets will be associated with six different base classifiers within a classifier ensemble as we explain in Sections 4.1, 4.2 and 4.5.

3.6 Feature normalisation

Feature vectors are typically normalised in such a way that they constitute scale, translation and rotation invariant representations of an image. How-

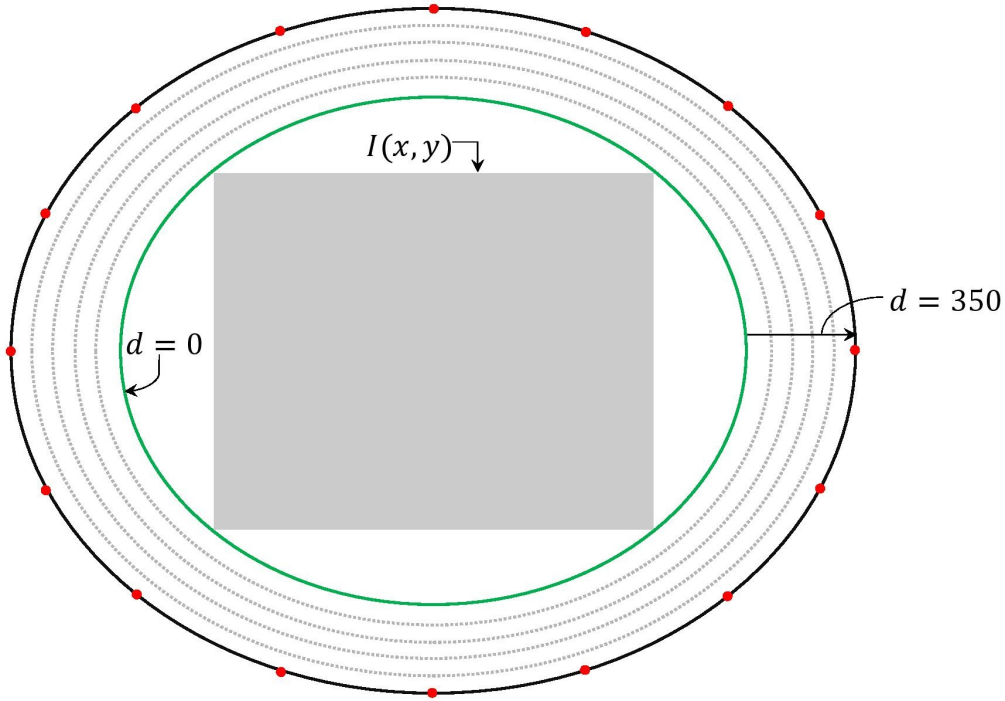


Figure 3.22: Conceptualisation of the concentric ellipses associated with different values of d . For a specific value of d , the origins of the respective MPTs are located at 16 evenly distributed positions on said ellipse. These positions are indicated by red dots for the specific ellipse associated with $d = 350$ pixels.

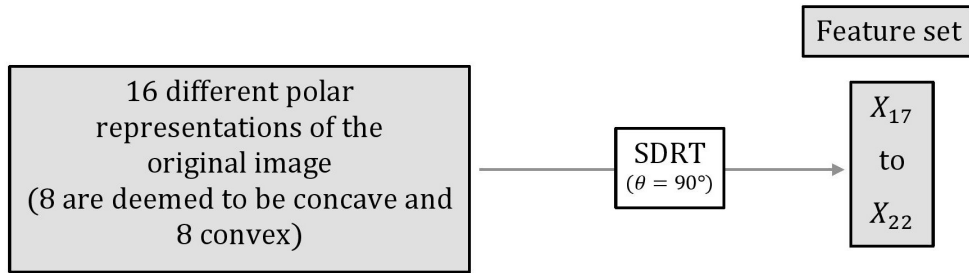


Figure 3.23: Overview of the proposed feature extraction protocol that is based on the detection of curved lines.

ever, within the context of detecting skilled forgeries among handwritten offline signatures, it is reasonable to assume that the forgeries are of decent quality. Said forgeries are therefore not expected to differ substantially from genuine reference samples with regards to their orientation. Within the scope of this thesis we therefore only aim to render the extracted feature vectors scale and translation invariant. Should the proposed systems however be applied to other types of data, we may opt to also enforce rotation invariance. In Section 6.2.1 we explain how rotation invariant features may be obtained

should this be the objective of another application. For now it suffices to note that feature sets X_{17} to X_{22} are already rotation invariant since the MPT converts rotational variations into translational variations. When the SDRT is therefore applied to the aforementioned MPT, rotational variations within the original image are effectively removed through a subsequent normalisation that ensures translation invariance.

Scale and translation invariance are achieved as follows. Since ideal (noiseless) signature data is considered for experimental purposes, the proposed systems first calculate an appropriate bounding box for the original image and crop said image accordingly. Once the relevant feature vectors have been extracted, the zero-valued entries are decimated from both ends of each feature vector. The dimension of each of the truncated feature vectors is subsequently adjusted to a specified value through nearest-neighbour interpolation. The resulting feature vectors that are associated with a specific feature set are subsequently placed into the columns of a matrix. Each entry within said matrix is finally divided by the standard deviation across all of its entries.

3.7 Concluding remarks

In this chapter we have demonstrated how textural information may be extracted from an image by appropriately modifying the SDRT and applying it to various MPTs. We also showed how the SDRT may be generalised for the purpose of generic curved line detection through the construction of appropriate MPTs. In the next chapter we explain how the corresponding feature sets extracted from different images may be appropriately matched.

Chapter 4

Feature matching, threshold selection and classifier ensembles

4.1 Introduction

In Chapter 3 we explained how four groups of feature sets may be extracted from each sample (image). The aforementioned four groups of feature sets are conceptualised in Figure 4.1. The average distance between each feature set extracted from a questioned sample and the corresponding feature sets extracted from genuine samples of the claimed class may subsequently be estimated. The proposed protocol for estimating the aforementioned *raw* distance associated with a specific feature set is discussed in Section 4.2.

In order to ensure that distances associated with samples from different classes are comparable, each raw distance is normalised on a class-specific basis. A normalised dissimilarity value is therefore obtained which is subsequently converted into a normalised *score*, as explained in Section 4.3.

Since intra-class variations have now been reduced, one may proceed to employ a single global threshold to each normalised score in order to classify a questioned sample. Each questioned sample will be classified as either positive, therefore concluding that it belongs to the claimed class, or negative, therefore concluding that it does not belong to the claimed class. The threshold selection protocol is discussed in more detail in Section 4.4.

When each feature set is associated with a distinct classifier, the normalised scores for a specific feature set from a number of positive and negative questioned samples are used to gauge the overall accuracy of the associated classifier. Various classifier ensembles may be subsequently constructed by assembling the classifiers associated with two or more feature sets. Ensemble construction is discussed in more detail in Section 4.5. The classifier combination protocol for classifiers within a specific ensemble is detailed in Section 4.6.

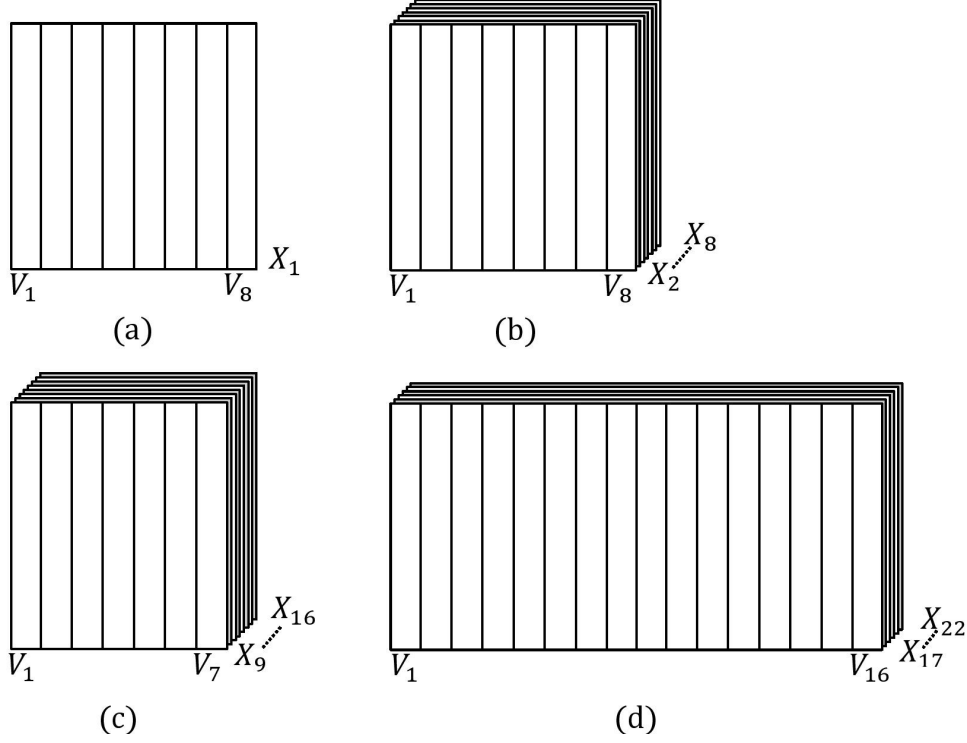


Figure 4.1: Conceptualisation of the four distinct groups of feature sets extracted by the systems proposed in this thesis. (a) A feature set that contains information extracted via the standard discrete Radon transform (SDRT). (b) A group of feature sets that are extracted by applying the modified discrete Radon transform (MDRT) to a normalised multiresolution polar transform (MPT). Each feature set is associated with a specific textural periodicity. (c) A group of feature sets that are extracted by applying the MDRT to a normalised MPT. Each feature set is associated with a specific textural orientation. (d) A group of feature sets that are extracted by applying the SDRT to a normalised MPT by only considering a projection angle of $\theta = 90^\circ$. Each of the aforementioned feature sets is geared towards the detection of specific curved lines.

4.2 Dynamic Time Warping (DTW)

Again consider Figure 4.1 that conceptualises the relevant four groups of feature sets and their dimensions. Each feature vector constitutes a column of a matrix. Note that the first group (see Figure 4.1 (a)) consists of a single feature set (X_1) that contains eight feature vectors (V_1 to V_8). The second group (see Figure 4.1 (b)) consists of seven feature sets (X_2 to X_8), where each feature set contains eight feature vectors (V_1 to V_8). The third group (see Figure 4.1 (c)) consists of eight feature sets (X_9 to X_{16}), where each feature set contains seven feature vectors (V_1 to V_7). The final group (see Figure 4.1 (d)) contains six feature sets (X_{17} to X_{22}), where each feature set is comprised of sixteen feature vectors (V_1 to V_{16}).

The distance (based on a specific feature set X) between a questioned sample I_q and a positive reference sample I_ω that is known to belong to the claimed class ω , may now be estimated. In order to achieve this, the average distance between the corresponding feature vectors extracted from these samples has to be obtained. This enables us to determine whether a questioned sample I_q should be classified as positive or negative. In Figures 4.2, 4.3, 4.4 and 4.5 examples of the aforementioned feature vectors are presented. Note that the raw DTW-based distances between the feature vectors depicted in these figures may differ substantially between *different groups* of classifiers, but within the context of a *specific* group, the distances are comparable. This is attributed to the fact that the features have already been appropriately normalised as explained in Section 3.6. Feature-specific normalisation within the context of a specific group of classifiers is therefore not required.

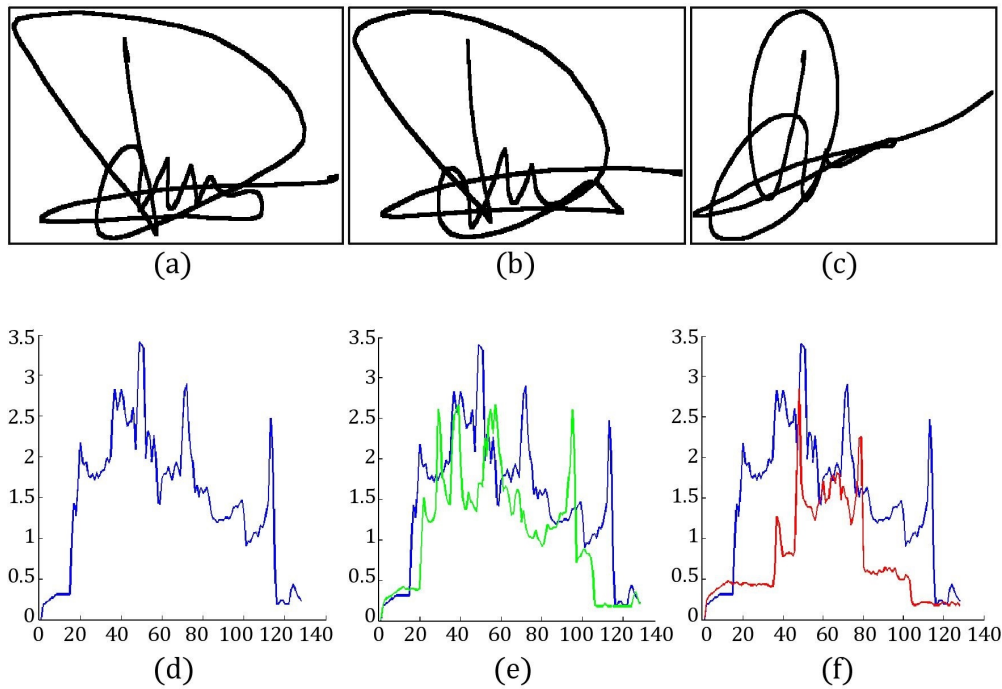


Figure 4.2: SDRT-based feature vectors associated with a projection angle of $\theta = 22.5^\circ$. (a) An example of a genuine reference sample I_ω . (b) An example of a positive questioned sample I_q^+ . (c) An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. (d) The feature vector associated with (a). (e) The feature vector associated with (b) is depicted in green, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 36.68. (f) The feature vector associated with (c) is depicted in red, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 63.09.

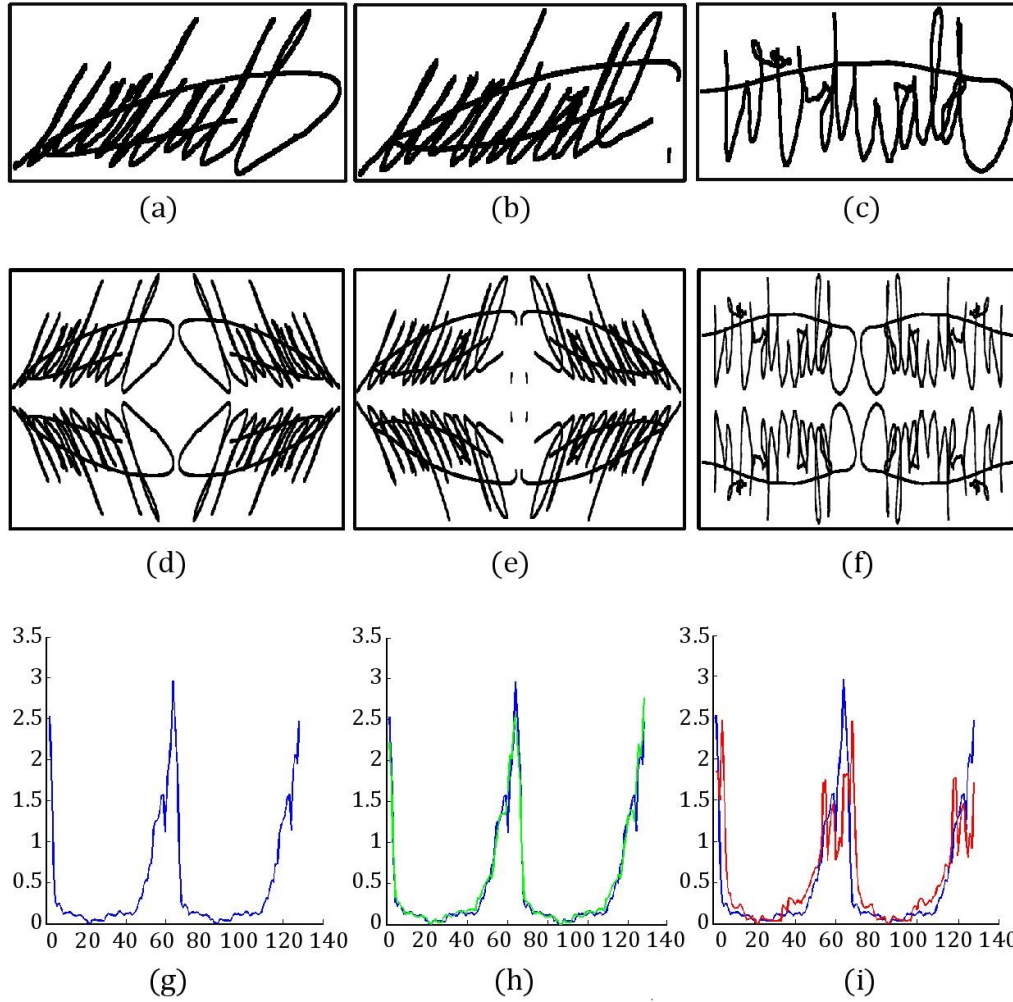


Figure 4.3: MDRT-based feature vectors associated with a projection angle of $\theta = 33.75^\circ$. These feature vectors are geared towards the detection of textural information based on specific *periodicities*. The origin of the MPT is located at the bottom right corner of the original image. **(a)** An example of a genuine reference sample I_ω . **(b)** An example of a positive questioned sample I_q^+ . **(c)** An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. **(d)** The composite image associated with (a). **(e)** The composite image associated with (b). **(f)** The composite image associated with (c). **(g)** The feature vector associated with (d). **(h)** The feature vector associated with (e) is depicted in green, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 6.58. **(i)** The feature vector associated with (f) is depicted in red, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 15.11.

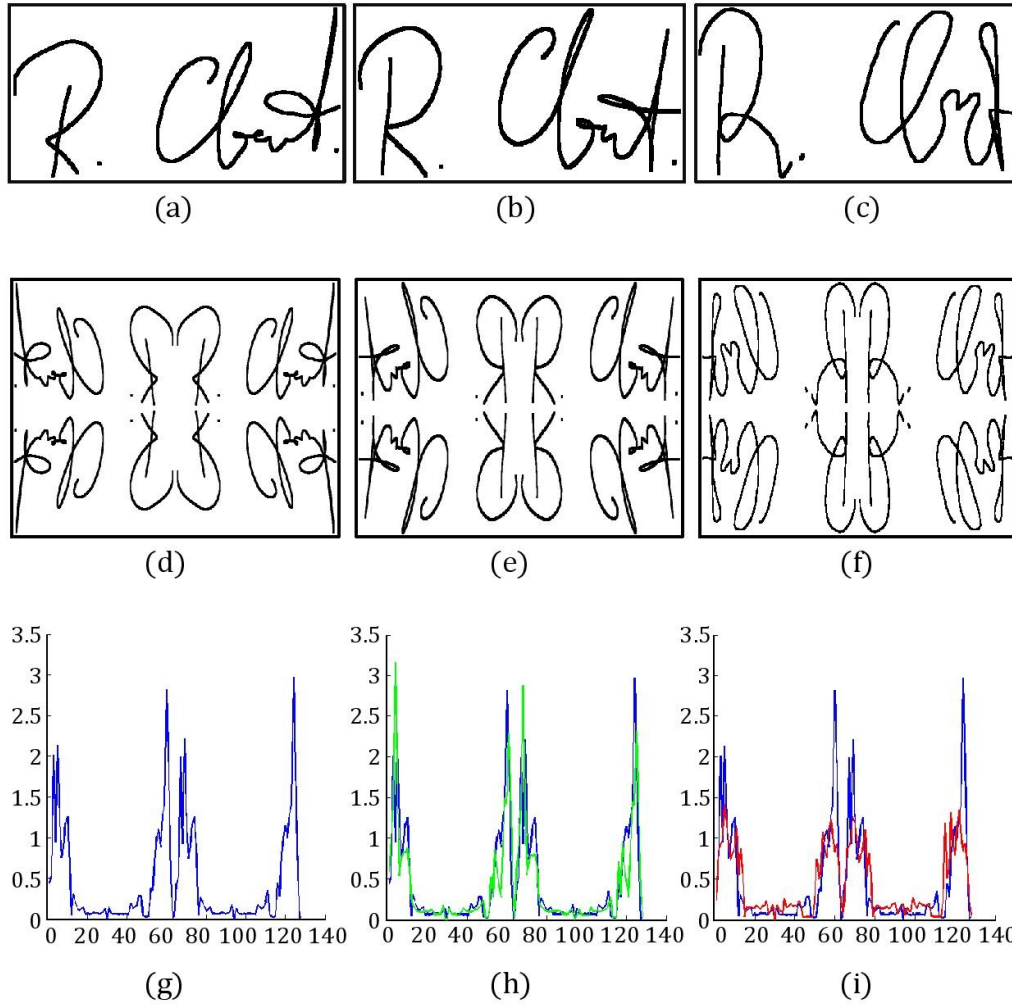


Figure 4.4: MDRT-based feature vectors associated with a projection angle of $\theta = 11.25^\circ$. These feature vectors are geared towards the detection of textural information based on specific *orientations*. The origin of the MPT is located at the bottom left corner of the original image. (a) An example of a genuine reference sample I_ω . (b) An example of a positive questioned sample I_q^+ . (c) An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. (d) The composite image associated with (a). (e) The composite image associated with (b). (f) The composite image associated with (c). (g) The feature vector associated with (d). (h) The feature vector associated with (e) is depicted in green, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 16.45. (i) The feature vector associated with (f) is depicted in red, while the feature vector associated with (d) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 19.53.

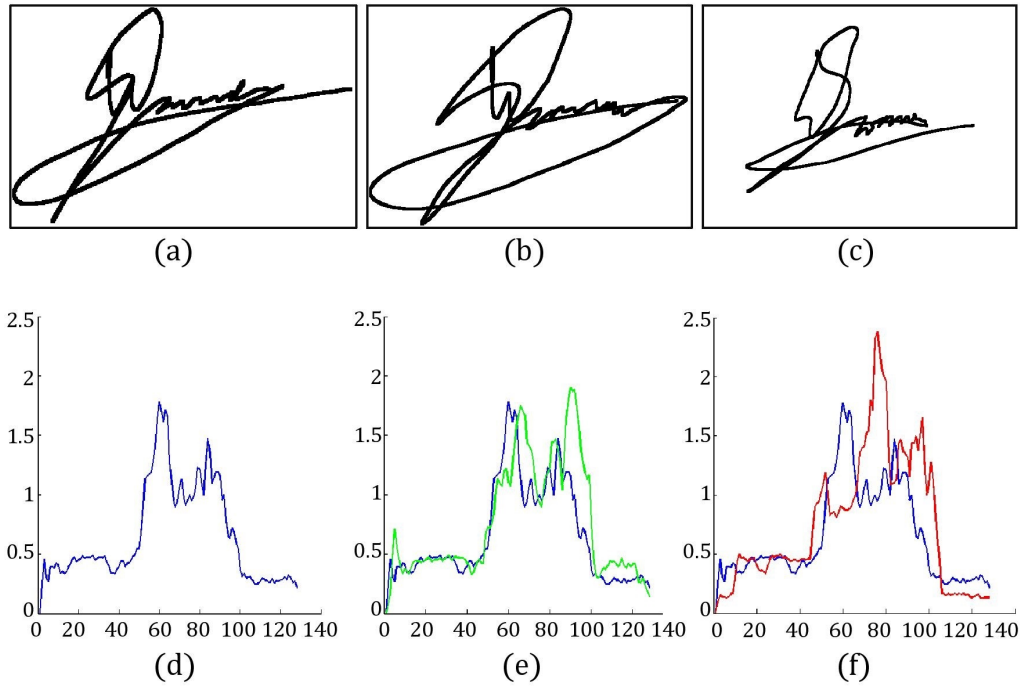


Figure 4.5: Feature vectors extracted by applying the SDRT to an MPT from a projection angle of $\theta = 90^\circ$. These feature vectors are geared towards the detection of curved lines. The origin of said MPT is located on the line that passes through the centre of the original image and its bottom right corner. The distance between the origin of the MPT and said corner of the image is given by $d = 210$ pixels. (a) An example of a genuine reference sample I_ω . (b) An example of a positive questioned sample I_q^+ . (c) An example of a negative questioned sample I_q^- , which constitutes a skilled forgery. (d) The feature vector associated with (a). (e) The feature vector associated with (b) is depicted in green, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in green and blue is 10.26. (f) The feature vector associated with (c) is depicted in red, while the feature vector associated with (a) is reproduced (in blue) for comparison purposes. The raw DTW-based distance between the respective feature vectors depicted in red and blue is 20.29.

The most straightforward and efficient method for estimating the distance between two feature sets may involve the calculation of the average Euclidean distance between the corresponding feature vectors. However, two distinct feature vectors rarely are appropriately aligned. The Euclidean distance between two unaligned feature vectors may therefore not constitute an accurate indication of the similarity between them.

We rather opt to employ a more sophisticated distance measure which is based on dynamic time warping (DTW). The proposed DTW-based distance measure non-linearly aligns the two feature vectors in question, based

on prominent features like peaks and valleys (see Figure 4.6). The Euclidean distance between the *aligned* feature vectors is subsequently calculated. This protocol therefore results in a more realistic estimate of the true distance between I_q and I_ω . DTW-algorithms have traditionally been employed within the context of human speech recognition (see Deller *et al.* (1999)). The reader is also referred to Keogh and Pazzani (2001), Henniger and Muller (2007), and Jayadevan *et al.* (2009) for detailed discussions on a number of DTW-algorithms.

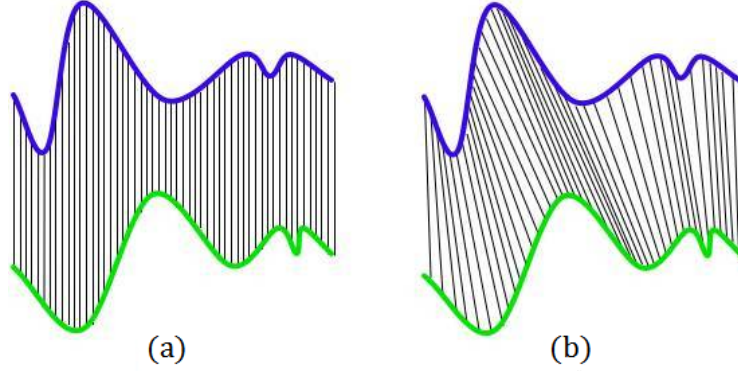


Figure 4.6: Conceptualisation of the difference between (a) the Euclidean distance and (b) a DTW-based distance (Swanepoel, 2015).

In order to ensure that the above-mentioned alignment of two feature vectors is not based on prominent features that are too far apart, the bandwidth associated with the DTW-algorithm is restricted to approximately 8 percent of the feature vector dimension.

It is assumed that N genuine reference samples (images) are available for each class ω , that is $I_\omega = \{I_\omega^{(1)}, \dots, I_\omega^{(N)}\}$. When a *specific* feature set X is extracted from each sample (image) in I_ω , the resulting feature sets are referred to as $\mathbf{X}_\omega = \{X_\omega^{(1)}, \dots, X_\omega^{(N)}\}$.

The distance metric $D(X^{(a)}, X^{(b)})$ denotes the average DTW-based distance between the corresponding feature vectors contained within feature sets $X^{(a)}$ and $X^{(b)}$.

The raw distance $D(X_q, \mathbf{X}_\omega)$ between a questioned feature set X_q and the N feature sets that constitute \mathbf{X}_ω and are known to belong to the claimed class ω , is subsequently determined. Said raw distance is therefore defined as the average DTW-based distance between the feature vectors that constitute X_q and the corresponding feature vectors that constitute \mathbf{X}_ω , that is

$$D(X_q, \mathbf{X}_\omega) = \frac{1}{N} \sum_{i=1}^N D(X_q, X_\omega^{(i)}). \quad (4.1)$$

4.3 Score normalisation

Once the raw distance, based on a specific feature set X , between a questioned image and N authentic reference samples has been obtained, a score is assigned to the questioned image. It is however essential that the raw distance, $D(X_q, \mathbf{X}_\omega)$, is first normalised on a class-specific basis in order to ensure that intra-class variation is minimised. The aforementioned normalisation is achieved as follows,

$$D_{\text{norm}}(X_q, \mathbf{X}_\omega) = \frac{D(X_q, \mathbf{X}_\omega) - \mu_\omega}{\mu_\omega}, \quad (4.2)$$

where

$$\mu_\omega = \frac{1}{\bar{N}} \sum_{i=1}^N \sum_{j>i}^N D(X_\omega^{(i)}, X_\omega^{(j)}) \quad (4.3)$$

and

$$\bar{N} = \frac{N(N-1)}{2}. \quad (4.4)$$

This normalisation strategy is similar to the well-known z -score, in which case the denominator in Equation 4.2 constitutes the standard deviation σ_ω of the raw distances between reference samples from class ω . Since N is typically small, we opt to rather estimate intra-class variation by considering the *average* raw distance μ_ω between reference samples from class ω . The reader is referred to Jain *et al.* (2005) for a more detailed discussion on the z -score and other score normalisation techniques.

The normalised distance $D_{\text{norm}}(X_q, \mathbf{X}_\omega)$ is subsequently converted into a normalised score $S_{\text{norm}}(X_q, \mathbf{X}_\omega)$ as follows,

$$S_{\text{norm}}(X_q, \mathbf{X}_\omega) = -D_{\text{norm}}(X_q, \mathbf{X}_\omega). \quad (4.5)$$

4.4 Threshold selection

Once a normalised score, based on feature set X , has been obtained for a questioned sample I_q , we aim to establish whether said sample belongs to the claimed class ω . For this purpose a global threshold τ is applied to the aforementioned normalised score. Said threshold is referred to as a *global* threshold since the *same* threshold is applied to all questioned samples irrespective of the claimed class ω . When $S_{\text{norm}}(X_q, \mathbf{X}_\omega) \geq \tau$, I_q is deemed to be a positive sample based on feature set X , otherwise it is deemed to be a negative sample.

An optimisation set is typically employed in order to obtain an estimate for the global threshold τ that results in an equal error rate (EER). Recall that the EER is achieved when the false acceptance rate (FAR) equals the false rejection rate (FRR). Within the context of offline signature verification, the optimisation set is typically comprised of a number of guinea-pig writers for

which skilled forgeries have been acquired in a controlled environment. The above-mentioned threshold is then used to authenticate questioned samples in an evaluation set. Within the context of offline signature verification, the evaluation set is typically comprised of *different* writers from those in the optimisation set and are representative of the general public. A classifier based on a specific feature set X may be evaluated by conducting a number of trials, as explained in Section 5.4. For each trial distinct sets of writers are employed for optimisation and evaluation purposes respectively. The proficiency of a specific classifier based on a specific trial may be subsequently quantified by its average error rate (AER), that is the average of its FAR and FRR. Recall that the global threshold τ for a specific trial is applied to the evaluation set. The overall proficiency of the above-mentioned classifier is finally quantified by its mean AER, that is μ_{AER} , across all trials.

4.5 Classifier ensemble construction

When each feature set X is associated with a unique classifier C , two or more classifiers may be assembled, thereby constructing a classifier ensemble. Four classifier ensembles are considered in this study.

The first ensemble E_{I} consists of the single base classifier C_1 associated with feature set X_1 . This ensemble constitutes the current benchmark system that is geared towards the detection of straight lines in the original image and utilises feature vectors based on projection profiles associated with the SDRT. Ensemble E_{I} is therefore defined as follows, $E_{\text{I}} = \{C_1\}$.

The second ensemble E_{II} consists of the seven base classifiers C_2 to C_8 associated with feature sets X_2 to X_8 . Each of these base classifiers is geared towards the detection of textural patterns with a distinct periodicity. Ensemble E_{II} is therefore defined as follows, $E_{\text{II}} = \{C_2, C_3, C_4, C_5, C_6, C_7, C_8\}$.

The third ensemble E_{III} consists of the eight base classifiers C_9 to C_{16} associated with feature sets X_9 to X_{16} . Each of these base classifiers is geared towards the detection of textural patterns with a distinct orientation. Ensemble E_{III} is therefore defined as follows, $E_{\text{III}} = \{C_9, C_{10}, C_{11}, C_{12}, C_{13}, C_{14}, C_{15}, C_{16}\}$.

The fourth and final ensemble E_{IV} consists of the seven base classifiers, $C_1, C_{17}, C_{18}, C_{19}, C_{20}, C_{21}$ and C_{22} , associated with feature sets $X_1, X_{17}, X_{18}, X_{19}, X_{20}, X_{21}$ and X_{22} . Each of these base classifiers is geared towards the detection of concentric curved lines that are associated with a distinct range of curvatures. Note that classifier C_1 , that is only geared towards the detection of straight lines in the original image, is included within E_{IV} . Ensemble E_{IV} , which may therefore be considered a generalisation of classifier C_1 , is defined as follows, $E_{\text{IV}} = \{C_1, C_{17}, C_{18}, C_{19}, C_{20}, C_{21}, C_{22}\}$.

4.6 Classifier combination

Since four classifier ensembles E_I , E_{II} , E_{III} and E_{IV} are now available, the base classifiers within each of these ensembles may subsequently be combined in order to gauge overall system proficiency. We opt to combine the classifiers within E_{II} , E_{III} and E_{IV} through both score averaging and majority voting. When the base classifiers within ensembles E_{II} , E_{III} and E_{IV} are combined through score averaging, the resulting combined classifiers are referred to as $C_{E_{II}}^{(s)}$, $C_{E_{III}}^{(s)}$ and $C_{E_{IV}}^{(s)}$, respectively. Similarly, when the base classifiers within ensembles E_{II} , E_{III} and E_{IV} are combined through majority voting, the resulting combined classifiers are referred to as $C_{E_{II}}^{(m)}$, $C_{E_{III}}^{(m)}$ and $C_{E_{IV}}^{(m)}$.

Score averaging. In the case of score averaging, the average score across all base classifiers within a specific classifier ensemble E is first obtained, after which a global threshold τ is applied to the average score. Each base classifier within E is therefore assigned an equal weight. For a specific trial, a global threshold τ is selected, which results in an EER when all the scores in the relevant optimisation set are considered. Said threshold is subsequently applied to all the scores in the relevant evaluation set. The mean μ_{AER} and standard deviation σ_{AER} of the AERs across all of the individual trials are subsequently reported.

Majority voting. In the case of majority voting, different global thresholds τ are applied to the scores associated with different base classifiers within a specific classifier ensemble E . For a specific trial, a global threshold is selected for a specific base classifier that results in an EER when all the scores in the relevant optimisation set are considered. Said threshold is subsequently applied to all the scores associated with the aforementioned classifier in the relevant evaluation set. A questioned sample is accepted when the majority of base classifiers within a specific ensemble decides to accept it, otherwise it is rejected. The mean μ_{AER} and standard deviation σ_{AER} of the AERs across all of the individual trials are subsequently reported.

When evaluated on Dolfig's data set, that contains authentic offline signatures and skilled forgeries from 51 writers, the statistics μ_{AER} and σ_{AER} estimated for the combined classifiers $C_{E_{II}}^{(s)}$, $C_{E_{III}}^{(s)}$, $C_{E_{IV}}^{(s)}$, $C_{E_{II}}^{(m)}$, $C_{E_{III}}^{(m)}$ and $C_{E_{IV}}^{(m)}$ may be compared to those for the benchmark system $E_I = \{C_1\}$. In this way the improvement in system proficiency (and the statistical significance of said improvement) that results from the research presented in this thesis may be quantified (see Section 5.5).

4.7 Concluding remarks

In this chapter it was demonstrated how normalised scores may be obtained from feature sets, which are extracted according to the protocol outlined in Chapter 3. Additionally, we elaborated on how said feature sets may be used

to construct classifier ensembles. The base classifiers within a specific classifier ensemble may subsequently be combined through score averaging or majority voting. In the following chapter a detailed description of the data set under consideration and the experimental protocol is provided, as well as an in-depth analysis of the experimental results and the contribution of this research.

Chapter 5

Experiments

5.1 Introduction

The *generic* systems introduced in Chapters 3 and 4 utilise novel strategies for the extraction of feature vectors geared towards the detection of textural information and curved lines within an image. In order to gauge the overall proficiency of the proposed generic systems, we implement said systems on an offline signature database. We discuss the database that is considered for experimental purposes in Section 5.2. In Section 5.3 we discuss the system parameters employed for the purpose of experimentation and in Section 5.4 we elaborate on the experimental protocol. Finally, in Section 5.5 the experimental results are reported and analysed.

5.2 Data

For the purpose of experimentation, we utilise signature data captured by Hans Dolfing for his PhD thesis (Dolfing, 1998). This data was originally captured online, but has subsequently been converted into ideal offline signatures with a fixed stroke width and zero noise (Coetzer *et al.*, 2004). The data set, that is referred to as Dolfing’s data set, is comprised of signatures associated with 51 individual writers. Said data set contains 15 authentic training samples and 15 authentic test samples for each writer. In addition to this, it also contains 30 over-the-shoulder and 30 home-improved skilled forgeries associated with each writer, with the exception of two writers for whom only 30 home-improved forgeries are available per writer.

For the purpose of this thesis, only five of the 15 authentic training samples per writer are employed for reference purposes. This ensures that the reported results are realistic in terms of what may be achieved in practice.

An over-the-shoulder forgery is produced when the forger witnesses the production of the actual signature and then imitates it, whereas a home-improved forgery is produced when the forger is in possession of a sample of the actual

signature and has sufficient time to practice imitating it. The offline versions of the over-the-shoulder and home-improved forgeries are collectively referred to as skilled forgeries within the context of this thesis. For the purpose of this thesis, only skilled forgeries are considered for experimental purposes. Random forgeries, which constitute imitations of a certain writer’s signature without any prior knowledge of said signature, are not considered.

5.3 System parameters

A set of predefined parameters is required for experimental purposes. Since these parameters influence the results, they are determined in such a way that the best possible results are obtained while still rendering the systems as efficient as possible.

A single parameter is predefined within the context of the MPT, namely the number of evenly distributed origins, that is N_{MPT} . A set of $N_{\text{MPT}} = 8$ origins is specified for the purpose of textural pattern detection (Section 3.4.4), while a set of $N_{\text{MPT}} = 16$ origins is specified for the purpose of curved line detection (Section 3.5).

Within the context of the SDRT, two parameters are predefined. The first relates to the number of projection angles, which is defined as $N_\theta = 8$. A set of eight equally distributed projections angles are therefore specified as explained in Section 3.2 so that $\theta \in [0^\circ, 180^\circ)$. An exception to this involves the special case where the SDRT is applied to the MPT from a single projection angle of $\theta = 90^\circ$ for the purpose of detecting curved lines. The second parameter relates to the feature vector dimension which is fixed at $V_{\text{dim}} = 128$. It was determined in Swanepoel (2015) that said parameter value is optimal in the sense that a larger value does not lead to a sufficient increase in system proficiency so as to justify the accompanying loss in efficiency.

The predefined parameters within the context of the MDRT relate to the set of projection angles, the feature vector dimension and the number of duplicates to be appended to the original MPT. The number of projection angles is defined as $N_\theta = 7$, but within the context of the MDRT, the equally distributed projection angles (as explained in Section 3.4.1) are restricted to the interval $\theta \in (0^\circ, 90^\circ)$. The feature vector dimension is again defined as $V_{\text{dim}} = 128$, and the number of duplicates to be appended is fixed at 3, as explained in Section 3.4.3.

For the scenario where the SDRT is applied to the MPT from a single projection angle of $\theta = 90^\circ$ for the purpose of curved line detection, only one additional parameter is predefined. This is the set of distances d from the edge of the original image and the origin of the relevant MPT, as discussed in Section 3.5. The set of distances is defined as $d = \{0 \text{ pixels}, 70 \text{ pixels}, 140 \text{ pixels}, 210 \text{ pixels}, 350 \text{ pixels}\}$.

5.4 Protocol

Dolfing’s data set is considered independently for each classifier, where each classifier coincides with a unique feature set, as explained in Sections 4.1 and 4.2.

A *global* threshold is applied to all questioned samples in order to classify them as either positive or negative. Questioned samples with a score larger than or equal to the global threshold are classified as positive, otherwise they are classified as negative. In order to find a realistic global threshold, the experimental data is separated into an optimisation set and an evaluation set. Said global threshold is determined by considering the optimisation set and selecting the threshold that results in an equal error rate (EER) (see Section 4.4). This threshold is subsequently applied to the evaluation set.

In order to ensure thorough experimentation and objective results, the proposed systems are implemented on Dolfing’s data set by conducting 30 independent trials. These trials are conducted as follows,

- The 51 writers in Dolfing’s data set are randomly ordered 10 different times, resulting in 10 different randomisations of the order of the writers.
- For each one of the aforementioned randomisations, the set of 51 writers is divided into three subsets of 17 writers each. A unique set of writers therefore constitutes a specific subset for each randomisation. Each one of these subsets is considered in turn as the evaluation set with the remaining two sets constituting the optimisation set.

The mean average error rate across all 30 trials, that is μ_{AER} , as defined in Section 4.4, is subsequently reported for each individual classifier C_1 to C_{22} , as well as for the combined classifiers $C_{E\text{II}}^{(s)}$, $C_{E\text{III}}^{(s)}$, $C_{E\text{IV}}^{(s)}$, $C_{E\text{II}}^{(m)}$, $C_{E\text{III}}^{(m)}$ and $C_{E\text{IV}}^{(m)}$.

5.5 Results

The achieved results are categorised into writer-specific results (Section 5.5.1) and results across all writers (Section 5.5.2). The writer-specific results indicate the number of writers for which a specific classifier or set of classifiers is the most proficient. The result across all writers enables one to gauge the overall proficiency of each individual classifier, as well as the proficiency of the combined classifiers. The aforementioned results may also be compared to those for similar existing systems in order to place the proficiency of the proposed systems into perspective. A statistical significance test is finally conducted in order to determine whether the presented results (for the classifiers proposed in this thesis) constitute a statistically significant improvement on the existing benchmark system. Note that each of the ensembles, as defined in Section 4.5, is associated with an independent system. The reported results are therefore achieved by considering each ensemble independently.

5.5.1 Writer-specific results

It is important to note that the results presented in this section are *unrealistic* within the context of a real-life scenario. For a specific classifier C , the optimal *local* threshold that results in an EER is selected for each classifier. The proficiency of said classifier in authenticating signatures associated with a *specific* writer is subsequently quantified by said EER. The protocol outlined in Section 5.4, that involves the application of a *global* threshold to normalised scores and the partitioning of the data into an optimisation and evaluation set, is therefore *not* adhered to. The sole purpose of this analysis is to demonstrate that the base classifiers associated with ensembles other than the benchmark $E_I = \{C_1\}$ may be better suited than E_I for authenticating signatures associated with *specific* writers. These results are presented in Figures 5.1, 5.2 and 5.3.

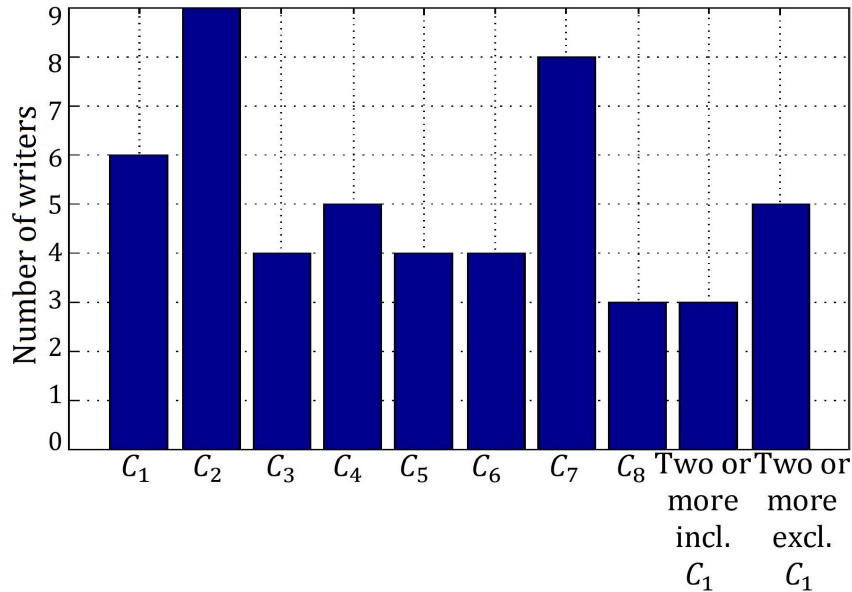


Figure 5.1: Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{II} = \{C_2, C_3, C_4, C_5, C_6, C_7, C_8\}$. The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar.

From Figure 5.1 it is clear that ensemble $E_I = \{C_1\}$ is the most proficient for 6 writers, and jointly the most proficient for another 3 writers. This however indicates that for each of the remaining 42 writers (84% of the total number

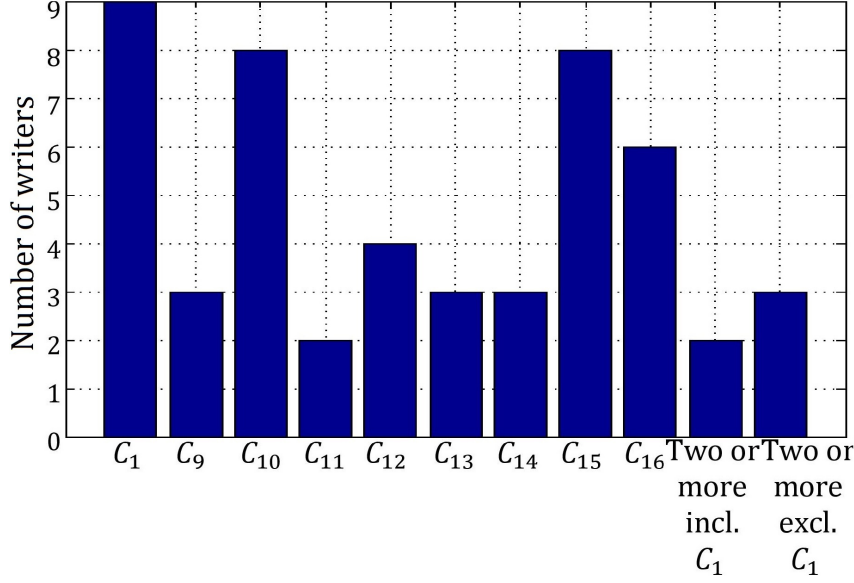


Figure 5.2: Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{III} = \{C_9, C_{10}, C_{11}, C_{12}, C_{13}, C_{14}, C_{15}, C_{16}\}$. The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar.

of writers) one or more of the base classifiers within ensemble E_{II} is the most proficient.

From Figure 5.2 it is clear that ensemble $E_I = \{C_1\}$ is the most proficient for 9 writers, and jointly the most proficient for another 2 writers. This however indicates that for each of the remaining 40 writers (80% of the total number of writers) one or more of the base classifiers within ensemble E_{III} is the most proficient.

From Figure 5.3 it is clear that ensemble $E_I = \{C_1\}$ is the most proficient for 17 writers, and jointly the most proficient for another 2 writers. This however indicates that for each of the remaining 32 writers (64% of the total number of writers) one or more of the other base classifiers within ensemble E_{IV} is the most proficient.

It may therefore be concluded that there is no scenario for which the SDRT-based benchmark system $E_I = \{C_1\}$ is the most proficient classifier for the majority of writers. As mentioned before, these results are not achievable in practice, but indicate that the implementation of the proposed systems may lead to an improvement in overall proficiency.

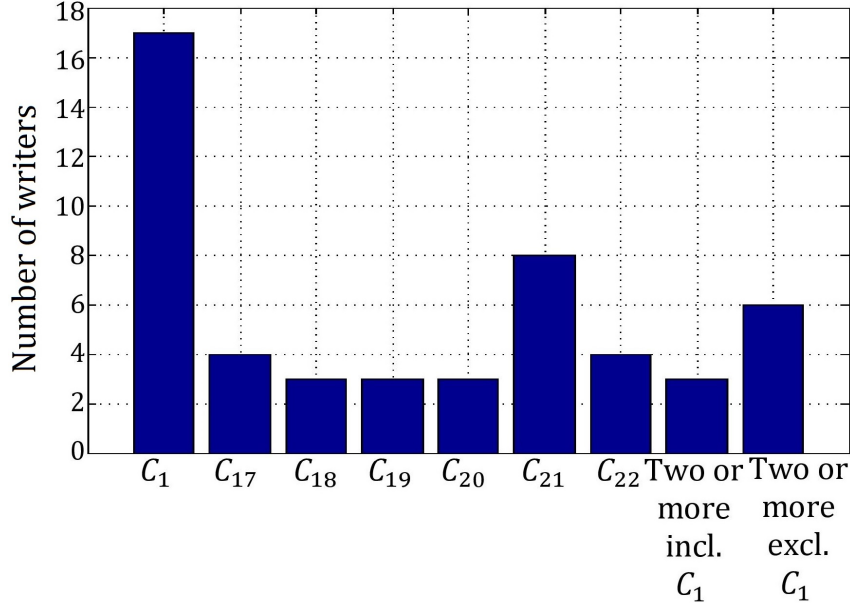


Figure 5.3: Writer-specific results for ensemble $E_I = \{C_1\}$ when compared to those in ensemble $E_{IV} = \{C_1, C_{17}, C_{18}, C_{19}, C_{20}, C_{21}, C_{22}\}$. Recall that classifier C_1 is included within ensemble E_{IV} . The height of each bar represents the number of individual writers for which a specific classifier is the most proficient. The number of writers for which two or more classifiers, that include C_1 , are jointly the most proficient, is represented by the height of the penultimate bar. The number of writers for which two or more classifiers, that do not include C_1 , are jointly the most proficient is represented by the height of the last bar.

5.5.2 Results across all writers

In this section *realistic* results are presented within the context of a real-life scenario across all writers. The protocol outlined in Section 5.4 is therefore strictly adhered to. Thirty independent trials are conducted by considering Dolfig’s data set. The scores are normalised according to the protocol defined in Section 4.3. During each trial, the data is partitioned into an optimisation and an evaluation set. A *global* threshold is selected that results in an EER across *all* of the writers in the optimisation set. Said threshold is subsequently applied to *all* of the normalised scores associated with *different* writers in the evaluation set. The mean AER, that is μ_{AER} , of the evaluation sets across all trials are reported. The results are reported for the *individual* classifiers as well as for the combined classifiers (using score averaging or majority voting) within the respective ensembles (see Tables 5.1, 5.2, 5.3 and 5.4).

The result reported for C_1 in Table 5.1, that is $\mu_{\text{AER}} = 9.51\%$, constitutes the proficiency of the SDRT-based benchmark system. We therefore aim to improve upon the aforementioned result.

The results reported in Table 5.2 are based on the detection of textural

Table 5.1: Result for ensemble E_I .

Classifier	C_1
$\mu_{\text{AER}}(\%)$	9.51

Table 5.2: Results for ensemble E_{II} .

Classifier	C_2	C_3	C_4	C_5	C_6	C_7	C_8	$C_{E_{II}}^{(s)}$	$C_{E_{II}}^{(m)}$
$\mu_{\text{AER}}(\%)$	10.40	10.10	8.68	8.75	10.22	9.61	10.57	8.22	6.10

patterns with specific *periodicities*. Note that the combined classifiers $C_{E_{II}}^{(s)}$ and $C_{E_{II}}^{(m)}$, with mean AERs of 8.22% and 6.10% respectively, are both more proficient than C_1 , for which the mean AER is 9.51%. Also, note that $C_{E_{II}}^{(m)}$ outperforms $C_{E_{II}}^{(s)}$, which may be attributed to the fact that score averaging may not be the optimal strategy for combining the individual base classifiers within an ensemble. All classifiers are awarded equal weights despite the fact that certain classifiers are clearly superior to others. In the case of majority voting the input of less proficient classifiers may be ignored if the final decision reflects the opinion of the more proficient classifiers.

Table 5.3: Results for ensemble E_{III} .

Classifier	C_9	C_{10}	C_{11}	C_{12}	C_{13}	C_{14}	C_{15}	C_{16}	$C_{E_{III}}^{(s)}$	$C_{E_{III}}^{(m)}$
$\mu_{\text{AER}}(\%)$	14.70	11.65	12.41	12.87	13.09	11.24	11.31	10.64	8.39	6.82

The results reported in Table 5.3 are based on the detection of textural patterns with specific *orientations*. Note that the combined classifiers $C_{E_{III}}^{(s)}$ and $C_{E_{III}}^{(m)}$, with mean AERs of 8.39% and 6.82% respectively, are both more proficient than C_1 , for which the mean AER is 9.51%. Also, note that $C_{E_{III}}^{(m)}$ outperforms $C_{E_{III}}^{(s)}$, which may be attributed to the same rationale as the one outlined within the context of Table 5.2.

Table 5.4: Results for ensemble E_{IV} .

Classifier	C_1	C_{17}	C_{18}	C_{19}	C_{20}	C_{21}	C_{22}	$C_{E_{IV}}^{(s)}$	$C_{E_{IV}}^{(m)}$
$\mu_{\text{AER}}(\%)$	9.51	12.27	11.39	11.03	10.80	11.05	10.63	10.14	8.40

The results reported in Table 5.4, are based on the detection of concentric curved lines. Note that the combined classifier $C_{E_{IV}}^{(s)}$, for which the mean AER is 10.14%, does not outperform C_1 . The combined classifier $C_{E_{IV}}^{(m)}$, for which the mean AER is 8.40%, does however outperform C_1 , for which the mean AER is 9.51%. The combined classifier $C_{E_{IV}}^{(m)}$ that employs majority voting again outperforms the combined classifier $C_{E_{IV}}^{(s)}$ that employs score averaging.

Furthermore, it is clear from Tables 5.1, 5.2, 5.3 and 5.4 that the combined classifiers associated with ensemble E_{II} are the most proficient within the context of either score averaging *or* majority voting. It may therefore be concluded that those features geared towards the detection of *textural* information based on specific textural *periodicities* are the most discriminative within the context of this thesis.

5.5.3 Comparison to existing research

In order to gain some perspective on how the results reported in Section 5.5.2 compare to that of existing systems, the reader is referred to Table 5.5. Recall that N represents the number of authentic reference signatures either utilised as templates (as is the case for the systems proposed in this thesis) or employed for the purpose of training statistical models (Coetzer *et al.* (2004), Panton and Coetzer (2010), and Swanepoel (2015)). For a detailed discussion on the training of statistical models, the reader is referred to Cohn *et al.* (1996). Recall that N_θ represents the number of projection angles, while N_{MPT} represents the number of origins specified for the MPT and V_{dim} represents the feature vector dimension.

Table 5.5: The mean AER for systems evaluated on Dolfing’s data set. The best results for the systems proposed in this study are presented in boldface.

System	N	N_θ	N_{MPT}	V_{dim}	$\mu_{AER}(\%)$
Coetzer <i>et al.</i> (2004)	15	128	-	512	12.20
Panton and Coetzer (2010)	15	128	-	512	8.89
Swanepoel (2015) (QDS)	5	8	-	128	5.55
Swanepoel (2015) (SVMS)	5	8	-	128	6.06
Patterns based on textural periodicity (this study)	5	7	8	128	6.10
Patterns based on textural orientation (this study)	5	7	8	128	6.82
Concentric curved lines (this study)	5	1	16	128	8.40

Note that both of the systems proposed by Swanepoel (2015) are *writer-*

independent, while the remaining systems, including those proposed in this thesis, are *writer-dependent*. Writer-independent systems have the advantage that skilled forgeries may also be used for model training by considering so-called guinea-pig writers (see Section 2.2.2). Within the context of writer-dependent systems, the guinea-pig writers (those within the optimisation set) are merely used to select an appropriate threshold. The more sophisticated and elegant score fusion strategies employed by Swanepoel (2015) may further contribute to the proficiency of the relevant systems. We intend to investigate similar score fusion strategies for future work (see Section 6.2.2). It is evident from Table 5.5 that the systems proposed in this thesis constitute an improvement on the current state-of-the-art within the context of *writer-dependent* systems.

5.5.4 Statistical significance testing

The results reported in Sections 5.5.2 and 5.5.3 demonstrate that the systems developed in this thesis outperform the current benchmark system $E_I = \{C_1\}$. However, in order to demonstrate that said improvements in system proficiency are statistically significant, a statistical significance test is required.

In order to conduct a t -test, two independent data sets with normal distributions are required. We opt to perform a two-sample t -test at a significance level of $\alpha = 0.05$. The level of significance indicates the probability that the null hypothesis is rejected when it is in fact true. For a detailed discussion on the t -test, the reader is referred to Cressie and Whitford (1986) and Dowdy *et al.* (2011).

Sample set A is defined as the set of 30 AERs obtained for ensemble E_I when 30 independent trials are conducted on Dolfing's data set as explained in Section 5.4. Sample set B is defined as the set of 30 AERs obtained for a combined classifier $C_E^{(s)}$ or $C_E^{(m)}$, as proposed in this thesis, which is obtained through score averaging or majority voting respectively. Furthermore, $\mu_{\text{AER}}^{(A)}$ is defined as the population mean for set A, while $\mu_{\text{AER}}^{(B)}$ is defined as the population mean for set B. The t -test is employed in order to determine whether the mean AER for ensemble E_I (that is $\mu_{\text{AER}}^{(A)}$) is significantly larger the mean AER associated with $C_E^{(s)}$ or $C_E^{(m)}$ (that is $\mu_{\text{AER}}^{(B)}$). A t -test is conducted for ensembles E_{II} and E_{III} , but not for ensemble E_{IV} since E_{IV} is not independent of E_I . The hypotheses considered are therefore as follows:

- Null hypothesis $H_0 : \mu_{\text{AER}}^{(A)} - \mu_{\text{AER}}^{(B)} \leq 0$
- Alternative hypothesis $H_A : \mu_{\text{AER}}^{(A)} - \mu_{\text{AER}}^{(B)} > 0$

This therefore constitutes a right-tailed t -test, where the alternative hypothesis indicates that a specific proposed system associated with sample set B achieves a lower mean AER than the benchmark system, therefore implying a statistically significant improvement in system proficiency.

Let σ_Y^2 and n_Y respectively denote the population variance and population size for an arbitrary sample set Y . The test statistic t_{stat} is subsequently obtained as follows,

$$t_{\text{stat}} = \frac{\mu_{\text{AER}}^{(A)} - \mu_{\text{AER}}^{(B)}}{Q}, \quad (5.1)$$

where

$$Q = \sqrt{\left(\frac{(n_A - 1)\sigma_A^2 + (n_B - 1)\sigma_B^2}{n_A + n_B - 2} \right) \left(\frac{1}{n_A} + \frac{1}{n_B} \right)}. \quad (5.2)$$

Note that within the context of the t -test conducted here, $n_A = n_B = 30$. Therefore, Q may be simplified as follows,

$$Q = \sqrt{\frac{\sigma_A^2 + \sigma_B^2}{30}} \quad (5.3)$$

In order to reject the null hypothesis at a significance level of $\alpha = 0.05$ the value of t_{stat} (as obtained through the t -test) has to exceed a critical t -value, t_{crit} , which is obtained from a so-called t -table. Therefore, in order to reject the null hypothesis and deem the improvement in system proficiency to be statistically significant the following must hold,

$$\begin{aligned} t_{\text{stat}} &> t_{\text{crit}} \\ &= t_{(\alpha, n_A + n_B - 2)} \\ &= t_{(0.05, 58)} \\ &= 1.671. \end{aligned} \quad (5.4)$$

The results reported in Tables 5.6 and 5.7 indicate that when the individual base classifiers within ensembles E_{II} and E_{III} are combined through majority voting, a statistically significant improvement in system proficiency over the benchmark system E_{I} is achieved. When the individual base classifiers within ensembles E_{II} and E_{III} are combined through score averaging, the improvement in system proficiency is however not statistically significant. This reiterates the conclusion reached earlier: score averaging does not constitute a sensible classifier combination strategy within the context of the systems proposed in this thesis.

Table 5.6: Statistics for a t -test, where sample set A contains the AERs associated with ensemble E_{I} and sample set B contains the AERs associated with ensemble E_{II} .

Classifier combination strategy	t_{stat}
Score averaging ($C_{E_{\text{II}}}^{(s)}$)	0.54
Majority voting ($C_{E_{\text{II}}}^{(m)}$)	2.69

Table 5.7: Statistics for a t -test, where sample set A contains the AERs associated with ensemble E_I and sample set B contains the AERs associated with ensemble E_{III} .

Classifier combination strategy	t_{stat}
Score averaging ($C_{E_{III}}^{(s)}$)	0.37
Majority voting ($C_{E_{III}}^{(m)}$)	1.91

5.6 Concluding remarks

Firstly, in this chapter writer-specific experiments were conducted on Dolfing’s data set. The results indicated that for various individual writers one or more of the base classifiers associated with the proposed systems outperform the existing SDRT-based benchmark system. This led us to believe that the proposed systems may also outperform the existing benchmark system across all writers in a realistic scenario.

Experimental results across all writers were subsequently reported for the individual classifiers, as well as for combined classifiers that employ score averaging or majority voting. We concluded that all of the combined classifiers, except for a combined classifier associated with ensemble E_{IV} (when score averaging is employed), outperform the SDRT-based benchmark system. We furthermore concluded that classifier combination through majority voting is consistently superior to score averaging.

Finally, it was shown that the majority voting-based combined classifiers associated with ensembles E_{II} and E_{III} *significantly* outperform the existing SDRT-based benchmark system associated with ensemble E_I . The superiority of the corresponding combined classifiers obtained through score averaging is however not statistically significant.

In the following chapter, avenues for potential future research are investigated.

Chapter 6

Conclusion and future work

6.1 Conclusion

In this thesis generic systems were developed that improve upon an existing benchmark system. Said benchmark system employs features that are based on the standard discrete Radon transform (SDRT) and is geared towards the detection of straight lines within an image. The proposed systems utilise novel feature extraction strategies that employ the SDRT in an *indirect* way. The first strategy involves the utilisation of a normalised modified discrete Radon transform (MDRT), which is applied to a multiresolution polar transform (MPT) for the purpose of detecting textural patterns, based either on their periodicity or their orientation. This is achieved via the construction of suitable spirals. The second strategy constitutes a specific scenario for the SDRT, where only one projection is calculated from an angle of 90° . This simplified version of the SDRT is also applied to a normalised MPT for the purpose of detecting various concentric curved lines.

For the purpose of template matching, a dynamic time warping-based (DTW-based) distance measure is employed. This facilitates the assignment of scores to questioned samples.

The proposed systems are evaluated by estimating their ability in detecting skilled forgeries within Dolfing's data set that contains ideal handwritten signatures from 51 writers. The results indicate that for each of the proposed ensembles, the proficiency of the combined classifier that utilises majority voting is superior to that of the SDRT-based benchmark system. Said superiority is *statistically* significant. Furthermore, when the proficiency of the proposed systems are compared to those of existing *writer-dependent* systems within the context of an ideal (noiseless) data set, the proposed systems are superior.

The objectives of this research, as outlined in Section 1.3, have therefore been achieved. The value and contribution of the proposed feature extraction protocol are further highlighted by a publication in peer-reviewed conference proceedings (Fick *et al.*, 2016). There is however room for improving the sys-

tems proposed in this thesis. Possible avenues for future research are therefore discussed in the following section.

6.2 Future work

Due to time constraints and restrictions on the scope of this thesis, certain research avenues were not pursued. A discussion on viable future research is therefore provided. Strategies for obtaining rotation invariant features are put forward in Section 6.2.1, should this be required for a specific data set or application. More sophisticated fusion strategies that may potentially further improve system proficiency, are presented in Section 6.2.2. The systems developed in this thesis were evaluated within the context of skilled forgery detection in a set of questioned offline handwritten signatures. The proposed systems are however generic and may therefore be applied to various different data sets and within a number of different authentication scenarios. Alternative data sets and authentication scenarios are discussed in Sections 6.2.3 and 6.2.4, respectively. In Section 6.2.5 possible alternatives to template matching (through DTW) are presented. Possible adjustments to the system parameters are presented in Section 6.2.6. Finally, ways in which the efficiency of the proposed systems may be improved are suggested in Section 6.2.7.

6.2.1 Rotation invariant features

In order to ensure that the proposed feature sets are suitable for other data sets or authentication scenarios, it may be required that they are rendered rotation invariant. The feature sets associated with the specific scenario where the SDRT is applied to an MPT from a projection angle of 90° , i.e. feature sets X_{17} to X_{22} , are already rotation invariant, as explained in Section 3.6. In this section we discuss how to obtain rotation invariant feature sets within the context of the SDRT-based benchmark system, as well as for scenarios where the MDRT is applied to an MPT.

Within the context of the SDRT-based benchmark system, each feature vector (column within the feature matrix X_1) coincides with a specific projection profile calculated from a predefined projection angle. In order to compensate for rotational variations, the DTW-based distance between a specific reference feature set and a *shifted* version of a questioned feature set may be calculated as conceptualised in Figure 6.1. When a questioned feature set is shifted by one column (feature vector) towards the right, the first feature vector is *replaced* by the last feature vector. The aforementioned wrap-around procedure is facilitated by the periodic nature of the projection profiles when calculated from projection angles in the interval $\theta \in [0^\circ, 360^\circ)$. Note that when the projection angles are restricted to the interval $\theta \in [0^\circ, 180^\circ)$, as is the case in this

thesis, the first feature vector is replaced by the last feature vector after the latter has been flipped in an up-down fashion (see Figure 6.1).

Whether the questioned feature set is shifted towards the left or right, is determined by the direction of the rotation (clockwise or counterclockwise) that is to be compensated for. The angle of rotation is determined by the number of columns by which the questioned feature set is shifted. In Figure 6.1, for example, the questioned feature set is shifted by one column towards the right. Within the context of the protocol outlined in Section 3.1, assuming that the projection angle increases when the column index increases, this implies that the scenario depicted in Figure 6.1 compensates for a counterclockwise rotation of 22.5° . The range of possible rotations is also determined by the range of the projection angles. Should the aim be to compensate for a rotation of 10° within the range of $[0^\circ, 360^\circ)$, 36 equally distributed angles within the interval of $\theta \in [0^\circ, 360^\circ)$ have to be specified.

Although this protocol constitutes an effective strategy for constructing rotation invariant features, it may impact negatively on the overall computational efficiency of the relevant systems.

In order to facilitate the construction of rotation invariant features within the context of the MDRT being applied to an MPT, an alternative strategy needs to be followed. The various subimages within a constructed composite image (see Section 3.4.2) have to be rotated on an individual basis. This may be achieved by rotating the original image (subimage within the constructed composite image) according to the rotation that has to be compensated for. Reflections of the rotated version of the original image are subsequently appended to the rotated original image.

Although this protocol will result in rotation invariant features, it once again negatively affects the computational efficiency of the relevant systems, since a new MPT has to be calculated for each rotation.

Since the MPTs of various orientations of the original image have already been calculated, the DTW-based distance between a reference feature set and a questioned feature set associated with a different orientation may simply be calculated in order to compensate for rotational variations. This does however limit the number of rotations that may be compensated for to the predefined orientations. Computational efficiency is however still negatively impacted, albeit to a lesser extent, due to the calculation of additional DTW-based distances.

6.2.2 More sophisticated classifier fusion strategies

When the *oracle* results for the proposed systems (as presented in Table 6.1) are considered, it is clear that there is room for improvement. The oracle results refer to those obtained through determining the best classifier for each individual writer in Dolfing's data set. Said classifier is then exclusively employed for authenticating the questioned signatures claimed to belong to the

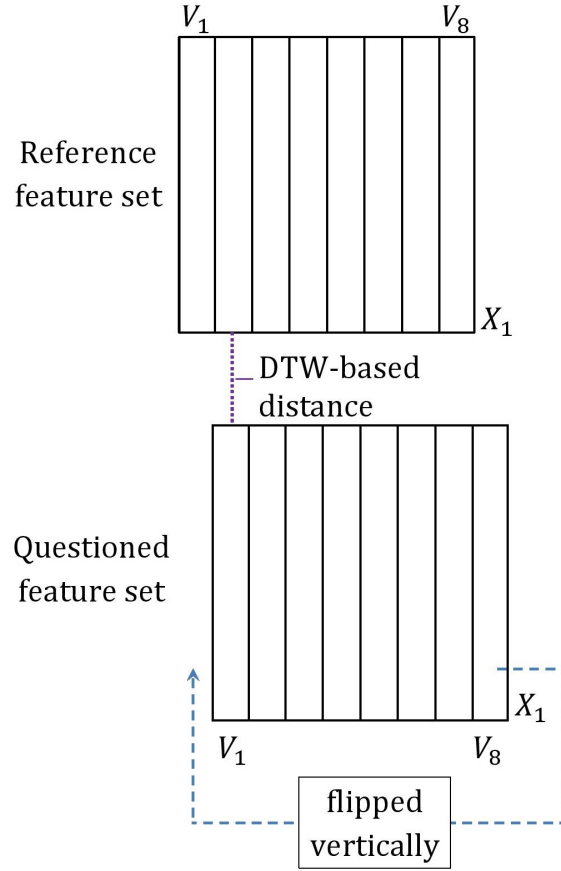


Figure 6.1: A conceptualisation of how rotational variations may be compensated for within the context of the SDRT-based benchmark system. The questioned feature set is shifted a specified number of columns towards the left or right with respect to the reference feature set. A wrap-around procedure is facilitated by the periodicity of the projection angle in the interval $\theta \in [0^\circ, 360^\circ)$.

relevant writer. These results constitute the best *obtainable* results for the proposed systems. Said oracle results are however not achievable in reality since the best classifier for each individual writer is not known. They simply hint at the possible improvements that may result from combining the proposed individual base classifiers in an optimal way. Recall that ensembles E_{II} , E_{III} and E_{IV} are defined in Section 4.5.

Table 6.1: Oracle results for ensembles E_{II} , E_{III} and E_{IV} .

Classifier ensemble	EER (%)
E_{II}	5.20
E_{III}	4.85
E_{IV}	6.65

As expected, for each of the classifier ensembles the oracle result is superior to the results reported in Section 5.5. Although these results are not realistic, they do indicate that more sophisticated fusion techniques may have a positive impact on the proficiency of the proposed systems.

Classifier fusion techniques may be categorised into feature level, score level and decision level fusion techniques. We now elaborate on score level and decision level fusion within the context of possible future improvements to the systems proposed in this thesis. Feature level fusion is however not discussed.

6.2.2.1 Score level fusion

Score level fusion techniques focus on combining the scores from various classifiers into a single score, after which the combined score is used for classification purposes. Said techniques typically involve the assignment of weights to the scores of the individual classifiers. In the book by Ross *et al.* (2006) score fusion is discussed in great detail. Score fusion is also employed and discussed by Swanepoel and Coetzer (2010) and Swanepoel (2015). Score fusion techniques can further be grouped into either static score fusion or dynamic score fusion techniques.

Static score fusion involves the assignment of a weight to the score of every individual classifier. Score averaging is an example of static score level fusion, where equal weights are assigned to the scores of the individual classifiers. As was shown in Section 5.5, score averaging is however not a good strategy for combining individual scores. It is therefore more sensible to assign a weight to each individual classifier according to its estimated proficiency when evaluated on samples within an optimisation set. The most proficient classifier is typically assigned the largest weight. The aforementioned weights are subsequently applied to the scores of questioned samples within an evaluation set.

Within the context of dynamic score fusion, weights are assigned to the scores of individual classifiers by extracting information from a specific questioned sample. This assignment therefore occurs during system implementation. The expected proficiency of a classifier may be quantified by some sort of response associated with said classifier. Within the context of the SDRT, the maximum projection value associated with a specific base classifier may, for example, constitute such a response.

6.2.2.2 Decision level fusion

Decision level fusion techniques combine the decisions of various individual classifiers in order to obtain an optimal, combined decision. Different weights may be assigned to the decisions from different classifiers. Alternatively, certain classifiers may be selected, effectively ignoring the decisions from the remaining classifiers. Decision level fusion strategies are discussed in detail in

Ruta and Gabrys (2000), Moreno-Seco *et al.* (2006) and Ponti (2011). Similar to score fusion, decision level fusion may be categorised as either static or dynamic (Valdovinos *et al.* (2005)).

Within the context of static decision level fusion, the decisions of the individual base classifiers are combined by assigning larger weights to more proficient classifiers. These weights are determined by considering an optimisation set and remain subsequently unchanged (static) for all questioned samples in the evaluation set. Alternatively, a subset of decisions associated with the most proficient individual classifiers is selected, effectively ignoring the decisions of the remaining classifiers. Majority voting constitutes a very specific example of a static decision level fusion strategy which simply accepts the decision of the majority of classifiers. Further investigation into decision level fusion is therefore warranted.

Within the context of dynamic decision level fusion, various strategies may be employed in order to determine the expected proficiency of an individual classifier based on the specific questioned sample being investigated. Weights are therefore assigned to the decisions of the individual base classifiers during system implementation.

The feasibility of all of the strategies discussed in this section for improving the proficiency of the systems proposed in this thesis should therefore be investigated.

6.2.3 Alternative data sets

The systems proposed in this thesis were evaluated by considering Dolfing's data set that contains ideal (noiseless) offline signatures with a fixed stroke width. Since these systems are however expected to be robust in the presence of moderate levels of noise, the following non-ideal (noisy) offline signature data sets may also be considered for evaluation purposes in the future:

- The MCYT-SignatureOff-75 data set that contains a total of 1125 genuine samples and 1125 skilled forgeries from 75 individual writers;
- The GPDS960 data set that contains offline signatures from 960 individual writers, with 24 genuine samples and 30 skilled forgeries for each writer.

Note that the proposed systems may be applied to binary or grey-scale images, while alternative existing feature extraction techniques, like those based on the Hough transform, are only applicable to binary images. The applicability of the proposed systems to RGB images may also be investigated in future research.

Furthermore note that, although image pre-processing (noise removal) is deemed to be outside the scope of this thesis, it should be incorporated as an additional step when the above-mentioned noisy data sets are considered for experimentation purposes. Even without pre-processing, the improvement in

system proficiency associated with the systems proposed in this thesis when compared to the SDRT-based benchmark system is expected to persist. However, since the systems proposed in this thesis were evaluated on ideal (noiseless) data, one can not expect the proficiency of said systems to be unaffected when evaluated on non-ideal data, unless noise is adequately removed.

The application of simple morphological operators or an adaptive median filter to the images in the noisy data sets mentioned earlier may be sufficient to remove noise in the majority of scenarios. The effectiveness of noise removal does however depend on the type and density of noise. Morphological thinning or the morphological dilation of a skeletonised signature constitute strategies that may be investigated for obtaining a uniform stroke width. The aforementioned adaptive median filter and morphological operators are discussed in detail in Gonzales and Woods (2010). It is however important to note that most morphological operators are only applicable to binary images. In the case of grey-scale images, appropriate binarisation has to be performed beforehand. A generic binarisation algorithm is detailed in Otsu (1979).

6.2.4 Alternative authentication scenarios

The proposed generic systems are well-suited for the recognition or verification of all two-dimensional images that contain prominent curved lines. The following authentication scenarios are therefore deemed suitable alternatives to offline signature verification (the application investigated in this thesis):

- Other biometric applications, like hand-vein recognition, iris recognition, ear recognition, etc;
- Health-related object detection;
- Vehicle-related object detection;
- Surveillance;
- Remote sensing.

In Figure 6.2 a number of images are presented that may, for example, be used in future research for the purpose of object detection and/or authentication. These examples are however not exhaustive. Note that, for some of these images pre-processing and/or proper object segmentation may be required before the systems proposed in this thesis are applied for authentication purposes.

6.2.5 Alternatives to template matching

In this thesis template matching is employed in order to quantify the difference between a reference sample and a questioned sample. This ultimately

The number of origins chosen for the MPT N_{MPT} may be increased in order to improve system proficiency. Since the calculation of the MPTs is currently the most costly component within the proposed systems, this may however lead to a significant decrease in computational efficiency. Alternatively, the number of projection angles N_θ may be increased within the context of the SDRT-based benchmark system or in scenarios where the MDRT is applied to an MPT. Since a maximum of *eight* projection angles is employed within the context of the systems proposed in this thesis, there is ample scope for improvement in system proficiency. Such an adjustment will however decrease computational efficiency since it necessitates the concatenation of additional duplicates of MPTs within the context of the MDRT. The set of distances d that is specified for the scenario where the SDRT is applied to an MPT from an angle of 90° , may also be adjusted. A larger set of distances will inevitably lead to a larger number of curvatures being detected. This will however also decrease computational efficiency, as is the case for the other two potential adjustments.

It is therefore clear that the potential adjustment of system parameters depends on the principal objectives of the authentication scenario. Since there is an inevitable trade-off between system proficiency and computational efficiency, the most important objective will dictate how the system parameters are specified.

6.2.7 Improved computational efficiency

Since the SDRT is deemed a very efficient feature extraction strategy, efforts in improving the efficiency of the systems proposed in this thesis will focus on the calculation of the various MPTs. In order to facilitate improved efficiency, the calculation of a smaller number of MPTs should be considered. The MPTs of the composite images (conceptualised in Section 3.4.2) may also be optimised. This may however prove to be a challenging task since we originally set out to extract as much information as possible from an input image.

It should therefore be clear that the work conducted in this thesis opens up a wide range of future research opportunities.

Bibliography

- Al-Omari, Y., Abdullah, S.N.H.S. and Omar, K. (2011). State-of-the-art in offline signature verification system. *International Conference on Pattern Analysis and Intelligent Robotics*, pp. 59-64.
- Ballard, D.H. (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2), pp. 111-122.
- Beylkin, G. (1987). Discrete Radon transform. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 35(2), pp. 162-172.
- Castellano, G., Bonilha, L., Li, L.M. and Cendes, F. (2004). Texture analysis of medical images. *Clinical radiology*, 59(12), pp. 1061-1069.
- Coetzer, J. (2005). *Off-line Signature Verification*. PhD thesis, Stellenbosch University.
- Coetzer, J., Herbst, B.M. and du Preez, J.A. (2004). Offline signature verification using the discrete Radon transform and a hidden Markov model. *EURASIP Journal on applied signal processing*, pp. 559-571.
- Coetzer, J., Swanepoel, J.P. and Sabourin, R. (2012). Efficient cost-sensitive human-machine collaboration for offline signature verification. *DRR*, pp. 82970J.
- Cohn, D.A., Ghahramani, Z. and Jordan, M.I. (1996). Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4, pp. 129-145.
- Cressie, N. and Whitford, H. (1986). How to use the two sample *t*-test. *Biometrical Journal*, 28(2), pp. 131-148.
- Deller, J.R., Proakis, J.G. and Hansen, J.H. (1999). *Discrete-Time Processing of Speech Signals*. Wiley-Blackwell; Reprint edition.
- Dolfing, J. (1998). *Handwriting recognition and verification - a hidden Markov approach*. PhD thesis, Eindhoven University of Technology.
- Dowdy, S., Wearden, S. and Chilko, D. (2011). *Statistics for research*, 3rd edition. John Wiley & Sons.
- Fick, C., Coetzer, J. and Swanepoel, J.P. (2016). Efficient curve-sensitive features for offline signature verification. *Pattern Recognition Association of South Africa*

- and *Robotics and Mechatronics International Conference (PRASA-RobMech)*, pp. 155-160.
- Gonzales, R. and Woods, R. (2010). *Digital Image Processing*, 3rd international edition. Prentice-Hall Inc.
- Hafemann, L.G., Sabourin, R. and Oliveira, L.S. (2015). Offline handwritten signature verification-literature review. *arXiv preprint arXiv:1507.07909*.
- Hai-peng, Y., Yi-xing, L. and Zhen-bo, L. (2005). Auto detection of wood texture orientation by Radon transform. *Journal of Forestry Research*, 16(1), pp. 1-4.
- Henniger, O. and Muller, S. (2007). Effects of time normalization on the accuracy of dynamic time warping. *International Conference on Biometrics: Theory, Applications, and Systems*, pp. 1-6.
- Hjouj, F. and Kammler, D.W. (2008). Identification of reflected, scaled, translated, and rotated objects from their Radon projections. *IEEE Transactions on Image Processing*, 17(3), pp. 301-310.
- Hou, W., Ye, X. and Wang, K. (2004). A survey of off-line signature verification. *International Conference on Intelligent Mechatronics and Automation*, pp. 536-541.
- Illingworth, J. and Kittler, J. (1988). A survey of the Hough transform. *Computer vision, graphics, and image processing*, 44(1), pp. 87-116.
- Impedovo, D. and Pirlo, G. (2008). Automatic signature verification: the state of the art. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 38(5), pp. 609-635.
- Jafari-Khouzani, K. and Soltanian-Zadeh, H. (2005). Rotation-invariant multiresolution texture analysis using Radon and wavelet transforms. *IEEE transactions on image processing*, 14(6), pp. 783-795.
- Jain, A., Nandakumar, K. and Ross, A. (2005). Score normalization in multimodal biometric systems. *Pattern Recognition*, 38(12), pp. 2270-2285.
- Jayadevan, R., Kolhe, S. and Patil, P. (2009). Dynamic time warping based static hand printed signature verification. *Journal of Pattern Recognition Research (JPRR)*, 4(1), pp. 52-65.
- Keogh, E. and Pazzani, M. (2001). Derivative dynamic time warping. *SIAM International Conference on Data Mining*, pp. 1-11.
- Kekre, H.B. and Bharadi, V.A. (2010). Gabor filter based feature vector for dynamic signature recognition. *International Journal of Computer Applications*, 2(3), pp. 74-80.
- Khan, S. and Dhole, A. (2014). A review on offline signature recognition and verification techniques. *International Journal of Advanced Research in Computer and Communication Engineering*, 3(6), pp. 6879-6882.

- Kumar, A., Chan, T.S. and Tan, C.W. (2012). Human identification from at-a-distance face images using sparse representation of local iris features. *In Biometrics (ICB), 5th IAPR International Conference*, pp. 303-309.
- Li, H., Zheng, H. and Wang, Y. (2007). Segment Hough transform – A novel Hough-based algorithm for curve detection. *Fourth International Conference on Image and Graphics*, pp. 471-477.
- Moreno-Seco, F., Inesta, J. M., De León, P. J. P. and Micó, L. (2006). Comparison of classifier fusion methods for classification in pattern recognition tasks. *SSPR/SPR*, pp. 705-713.
- Mukhopadhyay, P. and Chaudhuri, B.B. (2015). A survey of Hough Transform. *Pattern Recognition*, 48(3), pp. 993-1010.
- Nair, P.S. and Saunders, A.T. (1996). Hough transform based ellipse detection algorithm. *Pattern Recognition Letters*, 17(7), pp. 777-784.
- Nanni, L., Lumini, A. and Brahnam, S. (2010). Local binary patterns variants as texture descriptors for medical image analysis. *Artificial Intelligence in Medicine*, 49, 117-125.
- Nguyen, U.T., Bhuiyan, A., Park, L. A. and Ramamohanarao, K. (2013). An effective retinal blood vessel segmentation method using multi-scale line detection. *Pattern recognition*, 46(3), pp. 703-715.
- Nikam, S. B. and Agarwal, S. (2008). Texture and wavelet-based spoof fingerprint detection for fingerprint biometric systems. *First IEEE International Conference on Emerging Trends in Engineering and Technology ICETET'08*, pp. 675-680.
- Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1), pp 62-66.
- Panton, M. and Coetzer, J. (2010). Off-line signature verification using ensembles of local Radon transform-based HMMs. *Annual Symposium of the Pattern Recognition Association of South Africa*, pp. 201-206.
- Ponti Jr, M. P. (2011). Combining classifiers: from the creation of ensembles to the decision fusion. *Graphics, Patterns and Images Tutorials (SIBGRAPI-T)*, 2011 24th SIBGRAPI Conference, IEEE, pp. 1-10.
- Pun, C.M. and Lee, M.C. (2003). Log-polar wavelet energy signatures for rotation and scale invariant texture classification. *IEEE transactions on pattern analysis and machine intelligence*, 25(5), pp. 590-603.
- Ross, A.A., Jain, A.K. and Nandakumar, K. (2006). *Handbook of Multibiometrics*. International Series on Biometrics, 6. Springer.
- Ruta, D. and Gabrys, B. (2000). An Overview of Classifier Fusion Methods. *Computing and Information Systems*, 7(1), pp. 1-10.

- Seo, J.S., Haitsma, J., Kalker, T. and Yoo, C.D. (2004). A robust image fingerprinting system using the Radon transform. *Signal Processing: Image Communication*, 19(4), pp. 325-339.
- Shah, A.S., Khan, M.N.A. and Shah, A. (2015). An Appraisal of Off-line Signature Verification Techniques. *International Journal of Modern Education and Computer Science*, 7(4), pp. 67-75.
- Sivaraman, S. and Trivedi, M.M. (2013). Looking at vehicles on the road: A survey of vision-based vehicle detection, tracking, and behavior analysis. *IEEE Transactions on Intelligent Transportation Systems*, 14(4), pp. 1773-1795.
- Swanepoel, J. and Coetzer, J. (2010). Off-line signature verification using flexible grid features and classifier fusion. *International Conference on Frontiers in Handwriting Recognition*, pp. 297-302.
- Swanepoel, J.P. (2015). *Writer-independent Handwritten Signature Verification*. PhD thesis, Stellenbosch University.
- Toft, P. (1996). *The Radon Transform – Theory and Implementation*. PhD thesis, Technical University of Denmark.
- Valdovinos, R.M., Sánchez, J.S. and Barandela, R. (2005). Dynamic and static weighting in classifier fusion. *Iberian conference on Pattern recognition and image analysis*, pp. 59-66. Springer, Berlin, Heidelberg.
- Vankayalapati, H.D. and Kyamakya, K. (2009). Nonlinear feature extraction approaches with application to face recognition over large databases. In *Nonlinear Dynamics and Synchronization (INDS) 2nd International Workshop*, pp. 44-48.
- Wang, X., Xiao, B., Ma, J.F. and Bi, X.L. (2007). Scaling and rotation invariant analysis approach to object recognition based on Radon and Fourier-Mellin transforms. *Pattern Recognition*, 40(12), pp. 3503-3508.
- Wang, L. and Geng, X. (2009). *Behavioral Biometrics for Human Identification: Intelligent Applications: Intelligent Applications*. IGI Global.
- Wei, Z., Qiu, X., Sun, Z. and Tan, T. (2008). Counterfeit iris detection based on texture analysis. *19th International Conference on Pattern Recognition ICPR*, pp. 1-4.
- Wu, J.D. and Ye, S.H. (2009). Driver identification using finger-vein patterns with Radon transform and neural network. *Expert Systems with Applications*, 36(3), pp. 5793-5799.
- Wu, S., Weinstein, S.P., Conant, E.F., Schnall, M.D. and Kontos, D. (2013). Automated chest wall line detection for whole-breast segmentation in sagittal breast MR images. *Medical physics*, 40(4).
- Yadav, M., Kumar, A., Patnaik, T. and Kumar B. (2013). A survey on offline signature verification. *International Journal of Engineering and Innovative Technology*, 2(7), pp. 337-340.

- Yu, B. and Jain, A.K. (1997). Lane boundary detection using a multiresolution hough transform. *Proceedings of the International Conference on Image Processing*, 2, pp. 748-751.
- Zhang, X. and Zheng, N. (2010). Vehicle detection under varying poses using conditional random fields. *13th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pp. 875-880.