



FEDERAL UNIVERSITY OF SERGIPE
SCIENCE AND TECHNOLOGY CENTER
GRADUATE PROGRAM IN COMPUTER SCIENCE

**Analysis and evaluation of deep learning based
super-resolution algorithms to improve performance in
low-resolution face recognition**

Master's Thesis

Angelo Garangau Menezes



São Cristóvão – Sergipe

2019

FEDERAL UNIVERSITY OF SERGIPE
SCIENCE AND TECHNOLOGY CENTER
GRADUATE PROGRAM IN COMPUTER SCIENCE

Angelo Garangau Menezes

**Analysis and evaluation of deep learning based
super-resolution algorithms to improve performance in
low-resolution face recognition**

Master's Thesis presented to the Graduate Program in
Computer Science of the Federal University of Sergipe
as a partial fulfillment of the requirements for the De-
gree of Master of Science in Computer Science.

Supervisor: Carlos Alberto Estombelo-Montesco

São Cristóvão – Sergipe

2019

**FICHA CATALOGRÁFICA ELABORADA PELA BIBLIOTECA CENTRAL
UNIVERSIDADE FEDERAL DE SERGIPE**

M543a Menezes, Angelo Garangau
Analysis and evaluation of deep learning based super-resolution algorithms to improve performance in low-resolution face recognition / Angelo Garangau Menezes ; orientador Carlos Alberto Estombelo-Montesco . - São Cristóvão, 2019.
68 f. : il.

Dissertação (mestrado em Ciência da Computação) – Universidade Federal de Sergipe, 2019.

1. Computação. 2. Percepção facial. 3. Algoritmos computacionais. 4. Identificação biométrica. 5. Redes neurais (Computação). I. Estombelo-Montesco, Carlos Alberto orient. II. Título.

CDU 004




UNIVERSIDADE FEDERAL DE SERGIPE
PRÓ-REITORIA DE PÓS-GRADUAÇÃO E PESQUISA
COORDENAÇÃO DE PÓS-GRADUAÇÃO
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO


Ata da Sessão Solene de Defesa da Dissertação do
Curso de Mestrado em Ciência da Computação-UFS.
Candidato: ANGELO GARANGAU MENEZES

Em 12 dias do mês de dezembro do ano de dois mil e dezenove, com início às 15h00min, realizou-se na Sala de Seminário do DCOMP da Universidade Federal de Sergipe, na Cidade Universitária Prof. José Aloísio de Campos, a Sessão Pública de Defesa de Dissertação de Mestrado do candidato **ANGELO GARANGAU MENEZES**, que desenvolveu o trabalho intitulado: "*Analysis and evaluation of deep learning based super-resolution algorithms to improve performance in low-resolution face recognition*", sob a orientação do Prof. Dr. **Carlos Alberto Estombelo Montesco**. A Sessão foi presidida pelo Prof. Dr. **Carlos Alberto Estombelo Montesco** (PROCC/UFS), que após a apresentação da dissertação passou a palavra aos outros membros da Banca Examinadora, Prof. Dr. **Daniel Oliveira Dantas** (PROCC/UFS) e, em seguida, ao Prof. Dr. **CARLOS ALBERTO VILLACORTA CARDOSO** (UFS). Após as discussões, a Banca Examinadora reuniu-se e considerou o mestrando (a) **APROVADO** "(aprovado/reprovado)". Atendidas as exigências da Instrução Normativa 01/2017/PROCC, do Regimento Interno do PROCC (Resolução 67/2014/CONEPE), e da Resolução nº 25/2014/CONEPE que regulamentam a Apresentação e Defesa de Dissertação, e nada mais havendo a tratar, a Banca Examinadora elaborou esta Ata que será assinada pelos seus membros e pelo mestrando.

Cidade Universitária "Prof. José Aloísio de Campos", 12 de dezembro de 2019.


Prof. Dr. Carlos Alberto Estombelo Montesco
(PROCC/UFS)
Presidente


Prof. Dr. Daniel Oliveira Dantas
(PROCC/UFS)
Examinador Interno


Prof. Dr. Carlos Alberto Villacorta Cardoso
(UFS)
Examinador Externo


Angelo Garangau Menezes
Candidato

Acknowledgements

Firstly, I would like to say that I am thankful to God for creating this perfect simulation that we live, and for supporting me to get this far in this incredible adventure that we call life.

To my family, Roberto, Adilma, and Apolo, for all the support, love, and consideration that I have had all my life. You guys are the reason why I want to become a better version of myself every day.

To my advisor Prof. Dr. Carlos Estombelo, for believing and guiding me through the course of my master's program while also being an incredibly understanding person. Thanks for your friendship and for being hard on me when I needed it.

To the best people in the world that have inspired me, stayed by my side in the hardest moments, and have helped this thesis come to life directly with their support and love, Fernando Melo, Gracieth Cavalcanti, Barbara Sena, and Rita Macedo.

To Prof. Dr. André Carvalho and all the friends that I made in São Paulo while working at USP, for their amazing friendship, research insights, and support while doing some “balbúrdia” inside and outside the lab.

To Prof. Dr. Vijay Mago, for having introduced me to the field of data science and made me believe that with the right amount of effort, I would be able to learn anything and produce valuable research.

To Prof. Dr. Wilson Wang and Dr. Peter Luong, for having shaped me as a researcher and taught me how to overcome all the difficulties that academia could possibly have.

To all the great friends who understood my absence in certain moments and that I hope will always be with me Grace Kelly, Natalia Rosa, Thiago Charles, Raul Rodrigo, Duda Maia, Renan Albuquerque, Felipe Torres, Vinicius Araujo, Ronny Almeida, Davi Santana, Eriana Pinto, Manu Magno, and all the incredible others that I will probably have to pay a beer since their names are not here.

*“Quando disser sim para os outros, certifique-se
de não estar dizendo não para si mesmo.”
(Paulo Coelho)*

Resumo

Os cenários de vigilância e monitoramento estão propensos a vários problemas, pois não existe um controle sobre a distância dos possíveis suspeitos para a câmera e geralmente as tarefas envolvem avaliação de imagens em baixa resolução. Para tais situações, a aplicação de algoritmos de *super-resolution* (super-resolução) pode ser uma alternativa adequada para recuperar as propriedades discriminantes das faces dos suspeitos envolvidos.

Embora abordagens gerais de super-resolução tenham sido propostas para aprimorar a qualidade da imagem para a percepção no nível humano, os métodos de super-resolução biométrica buscam a melhor versão da imagem para “percepção” do computador, pois seu foco é melhorar o desempenho do reconhecimento automático. Redes neurais convolucionais e algoritmos de aprendizado profundo, em geral, têm sido aplicados a tarefas de visão computacional e agora são o estado da arte em seus vários subdomínios, incluindo classificação, restauração e super-resolução de imagens. No entanto, poucos trabalhos avaliaram os efeitos que os mais recentes métodos de super-resolução propostos podem ter sobre a precisão e o desempenho da verificação de faces em imagens de baixa resolução do mundo real.

Este projeto teve como objetivo avaliar e adaptar diferentes arquiteturas de redes neurais profundas para a tarefa de super-resolução de faces, impulsionada pelo desempenho do reconhecimento de faces em imagens de baixa resolução do mundo real. Os resultados experimentais em um conjunto de dados de monitoramento/vigilância e avaliação de presença universitária mostraram que arquiteturas gerais de super-resolução podem melhorar o desempenho da verificação de faces utilizando uma rede neural profunda treinada em faces de alta resolução para extração de características. Além disso, como as redes neurais são aproximadores de funções e podem ser treinadas com base em funções objetivo específicas, o uso de uma função de custo personalizada que foi otimizada para extração de características da face mostrou resultados promissores para recuperar atributos discriminantes em imagens de faces em baixa resolução.

Palavras-chave: Reconhecimento Facial em Baixa Resolução; Super-Resolução; Aprendizado Profundo; Redes Neurais Convolucionais.

Abstract

Surveillance scenarios are prone to several problems since they usually involve low-resolution footage, and there is no control of how far the subjects may be from the camera in the first place. This situation is suitable for the application of upsampling (super-resolution) algorithms since they may be able to recover the discriminant properties of the subjects involved.

While general super-resolution approaches were proposed to enhance image quality for human-level perception, biometrics super-resolution methods seek the best “computer perception” version of the image since their focus is on improving automatic recognition performance. Convolutional neural networks and deep learning algorithms, in general, have been applied to computer vision tasks and are now state-of-the-art for several sub-domains, including image classification, restoration, and super-resolution. However, no work has evaluated the effects that the latest proposed super-resolution methods may have upon the accuracy and face verification performance in low-resolution “in-the-wild” data.

This project aimed at evaluating and adapting different deep neural network architectures for the task of face super-resolution driven by face recognition performance in real-world low-resolution images. The experimental results in a real-world surveillance and attendance datasets showed that general super-resolution architectures might enhance face verification performance of deep neural networks trained on high-resolution faces. Also, since neural networks are function approximators and can be trained based on specific objective functions, the use of a customized loss function optimized for feature extraction showed promising results for recovering discriminant features in low-resolution face images.

Key-Words: Low-Resolution Face Recognition; Super-Resolution; Deep Learning; Convolutional Neural Networks;

List of Figures

Figure 1 – CNN Architecture Exemplified (DESHPANDE, 2017)	20
Figure 2 – Example of Inception Block. (Source: Author’s own)	21
Figure 3 – GoogleNet architecture. (SZEGEDY et al., 2015)	21
Figure 4 – Single Residual Block. (Source: Author’s own)	22
Figure 5 – Architecture of a GAN. (Source: Author’s own)	23
Figure 6 – Differences between normal and coordinate convolutions. (LIU et al., 2018)	24
Figure 7 – General classes of SR algorithms. (HUANG; LIU, 2015)	25
Figure 8 – Visual comparison of general interpolation methods: (a) Nearest Neighbor (b) Bilinear (c) Bicubic (d) Original HD image. - (Source: Author’s own) . .	26
Figure 9 – DL for SR algorithms related topics. (Adapted from Wang, Chen and Hoi (2019))	27
Figure 10 – Example of a generic pipeline for face recognition. (Source: Author’s own) .	28
Figure 11 – Example of face images in the VGGFace2 dataset. (Adapted from Cao et al. (2018))	36
Figure 12 – Example of face images in the CelebA dataset. (Adapted from Liu et al. (2015))	37
Figure 13 – Example of gallery images in the Quis-Campi dataset. (Adapted from Neves, Moreno and Proença (2017))	38
Figure 14 – Example of probe images in the Quis-Campi dataset. (Adapted from Neves, Moreno and Proença (2017))	38
Figure 15 – Example of gallery images for the UFS-Classroom Attendance dataset. (SÁ, 2019)	39
Figure 16 – Example of a probe image for the UFS-Classroom Attendance dataset. (SÁ, 2019)	39
Figure 17 – Example of a probe face image from the UFS-Classroom Attendance dataset saved in the three settings. (Adapted from Sá (2019))	40
Figure 18 – Example of a probe face image from the Quis-Campi dataset saved in the three settings. (Adapted from Neves, Moreno and Proença (2017))	40
Figure 19 – Pipeline for face verification in the ICB-RW. (Source: Author’s own)	41
Figure 20 – SRCNN architecture. (DONG; LOY; TANG, 2016)	44
Figure 21 – Subpixel CNN architecture. (Adapted from Dong, Loy and Tang (2016) and Shi et al. (2016))	44
Figure 22 – FSRCNN architecture. (DONG; LOY; TANG, 2016)	44
Figure 23 – SRGAN architecture. (JIAO; ZHAO, 2019)	44
Figure 24 – Watch-List setting for the ICB-RW. (NEVES; PROENÇA, 2016)	46

Figure 25 – Performance results for accuracy on ICB-RW and UFS Classroom 1 data.
Proposed architectures have an asterisk in their names.(Source: Author’s own) 52

Figure 26 – Performance results for accuracy on UFS Classroom 2 and UFS Classroom 3
data. Proposed architectures have an asterisk in their names. (Source: Author’s
own) 52

List of Tables

Table 1 – Number of students for each class	47
Table 2 – PSNR and SSIM validation results (2000 images from CelebA)	48
Table 3 – Rank-1 Accuracy (in %) for Face Recognition task (1xN) on 90 subjects from the Quis-Campi dataset (ICB-RW)	49
Table 4 – Accuracy (in %) for Face Recognition task in Classroom 1 (1xN)	50
Table 5 – Accuracy (in %) for Face Recognition task in Classroom 2 (1xN)	51
Table 6 – Accuracy (in %) for Face Recognition task in Classroom 3 (1xN)	51
Table 7 – Spearman Correlation PSNR/SSIM vs. Accuracy	52

List of abbreviations and acronyms

CoordConv	Coordinate Convolution
DL	Deep Learning
FPS	Frames per second
FSRCNN	Fast Super-Resolution Convolutional Neural Network
GAN	Generative Adversarial Network
GPU	Graphical Processor Unit
HR	High-Resolution
LFW	Labeled Faces in the Wild
LR	Low-Resolution
MSE	Mean Squared Error
PSNR	Peak signal-to-noise ratio
SR	Super-Resolution
SRCNN	Super-Resolution Convolutional Neural Network
SubCNN	Subpixel Convolutional Neural Network
SRGAN	Super-Resolution Generative Adversarial Network
SOTA	State-of-the-art
SSIM	Structural Similarity
UFS	Federal University of Sergipe

List of symbols

β Greek letter Beta

ϕ Greek letter Phi

$\mathbf{R}$ Real space

θ Theta

Contents

1	Introduction	15
1.1	Hypotheses	17
1.2	Objectives	17
1.3	Thesis Structure	18
2	Technical Background	19
2.1	Convolutional Neural Networks and Deep Learning	19
2.1.1	Residual Networks	20
2.1.2	Generative Adversarial Networks	22
2.1.3	Coordinate Convolutions	23
2.2	Super-Resolution	24
2.2.1	Operating Channels	25
2.2.2	Super-Resolution Benchmarking	26
2.2.3	Deep Learning for Image Super-Resolution	27
2.3	Face Recognition	28
2.3.1	Face Detection	28
2.3.2	Feature Extraction and Face Verification	29
2.4	Final Considerations	30
3	Related Work	31
3.1	Super-Resolution	31
3.2	Low-Resolution Face Recognition	33
3.3	Final Considerations	35
4	Methodology	36
4.1	Datasets	36
4.1.1	VGGFace2	36
4.1.2	CelebA	36
4.1.3	Quis-Campi Dataset (ICB-RW)	37
4.1.4	Federal University of Sergipe Classroom Attendance Dataset	38
4.2	Data Pre-Processing	38
4.3	Transfer Learning	40
4.3.1	Face Feature Extraction	40
4.3.2	Face Verification	41
4.3.3	Face Loss	42
4.4	Final Considerations	42

5	Experimental Results	43
5.1	Experiments	43
5.1.1	Task 1 - Face Super-Resolution	43
5.1.2	Task 2 - Watch-List ICB-RW (1x5 Problem)	46
5.1.3	Task 3 - Attendance Evaluation (1xN Problem)	47
5.2	Results Evaluation	47
5.2.1	Task 1 - Face Super-Resolution	48
5.2.2	Task 2 - Watch-List ICB-RW (1x5 Problem)	49
5.2.3	Task 3 - Attendance Evaluation (1xN Problem)	50
5.2.4	Hypotheses Discussion	52
5.3	Final Considerations	53
6	Conclusions	55
	Bibliography	57
	Appendix	63
	APPENDIX A Perceptual Results of SR Algorithms for 4x Upscaling	64
	APPENDIX B Average Training Losses for the SR Algorithms	69

1

Introduction

An essential ability in human beings that groups them as social animals is face perception. Infants tend to prefer to look at faces at a very early age, and across the lifespan, most people spend more time looking at faces than at any other type of object ([JOHNSON et al., 1991](#)).

Faces provide a wealth of information that facilitates social communication since humans are able to recognize the identity of other people and interpret their emotional state by analyzing the facial expression and pose. More specifically, regarding identity recognition, there is behavioral and neural evidence that such a feature has its basis on the perception of aspects of facial structure that are invariant across changes ([GOBBINI; HAXBY, 2007](#); [HAXBY; HOFFMAN; GOBBINI, 2000](#)).

Face perception is also related to a high-level visual and memory process that involves the retrieval of the memory of faces and the identity information stored in memory (i.e., person semantic knowledge). This process is developed in such a robust way in human brains that some people are able to recognize others by situations where there are only a few resembling features of a person, such as in caricature drawings and photos with low-resolution ([CHANG et al., 2017](#)). The field of research that describes and evaluates the reliable methods for automatic identification of subjects based on their physiological and behavioral characteristics is usually called biometrics ([NGUYEN et al., 2018](#)).

As an example of how face biometrics has become an important matter in modern society, situations in surveillance that employ the verification of a watch-list of subjects through CCTV footage have become quite regular for world security standards in airports, malls, and other crowded places. However, as sometimes they do not involve automation, they might become a weak spot as they require an impressive amount of manual work to check the live feed or saved data of several cameras ([RASTI et al., 2016](#)). This is one of the reasons why countries are spending a large number of resources to rapidly grow their technology market related to surveillance in order to have intelligible solutions specifically designed to their

needs (FELDSTEIN, 2019).

Even though computers have shown a great ability to also deal with image and face recognition in the last decade, in situations where low-resolution (LR) inputs are employed, they tend to fail as much as humans when trying to identify an individual or reconstruct a higher-resolution representation of the same subject (NGUYEN et al., 2018). These occurrences are the majority in surveillance scenarios since the cheapest and most commonly used cameras can only provide low-quality video footage, and there is no control for the distance between the subjects of interest and the device (RASTI et al., 2016).

These recognition faults mainly occur because when the resolution drops, the amount of information available for identifying or verifying a subject decreases as well. That leads to a severe degradation for both human perception and machine interpretation. Since there is no standard resolution that can be set for making recognition available (NGUYEN et al., 2018), the development of image upscaling algorithms, commonly known as super-resolution (SR) algorithms, has become an intensive area of research. An example of that is the fact that the pioneering work of this group of algorithms dates back to 1974, when Gerchberg (1974) showed that the resolution of a data object could be significantly improved through error energy reduction. Thenceforth, researchers have put a massive effort into investigating SR and its possible range of applications, even knowing that it is fundamentally an ill-posed problem since the details presented in the LR samples are usually not enough to provide a robust reconstruction of the original high-resolution (HR) image (TIAN; MA, 2011).

Deep Learning (DL) algorithms started to be used to solve tasks regarding image classification and reconstruction due to their computational cost being now facilitated by advances in hardware and parallel processing (KRIZHEVSKY; SUTSKEVER; HINTON, 2012). This group of techniques has become the state-of-the-art (SOTA) rapidly in a great variety of tasks regarding images both for accuracy and applicability (LECUN; BENGIO; HINTON, 2015). Also, they have shown excellent performance in image restoration tasks that are related to biometrics such as iris, fingerprint, and face super-resolution for improving recognition performance (RIBEIRO; UHL, 2017; LI; FENG; KUO, 2018; KIM et al., 2019).

Most of the SR solutions for LR face recognition have relied on the use of convolutional neural networks (CNNs) optimized by a pixel loss (NGUYEN et al., 2018). Nevertheless, there exists nowadays a large pool of network designs and learning strategies that are applied to solve similar computer vision problems (HARIS; SHAKHNAROVICH; UKITA, 2018; LIU et al., 2018). Since the goal of SR for face biometrics is to optimize face recognition performance while keeping reasonable perceptual quality, replicating successful strategies from similar computer vision tasks can be a worth research direction. One example of a different strategy that some similar works have applied is the use of different types of convolution operators and customized loss functions to increase performance (WANG; SHE; WARD, 2019; WANG; CHEN; HOI, 2019).

One of the current issues with SR solutions to the LR face recognition problem is that, researchers often train their SR deep learning models reporting their accuracy results only on the downsampled version of the same or other HR frontal image dataset (OUYANG et al., 2018; ABELLO; JR., 2019). However, it is known that such task becomes more challenging when faces are captured in an unconstrained environment where they can be subject to blurring, motion, non-frontal pose, and other situations that hinder recognition.

The origin of the analysis to be presented in this thesis is related to the lack of recent studies of if and how the state-of-the-art deep learning SR techniques may assist face biometrics in real-world low-resolution scenarios, taking into consideration also different network architectures, learning strategies, and their real applicability and scalability.

1.1 Hypotheses

For the development of this thesis and the proposal of experiments, the following specific hypotheses were elaborated:

1. The relationship between image quality metrics and accuracy performance is not significant.
2. The use of a specific convolution operator that take into account position information (CoordConv) can effectively improve metric performance over normal convolution operators when dealing with super-resolution.
3. Application of a loss function based on face identity for an upscaling network (FaceLoss) can influence the verification results positively in a face recognition pipeline using DL models.

1.2 Objectives

Taking into account all the possible challenges regarding the discussed topics, the general objective of this thesis is to evaluate the efficiency of a face recognition pipeline in real-world low-resolution scenarios and check whether the recently developed SR algorithms and their variants are capable of enhancing recognition performance in these situations.

The specific objectives are listed below:

- Evaluation of the possibility of a correlation between image quality metrics and face verification accuracy in a LR recognition pipeline as considered by hypothesis 1.

- Evaluation of different SOTA neural network architectures, also involving different convolution operators as proposed in hypothesis 2, for the super-resolution task driven by face biometrics performance involving faces in real-world LR datasets.
- Evaluation of an adapted loss function that optimizes the DL model for better face feature extraction while keeping the SR upsampling characteristic as suggested by hypothesis 3.

1.3 Thesis Structure

In order to make an easier read, this thesis brings the technical background before the related work section since the discussed topics are from recent research, and a prior overview can be useful for a better comprehension of the concepts. Therefore, this manuscript was structured with the following sections:

- Chapter 1 - Introduction
- Chapter 2 - Technical Background
- Chapter 3 - Related Work
- Chapter 4 - Methodology
- Chapter 5 - Experiments
- Chapter 6 - Results
- Chapter 7 - Final Considerations

2

Technical Background

This section gives a brief technical background overview for the topics discussed in this thesis in order to provide the basics that validate the proposed experiments and hypotheses.

2.1 Convolutional Neural Networks and Deep Learning

Deep learning (DL) is a branch of machine learning that is capable of learning the data representation through the use of a structure of hierarchical layers, similar to the way the brain handles new information. Its concept is mainly applied to supervised learning problems (e.g., where there is a need for mapping an input vector to an output vector), and its core is based on the math behind Artificial Neural Networks ([LECUN; BENGIO; HINTON, 2015](#)).

Deep Neural Networks can have different architectures based on the nature of the data that is used as input. When image data needs to be processed as input, CNNs have been ideally applied by academia and industry because of its interior architecture properly set to work with high dimensional data and extract its more discriminating features. ([LECUN; BENGIO; HINTON, 2015](#); [SHI et al., 2016](#))

A typical structure of a CNN can be seen by Figure 1 where an image is used as input, and the network needs to predict a label for it. The first operation that happens inside the network is on the convolutional layer, where a moving window is applied to a small pixel grid of the image. This moving window, commonly called a kernel, works as a “filter” and its task is to multiply its weight values by the original pixel values. All these multiplications are summed up to one number that is going to be placed on the matrix used as input on the following layer.

The CNN per se consists of several stacked convolutional networks mixed with nonlinear and pooling layers that work as feature extractors. Usually, the nonlinear layer is added after each convolution operation, which brings a nonlinear property characteristic to the network through the use of an activation function. The pooling layer will then be placed after the nonlinear layer




Figure 1 – CNN Architecture Exemplified (DESHPANDE, 2017)

working directly with the width and height of the image in order to perform a downsampling operation. This step reduces the image data to a more compressed version containing only details that were processed and identified by the previous filter (convolutional) layer. After a series of “feature extraction” layers, a fully connected layer is generally stacked upon them in order to map the extracted features to a fixed output.

The learning phase of a CNN happens on the update of the weights presented on every convolutional layer and the weights for the fully connected one. The first often allows the network to identify edges, contours, and shapes that characterize the image while the second is accountable for the classification or regression step. The training is usually performed using variants of gradient-based optimization methods via backpropagation (KRIZHEVSKY; SUTSKEVER; HINTON, 2012; LECUN; BENGIO; HINTON, 2015).

2.1.1 Residual Networks

When training large image classifiers, usually there is a considerable variation in the location and size of the object of interest. In order to have a robust feature extractor that identifies features that are globally or locally distributed on the image, the use of different kernel sizes may be needed. With this in mind, Szegedy et al. (2015) proposed GoogleNet using large blocks that contained different convolution operators with several kernel sizes. One representation of such block is shown in Figure 2.

Using several stacks of blocks in a very deep network, they were able to achieve 93.3% top-5 accuracy on the ImageNet competition with much less computation than the state-of-the-art (SOTA) at that time, VGG16. The final architecture of GoogleNet can be seen in Figure 3.

Nevertheless, as academia started to implement and test different types of deep architectures, the problem of vanishing gradients became popular. This issue appears because certain activation functions squish an ample input space into the range between 0 and 1. Then, sometimes even when a large change arrives in the input, the output is going to have only a




Figure 2 – Example of Inception Block. (Source: Author's own)




Figure 3 – GoogleNet architecture. (SZEGEDY et al., 2015)

minor change, and consequently, the gradients become too small for updating the weights when backpropagated (LECUN; BENGIO; HINTON, 2015).

One solution that researchers found to this problem was to use skip connections. These connections, as shown by Figure 4, are used to feed posterior layers the same input that previous layers had, which makes the network skip the training of a few layers and learn only the residual between the input and the output (HE et al., 2016).

This structure gave the name for the group of residual networks, commonly known as ResNets, and influenced researchers to go even deeper since networks consequently could have more layers and still train in sufficient time. One of the examples of such structure in SOTA applications is in the work of Szegedy et al. (2017), where inception and residual blocks are combined to create robust feature extractors.




Figure 4 – Single Residual Block. (Source: Author’s own)

2.1.2 Generative Adversarial Networks

Generative Adversarial Networks (GANs) were proposed by [Goodfellow et al. \(2014\)](#) in order to sidestep the common difficulties that involve deep generative models such as approximating intractable probabilistic computations that arise in maximum likelihood estimation and leveraging the benefits of piecewise linear units in the generative context.

In this architecture, a discriminator network $D(x)$, where x is an image, is optimized for distinguishing whether the given input is fake or not, while a generator network $G(x)$, where x can be random noise or even another image, is optimized to generate fake image samples that follow the same distribution of the real image and fool the discriminator from discerning which one is the real ([WANG; SHE; WARD, 2019](#)). Therefore, in this context, the output for the discriminator network is always a label (real $\rightarrow 1$, fake $\rightarrow 0$), and for the generator is always an image. The general idea presented in the learning process is shown in Figure 5.

In other words, D is trained to maximize the probability of assigning the same correct label for both generated and real images, while simultaneously G is trained to minimize $\log(1 - D(G(z)))$. [Goodfellow et al. \(2014\)](#) described their optimization, also known as adversarial training, as the play of a minimax game with value function $V(D, G)$:

$$\min(G) \max(D) \rightarrow V(D, G) = E_{x \sim p_{data}(x)}[\log D(x)] + E_{z \sim p_z(z)}[\log(1 - D(G(x)))] \quad (2.1)$$

given $p_z(z)$ as the input noise and considering E_x and E_z the error associated with discriminator and generator, respectively.

Since then, GANs attracted growing interests in the research community due to their




Figure 5 – Architecture of a GAN. (Source: Author's own)

applicability and versatility. They have been applied to various domains such as natural language processing, time-series synthesis, and computer vision (YANG et al., 2017; DONAHUE; MCAULEY; PUCKETTE, 2018; BAO et al., 2017). In the latter area, they have become the SOTA for several applications such as image-to-image translation, image inpainting, and image SR (MA et al., 2018; YU et al., 2018a; LEDIG et al., 2017).

However, since generator and discriminator need to achieve Nash equilibrium during training where neither generator nor discriminator can become too specialist in its task, GANs suffer from major challenges when training such as non-convergence, mode collapse, and diminished gradient (WANG; SHE; WARD, 2019). Consequently, they are highly sensitive to hyperparameters. In addition, for obtaining good results with them, their loss functions need to represent well the real optimization problem involved in the task (JOHNSON; ALAHI; FEI-FEI, 2016).

2.1.3 Coordinate Convolutions

The convolution operator is widely used in image processing, after learning the ideal filter weights, due to its ability to extract features of content from the training set that may not be in the same angle or place all the time. Such learned characteristic is called translation invariance. However, Liu et al. (2018) noted that also due to this feature, regular convolutions in CNNs could perform poorly in tasks that involve coordinate transforms. One example of this problem is the mapping between coordinates in (x, y) cartesian space to coordinates in the pixel space features, where even state-of-the-art architectures would not be able to obtain more than 90% of testing accuracy.

For dealing with problems that require varying degrees of translation dependence or complete translation invariance, [Liu et al. \(2018\)](#) proposed an operator called CoordConv, which works by giving the normal convolution operator access to its own input coordinates through the use of extra coordinate channels. This allows the network to check and work with the exact location of pixels inside its grid. This operator allows the network to learn either complete translation invariance or varying degrees of translation dependence, as required by position regression tasks. Their result in the same given position regression task presented perfect generalization, being 150 times faster, and having 10–100 times fewer parameters. The difference between a standard convolution operator to a CoordConv can be visualized in Figure 6.




Figure 6 – Differences between normal and coordinate convolutions. ([LIU et al., 2018](#))

Since their launch, researchers have explored different applications and scenarios where normal convolutions can be switched to CoordConv for improving performance ([UPADHYAY; SINGHAL; SINGH, 2019](#); [XU; CHEN; JIA, 2019](#)). Nonetheless, only [Zafeirouli et al. \(2019\)](#) so far in literature have reported the improvements that CoordConvs may provide over the use of regular convolutions for SR, which makes it an interesting research direction.

2.2 Super-Resolution

Super-resolution can be described as an attempt to generating a higher resolution image out of a lower resolution input. Throughout this domain, researchers have applied different strategies to reconstruct the HR image, which culminated in different classes of SR algorithms being developed depending on a variety of conditions ([HUANG; LIU, 2015](#)). Some of the categories involving SR are shown in Figure 7.

The general principle of supervised SR is that a LR image I_{LR} is the result of a degradation process that was applied to its HR version I_{HR} as in:

$$I_{LR} = D(I_{HR}) \quad (2.2)$$




Figure 7 – General classes of SR algorithms. (HUANG; LIU, 2015)

The degradation function D is naturally unknown, but researchers usually associate it with blur, motion, warp, and noise (NGUYEN et al., 2018). Therefore, the goal of the SR algorithm is to learn the inverse mapping in such a way that, from a LR input, its HR can be achieved as in:

$$I_{HR} = F_{SR}(I_{LR}; \theta) \quad (2.3)$$

where F_{SR} is the SR function and θ its parameters.

The most common and used techniques for upscaling images are the ones based on interpolation such as bicubic, bilinear, or nearest neighbor since their time cost is low, which makes them ideal for real-time applications. An illustration of the results when zooming an image (4x) with each technique is presented in Figure 8. Although the bicubic interpolation has a higher time complexity, it is the default method for upscaling images in software such as MATLAB and Photoshop. (PURKAIT; PAL; CHANDA, 2014; VEDADI; SHIRANI, 2014)

2.2.1 Operating Channels

The human evaluation of the degradation degree in a LR image is based on the perception of the RGB channel. However, when applying SR methods to images, some researchers instead use the YCbCr color space representation. In this space, images are depicted in Y, Cb, Cr channels, denoting the luminance, blue-difference, and red-difference chroma components,



Figure 8 – Visual comparison of general interpolation methods: (a) Nearest Neighbor (b) Bilinear (c) Bicubic (d) Original HD image. - (Source: Author's own)

respectively (WANG; CHEN; HOI, 2019). Some works report that using only the Y channel may bring better results than when working with the addition of Cb and Cr channels since they are more blurry than the Y channel by nature, and therefore are less affected by the downsampling process (DONG et al., 2015). There is no consensus in academia for which channels are better for training and evaluating SR; nevertheless, the most recent architectures tend to operate on RGB channels (LEDIG et al., 2017; CHEN et al., 2018).

2.2.2 Super-Resolution Benchmarking

Even though different works have presented several ways of benchmarking and measuring their image SR results regarding their specific field of application, the most common objective measurements of image quality are Peak Signal to Noise Ratio (PSNR) and Structural Similarity (SSIM). (TIAN; SUZUKI; KOIKE, 2010). PSNR is an estimation of quality based on the mean squared error (MSE) of pixels for every channel between the HR image generated and the ground truth, as can be seen in Equations 2.4 e 2.5.

$$PSNR = 10 \log_{10} \left(\frac{S^2}{MSE} \right) \quad (2.4)$$

$$MSE = \frac{\sum_{n,m} (x_{mn} - y_{mn})^2}{m * n} \quad (2.5)$$

where: S is the maximum value in the input image data type; n is the number of pixels; m the number of channels; x_{mn} and y_{mn} represent the pixel value described in n with the channel m for the generated and original images respectively.

SSIM is a measurement that considers the visual degradation in quality with more importance through analysis of the homogeneity and phase coherence of the gradient magnitude on the original and reconstructed image. This similarity is based on structure, brightness, and contrast

of the images. (BEGIN; FERRIE, 2006; REIBMAN; BELL; GRAY, 2006) Its mathematical formulation can be seen in Equation 2.6.

$$SSIM = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (2.6)$$

where: μ_x and μ_y represent the average intensity value of a linked windows for the original and reconstructed image; c_1 and c_2 denote the brightness of two images; σ_x and σ_y formulate the variance of the two sets of intensity for both images; σ_{xy} presents the correlation between these two sets.

2.2.3 Deep Learning for Image Super-Resolution

Deep learning solutions for SR fits into the “learning-based” category show in Figure 7. In the last few years, DL methods have become the most explored approach for performing SR tasks since they early showed SOTA performance in various benchmarks and competitions (AGUSTSSON; TIMOFTE, 2017; TIMOFTE et al., 2018). In special, the single image super-resolution (SISR) problem has been the most fundamentally tackled problem within SR, since researchers can make use of already available large datasets scrapped from the internet to train their models (LIU et al., 2015; CHEN et al., 2018).

A variety of methods have been used and incorporated for solving the SR problem, ranging from simpler approaches involving only convolutional layers, to more sophisticated ones with the use of residual blocks, recursive learning and different losses (WANG; CHEN; HOI, 2019). An overview of the most related directions that researchers have taken when considering working with DL in SR can be analyzed in Figure 9.




Figure 9 – DL for SR algorithms related topics. (Adapted from Wang, Chen and Hoi (2019))

For proposing the new architectures assessed in this thesis, different network designs and learning strategies presented in Figure 9 were considered, such as the use of “Residual Learning”

and “Content Loss”. A more deep review of the works which influenced the directions taken in this manuscript is presented in Section 3.

2.3 Face Recognition

The basic steps that involve a general face recognition pipeline are defined in Figure 10 and described in sequence.




Figure 10 – Example of a generic pipeline for face recognition. (Source: Author’s own)

2.3.1 Face Detection

Face detection in an image is the first step in a recognition pipeline because it eliminates unnecessary information from the image. In this way, if the algorithm finds one or more faces, they are extracted from the original image so that they can be analyzed separately (MUTTU; VIRANI, 2015).

The training phase of these algorithms happens with the use of several images containing faces and others without them. Even though this problem presents itself as a simple binary classification, several face detection algorithms need to be trained exhaustively so that they can give good results (ZHANG et al., 2016). Two measures are responsible for evaluating the quality of face detection algorithms (VEZHNEVETS, 2002):

- False positive: Represents the number of objects that were detected wrongly as faces.
- False negative: Represents the number of faces that were not detected.

Face detection algorithms are usually divided into four different groups: knowledge, feature, template, and appearance-based models (ZAFEIRIOU; ZHANG; ZHANG, 2015). However, as the amount of available data has increased over the years for training such algorithms, the appearance-based methods have overcome the other solutions since they generalize face

models from a set of representative samples. A common core for the SOTA algorithms proposed in this group of techniques is the use of CNNs since they derive problem-specific feature extractors from the training examples automatically, without making any assumptions about the features to extract or the areas of the face patterns to analyze due to their spatially invariant characteristic (ZHANG; ZHANG, 2010).

2.3.2 Feature Extraction and Face Verification

The “real” recognition step in a face recognition pipeline consists of the representation and extraction of facial features of an image. These features are then input into a mathematical model, which is meant to specify whether the presented face matches one or any previously stored face (CROSSWHITE et al., 2018).

The implementation of recognition systems can range from low-throughput to process-intensive methods where, for example, GPUs are required. Some more straightforward methods can make use of metric learning approaches or principal component analysis for dimensionality reduction. On the other hand, the most sophisticated ones are usually based on analysis of probability densities, manifold learning, and deep neural networks, among other methods with a higher computational cost (WANG; DENG, 2018).

For extracting discriminative features of an image that only contains a face (after the pre-processing step), models based on CNN have been the ones most used by SOTA approaches. This architecture is suitable for feature extraction because it takes advantage of local connections to extract the spatial information effectively. Also, their shared weights significantly reduce the number of parameters for training the network, which consequently reduces its size (CHEN et al., 2016). An effective way to create accurate face recognition models is through the application of Transfer Learning (LECUN; BENGIO; HINTON, 2015) using available pre-trained models. These models are often trained in datasets with millions of faces and, through the use of their intern representations, it is possible to extract discriminative features of an input face directly (CAO et al., 2018).

Once extracted all the features of the involved subjects, the system needs to decide whether the person is whom he/she claims to be. This step is called face verification and different machine learning approaches can be employed to perform it depending on how many dimensions the obtained feature space may have (FACELI et al., 2011). These approaches can be differentiated by how their functions create the decision boundaries on the feature hyperplane. However, in the context of face recognition, when only one or few training samples are provided, methods based on distance metrics have shown the best results regarding computational complexity and accuracy (NGUYEN; BAI, 2010; SCHROFF; KALENICHENKO; PHILBIN, 2015).

2.4 Final Considerations

In this chapter, an overview of the main topics discussed in this thesis was provided. It is important to reinforce that most of the trends regarding DL in the fields of face recognition and super-resolution have only emerged in the past five years through empirical experimentation with different architectures. This statement indicates that most of the theory behind why these models have performed better than others is still in the development phase and will probably lead in more exploration and changes in the following years.

In the next chapter, different SOTA works with respect to SR and LR face recognition are discussed. Their evaluation was essential to extract meaningful insights for proposing the hypotheses and objectives of this thesis.

3

Related Work

This chapter presents some of the related works evaluated during the development of this thesis.

3.1 Super-Resolution

[Baker and Kanade \(2000\)](#) proposed the first SR work to be applied to faces in 2000. They created an algorithm that was used to learn priors on the spatial distribution of the image gradient for frontal images of faces. At that time, they stated that the high-frequency details inferred by the probabilistic models were “hallucinated” by the model.

The work of [Tian, Suzuki and Koike \(2010\)](#) presented objective and subjective measures for evaluating how SR impacts different image processing and computer vision tasks. Their findings reflected the conflicts between objective and subjective measures since the former tends to penalize the model that enhanced the image according to computer vision standards, and the latter tends more to changes that improve the image quality based on the human vision system.

[Dong et al. \(2015\)](#) were the first to propose the use of CNNs for the SR problem. Their architecture was called “Super-Resolution Convolutional Neural Network” (SRCNN) and provided superior accuracy compared with other SOTA example-based methods at the time. In their work, LR images are pre-upsampled using traditional methods (e.g., bicubic interpolation) to the desired size, and then a deep CNN with three layers is applied to the coarse image for reconstructing the high-frequency details. This work became later the baseline for all works that involve DL based algorithms in SR. One advantage of this method (and all pre-upsampling methods) is that they can take input images of any arbitrary size and perform the SR task. The downsides may be the introduction of noise and blurring and, since most operations are performed with images in a high-dimensional space, time and memory costs can be higher than other frameworks.

[Dong, Loy and Tang \(2016\)](#) designed a compact hourglass-shape CNN structure using the basic SRCNN structure for faster inference and improved accuracy in SR called FSRCNN. They proposed an architecture with a deconvolution layer at the end of the network for mapping the original LR image directly to the super-resolved output, an iterative up-and-down sampling in the mapping layers, and the use of smaller filter sizes with more mapping layers. The results pointed out an increase in performance of over 40x for inference time while presenting a superior restoration quality when compared against the naive SRCNN architecture.

For the work of [Shi et al. \(2016\)](#), the authors presented a strategy to solve the necessity to upscale the LR with interpolation methods or using a single filter before feature extracting and mapping. They presented a modified CNN architecture with an efficient sub-pixel convolutional layer for “post-upsampling” where the feature extraction could happen in the LR space before being upsampled. This architecture was capable of performing real-time SR in 1080p videos on a single K2 GPU.

In the work of [Ledig et al. \(2017\)](#), a deep generative adversarial network using residual convolutional blocks was applied for image SR. Their approach achieved SOTA results in upscaling photo-realistic natural images by a factor of 4. To accomplish such results, instead of only optimizing the network by image similarity in pixel space, the authors proposed a perceptual loss function, which consisted of an adversarial loss and a content loss. The adversarial loss was responsible for pushing the upsampled solution to the natural image manifold using a discriminator network trained to differentiate between the super-resolved images and the original photo-realistic ones. Besides, they proposed the use of a content loss motivated by perceptual similarity. This similarity was calculated from the comparison of extracted semantic features from an ImageNet pre-trained network. Also, they evaluated the impact of applying several image losses together, such as adversarial, content, mean-squared-error, and total-variance, which inspired this thesis in investigating a different task-specific learning strategy.

The work of [Haris, Shakhnarovich and Ukita \(2018\)](#) presented an approach to detect objects in LR images using an end-to-end strategy with the training of a CNN to perform the SR steps, and also aid detection. In this approach, a specific multi-objective loss function was developed for CNN training, where individual weights for each part of the loss were used in order to optimize the learning process based on each desired task. The goal behind the work was to assess how much improvement in resolution would assist a recognition/detection task in the input image.

[Chen et al. \(2018\)](#) presented an end-to-end approach to perform SR on face images using prior geometric face features as prior information. The authors divided the training process into several stages where different encoder-decoder network architectures were applied to extract geometric features from the faces to aid the task of SR. The deep network produced in this work, FSRNet, presented results that today are the SOTA for SR in face images. However, when dealing with real-world SR of LR face images in the wild, obtaining priors is a hard and computationally

expensive task that hinders its implementation in the face biometric context.

Zafeirolouli et al. (2019) proposed an efficient, lightweight model leveraged by the benefits of a recursive progressive upsampling architecture to tackle the SR problem. This work recognized that SR tasks involve spatial representations and transformations, and exploited the pixel position information to reinforce the reconstruction task using the CoordConv operator. They obtained comparative results with SOTA implementations in four SR benchmarks. More importantly, their results also showed accuracy performance improvements with the use of the coordinate convolutional layer for the SR task while keeping low computational complexity, which motivated the application of this operator for proposing the new architectures described in Chapter 5.

3.2 Low-Resolution Face Recognition

Hennings-Yeomans, Baker and Kumar (2008) presented an approach for simultaneous SR and face feature extraction for recognition of LR faces by treating face features (e.g., Eigenfaces, Fisherfaces) as prior information in the SR method. They evaluated their approach against matching gallery and probe images in the LR and applying the pure SR approach to check for matches in the high-dimensional domain. They concluded that their approach could produce better recognition performance since the focus of the SR shifted to recognition instead of reconstruction. This particular work inspired this thesis for using features extracted from the face as an optimization strategy for the SR models.

The work of Rasti et al. (2016) proposed a system that super-resolves a face image before the face feature extraction and recognition phases. They used a deep CNN to upscale the image followed by a Hidden Markov Model and Single Value Decomposition based face recognition model. They experimented in two general and one small surveillance database and pointed out that such upscaling phase could result in a 6 to 10% increase in performance for face recognition. The increase in accuracy performance reported in this paper influenced the elaboration of the hypothesis that DL based SR could assist positively face recognition in real-world LR data.

Berger, Peyrard and Baccouche (2016) proposed a two-step neural approach for face SR with the focus of improving face recognition. They employed a generic SR CNN network based on the work of Peyrard, Mamalet and Garcia (2015) trained on the Labeled Faces in the Wild (LFW) dataset (HUANG et al., 2008) and then refined the HR output with localized SR steps using autoencoders trained in patches of the images on the LFW. The localized SR step focused on locally reconstructing image patches at crucial face landmark points (e.g., eyes, nose, mouth) via dictionary learning. They claimed that the image reconstruction had a +2.80dB improvement, while the recognition performance also had a 3.94% increase compared to the same results on x4 bicubic interpolation. However, they lacked tests in real-world LR datasets in order to check if their model would be able to keep the high performance within the wild data. Also, as having

two networks in cascade is computationally expensive, this CNN architecture may not be ideal for real-world surveillance situations.

Wang et al. (2016) presented an attempt to deal with the problem of very low-resolution recognition, where the region of interest could be smaller than 16x16 pixels. Their approach achieves feature enhancement and recognition simultaneously through the use of a deep SR network for pre-training with a carefully selected loss function for matching between LR and HR face images. The recognition step employed a deep neural network for classification trained on a different dataset that shared similar features with the one used for evaluation. They report a rise of 1.71% in top-1 accuracy on the famous UCCS surveillance dataset, which is no longer publicly available.

Aghdam et al. (2019) explored factors for improving LR face recognition using DL classifiers in two real-world surveillance datasets. Instead of focusing on SR approaches, they proposed two strategies to overcome the lack of fine information on the face images: increase the crop on a detected face before upsampling the image to match the input size of the classifier and match the resolution between gallery and probe images. For classification, they evaluated several ResNet-50 and SENet-50 architectures for feature extraction trained on VGGFace2 and the MS-Celeb-1M datasets. Along with a nearest neighbor classifier, they were able to achieve SOTA results in Rank-1 verification for the ICB-RW and SCFace surveillance datasets. Their work inspired this thesis in also investigating different face crop sizes and using the ICB-RW database as benchmark for LR face recognition.

ElSayed et al. (2018) evaluated the effects that SR and face alignment may have on accuracy for LR face recognition using an unsupervised approach. They proposed experiments where a LR version of the LFW dataset was frontalized and fed to a simple SR network based on SRCNN. Later they made use of an unsupervised recognition model using speed up robust features and local binary features. They tested only on the LFW data and reported that SR and face alignment increased the recognition performance.

Ataer-Cansizoglu et al. (2019) proposed a two-stage architecture for simultaneous feature extraction and super-resolution. They trained a VGG-based deep face recognition network to be used as a feature extractor and trained an SR network to decrease the L1 distance between the features extracted from the VGG network for real and generated images. The evaluation procedure presented two DL based SR networks and showed that this setup increases recognition performance. However, they only evaluated the results in LR frontal images that were acquired after the downsampling of four HR datasets.

Li et al. (2019) presented results for LR face recognition in the wild with the evaluation of different deep learning SR network architectures in two originally LR datasets. They trained a VGG network using LR versions and HR versions of images from a HR dataset to extract features and then applied different classifiers for validation. Their best results were obtained by pre-training an SR architecture based on GAN in LR images of datasets with similar features to

the ones used for evaluation. This trick helped their models reach close to SOTA recognition performance in the used datasets.

[Abello and Jr. \(2019\)](#) explored the use of a loss function defined as the L2 error between face features from an super-resolved face and the ground truth for improving LR face recognition. For feature extraction, they used the pre-trained network with Inception architecture proposed by [Schroff, Kalenichenko and Philbin \(2015\)](#). They reported that this loss was able to give improvements for both image quality and recognition performance. Nevertheless, they only tested their system in LR versions of the HR dataset used for training.

3.3 Final Considerations

In this chapter, several works that present the SOTA for SR and LR face recognition are described with their results. After analyzing what the trends present in the SOTA were, several architectures for SR were chosen for evaluation in real-world LR images. Also, the discussed related works led this thesis to a research direction that included the assessment of the use of different convolution operators and loss functions for better image quality and recognition accuracy.

For the next chapter, the methodology behind the experiments proposed in this thesis for assessing the hypotheses established in Chapter 1 are presented.

4

Methodology

This chapter describes the materials and methods used for experimenting with different architectures for super-resolving images and improving face recognition.

4.1 Datasets

4.1.1 VGGFace2

The VGGFace2 is an in the wild dataset of faces that contains 3.31 million images from 9131 celebrities downloaded from Google Image Search and shows significant variations in pose, age, lighting, and background ([CAO et al., 2018](#)). One advantage of using this dataset for training a robust image classifier/feature extractor is the fact that approximately 20% of its images have pixel resolution lower than 50 pixels, which leads the model to have a better feature representation for low-resolution face images ([AGHDAM et al., 2019](#)). A set of five images from this dataset can be seen in Figure 11.



Figure 11 – Example of face images in the VGGFace2 dataset. (Adapted from [Cao et al. \(2018\)](#))

4.1.2 CelebA

The CelebFaces Attributes Dataset (CelebA) is a large-scale face attributes dataset with over 200,000 images of 10,117 celebrities around the world ([LIU et al., 2015](#)). It presents a

vast diversity of poses and background clutter. Also, it is commonly used for training SOTA SR networks because of its rich features and size (CHEN et al., 2018; YU et al., 2018b; KIM et al., 2019). The first 18,000 images of this dataset were used for training, and the following 2,000 were used for validating the results for the SR networks according to the specified image quality metrics. A set of 5 samples from this dataset is shown in Figure 12



Figure 12 – Example of face images in the CelebA dataset. (Adapted from Liu et al. (2015))

4.1.3 Quis-Campi Dataset (ICB-RW)

The Quis-Campi dataset is a growing biometric database comprising 3000 images from 320 subjects automatically acquired by an outdoor visual surveillance system, with subjects on-the-move and at-a-distance (up to 50 m). The system used for image acquisition has a master wide camera for subject detection and tracking, and a slave pan-tilt-zoom (PTZ) camera, as the foveal sensor, for extracting the facial region at a high-magnification state.

In the context of face verification, they supply three high-quality images of the subject in a controlled environment to be used as gallery data and several images of the same subject on the move inside an university campus to be used as probe data. One strong feature of this dataset is that all probe images present variation in illumination, pose, focus, expression, motion-blur, and occlusion (NEVES; MORENO; PROENÇA, 2017).

Part of this dataset was published to promote the International Challenge on Biometric Recognition in the Wild (ICB-RW) competition, and that is why most of the results present the same benchmarking setup used in the competition. The ICB-RW challenge provided three face images to be used as gallery and five probe images for each of 90 subjects.

As the goal was to evaluate the performance of the proposed network architectures against other works in a real-word LR scenario, the Quis-Campi dataset was adapted to resemble the ICB-RW challenge to the maximum once the latter was not available. Therefore, since not all the subjects on the dataset had enough images to be selected, the first 90 subjects out of the 320 that had available the three gallery and five probe images were picked. This approach to making an equivalent representation of ICB-RW was considered instead of taking 90 random samples out of the 320 since, at the time of the 2016 competition, the Quis-Campi dataset did not have all of their subjects registered, and according to Neves and Proença (2016), the new samples were

registered and added automatically to the database. The images chosen for one of the subjects can be seen in Figures 13 and 14.



Figure 13 – Example of gallery images in the Quis-Campi dataset. (Adapted from [Neves, Moreno and Proença \(2017\)](#))




Figure 14 – Example of probe images in the Quis-Campi dataset. (Adapted from [Neves, Moreno and Proença \(2017\)](#))

4.1.4 Federal University of Sergipe Classroom Attendance Dataset

This dataset was formulated with the goal of creating an automated attendance system for classes within the computer science department of the Federal University of Sergipe (UFS) ([SÁ, 2019](#)). The dataset is composed of one high-resolution frontal image of each student, referred to as a gallery image, and three probe images of a whole class taken in slightly different angles with a 1.2MP Webcam. For this thesis, three classes with different amount of students were used for the evaluation since this dataset presents a challenging LR uncontrolled environment ideal for testing real-world face recognition pipelines. An example of the gallery and probe images present on the dataset can be seen in Figures 15 and 16, respectively.

4.2 Data Pre-Processing

For detecting and extracting the faces from the presented datasets, a pre-trained Multi-task Cascaded Convolutional Neural Network (MTCNN) was used since it has shown SOTA results on a variety of benchmarks for face detection and face alignment while keeping real-time results ([ZHANG et al., 2016](#)).



Figure 15 – Example of gallery images for the UFS-Classroom Attendance dataset. (SÁ, 2019)



Figure 16 – Example of a probe image for the UFS-Classroom Attendance dataset. (SÁ, 2019)

For training and evaluating the SR networks, every image from the CelebA dataset was scaled to a $[0,1]$ range, and then underwent a process of “crappification” to create an LR pair to be used since a paired supervised learning approach was adopted. Each cropped face was resized to 160×160 and saved as the HR sample. Then, for obtaining a “crappy” version of it, the same cropped face was resized to 40×40 pixels, and saved using JPEG compression with a quality factor that varied randomly from 10 to 70 (where 1 is the minimum, 75 the standard, and 95 the maximum quality) as advised by Howard et al. (2018) during the *FastAI* course. This compression approach helps to simulate the data distribution that may be present in real-world LR surveillance footage. The resolution was chosen according to the input size of the deep learning model used for feature extraction ($160 \times 160 \times 3$), and due to limited computational resources, only the $4\times$ SR upscaling setting was experimented.

As the primary task was to evaluate how SR may influence verification performance in

real-world LR in-the-wild scenarios and faces detected in the Quis-Campi and UFS-Classroom Attendance datasets had a large variation in size due to differences in data acquisition equipment, every face detected (gallery and probe) was saved as an image in three different settings: without any change to the size, with bicubic resizing to 40x40 pixels, and with bicubic resizing to 40x40 pixels with a 1.3 crop margin to the borders. This last setup is able to increase the amount of information within the image and, added to the employed resolution matching step, can increase recognition performance for LR samples as validated by [Aghdam et al. \(2019\)](#). An example of the three cases for each dataset can be seen in Figures 17 and 18.




Figure 17 – Example of a probe face image from the UFS-Classroom Attendance dataset saved in the three settings. (Adapted from [Sá \(2019\)](#))




Figure 18 – Example of a probe face image from the Quis-Campi dataset saved in the three settings. (Adapted from [Neves, Moreno and Proença \(2017\)](#))

4.3 Transfer Learning

4.3.1 Face Feature Extraction

Inspired by the work of [Schroff, Kalenichenko and Philbin \(2015\)](#), a pre-trained network to extract feature embeddings of the faces for further comparison was employed. The chosen deep network architecture was the Inception-ResNet-V1 ([SZEGEDY et al., 2017](#)) trained on the VGGFace2 dataset. This network was able to achieve the SOTA accuracy of 0.9965 on the LFW benchmark. Compared to the Inception network architecture employed similarly in the Facenet paper, the Inception-ResNet-V1 network achieves faster convergence without adding additional

computation complexity due to its residual connections. This network was trained for mapping an $160 \times 160 \times 3$ image ($\mathbb{R}^{H \times W \times C}$) to a vector ($\phi(\text{Image})$) in a feature space of 512×1 dimensions ($\mathbb{R}^{512}$).

4.3.2 Face Verification

The verification is performed by applying the nearest neighbor algorithm to check the distance among embeddings for the selected probe and gallery images. The metric employed for verification of closeness is the cosine similarity used previously by the winner of the ICB-RW competition (NEVES; PROENÇA, 2016), and also by Aghdam et al. (2019). A description of the metric can be seen in Equation 4.1.

$$\text{Nearest Neighbor} = 1 - \frac{\phi(I_{Face1}) \cdot \phi(I_{Face2})}{\|\phi(I_{Face1})\|_2 \|\phi(I_{Face2})\|_2} \quad (4.1)$$

Most approaches that deal with the LR face recognition problem, as it was reviewed in Chapter 3, try to train from scratch a robust network that may be able to overcome the difficulties of such task. However, this study approaches the problem from a different point of view since it takes advantage of pre-trained large and robust classifiers with the addition of a specially designed upscaling step including an SR network previous to the feature extraction and verification steps.

A general description of the whole pipeline for verification after the upsampling task, using the ICB-RW challenge as an example, can be seen in Figure 19.




Figure 19 – Pipeline for face verification in the ICB-RW. (Source: Author's own)

4.3.3 Face Loss

Mean-squared-error loss optimizes the response of an SR network for generating images with better quality, but it does not take into account if the recognized person has kept the same unique features that may differentiate this person from the others. As a task-driven approach for SR was meant to be developed, the face identity loss commonly used for face normalization (COLE et al., 2017) and 3D face reconstruction (GECER et al., 2019) was adopted to guide the SR process for better face recognition accuracy.

The SOTA face feature extractor (Inception-ResNet-V1) was used for checking if the distance between the embedding of the real and the super-resolved image was decreasing within each epoch. To accomplish that, the cosine similarity measure (Equation 4.1) of the face embeddings extracted from the image pairs was added to the standard image loss used for training each SR network. This customized loss ensures that the reconstruction made by the SR network resembles the target identity under various conditions in the feature space. The definition for this loss can be seen in Equation 4.2.

$$FaceLoss = 1 - \frac{\phi(I^{SR}) \cdot \phi(I^{HR})}{\|\phi(I^{SR})\|_2 \|\phi(I^{HR})\|_2} \quad (4.2)$$

This approach adopts the same concept of the task-driven loss presented by Haris, Shakhnarovich and Ukita (2018), but it kept the focus on a more robust and general task (face recognition). In addition, the Face Loss presented here differs from the recent work of Abello and Jr. (2019) since they applied L2 error, which is more susceptible to outliers, on the feature vectors for the optimization of their SR network. Moreover, in the end, they neither provided an ablation study nor validated their approach for accuracy improvement in real-world LR data.

4.4 Final Considerations

This chapter described the datasets and the steps taken towards finding the best DL architecture for super-resolving images with the goal of improving face recognition. To ensure the models work in real-world LR data, two challenging LR datasets (surveillance and attendance assessment) were used for evaluation. Also, for making sure the super-resolved face images have discriminant features regarding identity, a custom loss function was proposed.

For the next chapter, the experiments of this thesis are presented and explained. Then, their results are broadly discussed according to the hypotheses idealized in Chapter 1.

5

Experimental Results

5.1 Experiments

This section provides an overview of the performed experiments with their respective parameterization.

5.1.1 Task 1 - Face Super-Resolution

For performing this task, 4 of the CNN architectures discussed in Section 3 were modified and implemented following the implementation details on their papers in order to evaluate the best choice for a face verification pipeline. These models were chosen based on previously reported results and computational complexity. In this thesis, 7 different new architectures were proposed for evaluation. The variants that have the name “Coord” kept the same original architecture but had the first “Conv2d” layer switched to a “CoordConv”. The variants with “FaceLoss” had their training with the addition of the customized loss function presented in Section 4.3.3. The implemented models for SR described below can be checked on https://github.com/angelomenezes/Pytorch_Face_SR.

Network architectures from literature that were evaluated:

1. SRCNN ([Dong et al. \(2015\)](#)) shown in Figure 20.
2. Subpixel CNN ([Shi et al. \(2016\)](#)) shown in Figure 21.
3. FSRCNN ([Dong, Loy and Tang \(2016\)](#)) shown in Figure 22.
4. SRGAN ([Ledig et al. \(2017\)](#)) shown in Figure 23.




Figure 20 – SRCNN architecture. (DONG; LOY; TANG, 2016)




Figure 21 – Subpixel CNN architecture. (Adapted from Dong, Loy and Tang (2016) and Shi et al. (2016))




Figure 22 – FSRCNN architecture. (DONG; LOY; TANG, 2016)




Figure 23 – SRGAN architecture. (JIAO; ZHAO, 2019)

Network architectures with modifications proposed in this thesis for evaluation:

1. SRCNN Coord
2. Subpixel CNN Coord
3. FSRCNN Coord
4. FSRCNN Coord FaceLoss
5. SRGAN Coord
6. SRGAN FaceLoss
7. SRGAN Coord FaceLoss

- **Hyperparameters** → SRCNN and SRCNN Coord

- Batch size: 64
- Number of epochs: 50
- Loss: MSE
- Optimizer: Adam
- Learning Rate: 0.01 with decay to 10% of the current learning rate every 15 steps

- **Hyperparameters** → Subpixel CNN and Subpixel CNN Coord

- Batch size: 32
- Number of epochs: 50
- Loss: MSE
- Optimizer: Adam
- Learning Rate: 0.01 with decay to 20% of the current learning rate every 15 steps

- **Hyperparameters** → FSRCNN, FSRCNN Coord and FSRCNN Coord FaceLoss

- Batch size: 32
- Number of epochs: 50 (30 for FSRCNN Coord FaceLoss)
- Loss: MSE and variant that added FaceLoss
- Optimizer: Adam
- Learning Rate: 0.001 with decay to 20% of the current learning rate every 15 steps

- **Hyperparameters** → SRGAN, SRGAN FaceLoss, SRGAN Coord and SRGAN Coord FaceLoss
 - Batch size: 32
 - Number of epochs: 30
 - Loss: MSE + Adversarial Loss + Perceptual Loss and variant that added FaceLoss
 - Optimizer: Adam with $\beta_1 = 0.5$ and $\beta_2 = 0.999$
 - Learning Rate: 0.001 with decay to 20% of the current learning rate every 15 steps

5.1.2 Task 2 - Watch-List ICB-RW (1x5 Problem)

As proposed by [Neves and Proença \(2016\)](#), the ICB-RW challenge used part of the Quis-Campi dataset to evaluate the average Rank-1 identification of a suspect against the “watch-list” subjects. For each probe image, the model had to output a similarity score related to each of five possible suspects. An example of this setup is presented in Figure 24.




Figure 24 – Watch-List setting for the ICB-RW. ([NEVES; PROENÇA, 2016](#))

For this experiment, each individual had its frontal gallery image and a random probe image selected along with four random probe images of different subjects. The challenge is to obtain the highest number of matches according to the smallest distance given by the nearest neighbor algorithm.

5.1.3 Task 3 - Attendance Evaluation (1xN Problem)

For the task of evaluating the attendance inside a classroom, the identity of every student needs to be checked against all entries on the attendance list. The number of students in each class can be seen in Table 1.

Table 1 – Number of students for each class

	N° of Students
Class 1	15
Class 2	16
Class 3	12

This experiment followed the same verification principle of Task 2. Nevertheless, it is a more challenging situation since this recognition task takes into account all the subjects in the classroom (1 vs. ALL problem).

Each experiment described in this section was either run in a personal computer with an Intel i7-6500U with 16 GB of memory and GeForce GTX 950M (4 GB) or in a Google Cloud instance with a Skylake processor (8 vCPUs and 52 GB memory) with an NVIDIA Tesla P100 (16 GB). All models were implemented and evaluated using Python and the Pytorch library.

5.2 Results Evaluation

When training the SRGAN and its variants, if the discriminator network had its weights updated in the same frequency of the generator, its loss would quickly converge to zero, and both networks would not have any gradients for learning along the epochs. Therefore, in order to make the learning happen, two extra steps were followed according to the tips given on the *FastAI* course (HOWARD et al., 2018):

- **Generator Pre-Train:** the generator network was pre-trained for 5 epochs using only the MSE loss in order to have some advantage initially against the discriminator.
- **Smart update for the Discriminator:** across the epochs, the discriminator network was only trained (had its weights updated) when its loss was above a threshold (0.5). This step ensures that the network is learning gradually to assess the output of the generator since the update of weights for the discriminant only happened when it was making more “mistakes” according to the right labels for each input. All the training losses can be evaluated in Appendix B.

The results for each proposed task are shown and discussed in the following subsections.

5.2.1 Task 1 - Face Super-Resolution

The validation results regarding image quality metrics and inference results for the SR architectures are presented in Table 2. The PSNR was calculated on the RGB channels, and the average time for inference was computed as the average of 10 runs of each algorithm. The perceptual results can be evaluated in the Appendix A.

As can be seen in Table 2, except for the “SRGAN Coord FaceLoss”, all the presented architectures achieve real-time performance in a small GPU (GeForce GTX 950M). The best algorithm regarding quality metrics was the FSRCNN with the coordinated convolution operator. Nonetheless, the SRGAN and its variants, even with a low PSNR, had the best human perceptual quality as can be seen by the perceptual clarity of the outputs in the Appendix A. Their better performance may be related to the fact that they make use of different losses, which optimizes for a less blur and more textured output (Adversarial and Perceptual Loss).

Table 2 – PSNR and SSIM validation results (2000 images from CelebA)

Validation metric results for 4x Upscaling					
	PSNR	SSIM	Avg. Time for Inference (s)	Inference FPS in GPU	Trained on Channel
SRCNN	27.95	0.7973	0.0100	100.00	Y
SRCNN Coord	27.98	0.7966	0.0153	65.36	Y
SubCNN	28.08	0.8003	0.0035	285.71	Y
SubCNN Coord	28.13	0.8022	0.0048	208.33	Y
FSRCNN	28.45	0.8104	0.0046	217.39	RGB
FSRCNN Coord	28.88	0.8175	0.0048	208.33	RGB
FSRCNN Coord Face Loss	28.78	0.8151	0.0047	212.77	RGB
SRGAN	27.02	0.8077	0.0336	29.76	RGB
SRGAN Coord	27.28	0.8078	0.0382	26.18	RGB
SRGAN FaceLoss	26.63	0.8083	0.0384	26.04	RGB
SRGAN Coord FaceLoss	26.69	0.7984	0.0404	24.75	RGB
Bicubic	27.93	0.7881	0.0012	833.33	-

→ **Red color** highlights the architectures proposed in this thesis.

Architectures that were modified with the CoordConv operator improved its PSNR 100% of the times with an average increase of 0.16 dB. Their SSIM presented fluctuations, and no significative gains could be measured. This situation may be explained by the fact that all architectures were optimized to decrease the difference in pixels according to the MSE loss, which directly improves the PSNR of the image but often smooths and blurs the output. Such blur decreases the perceptual quality of the image, and consequently, does the same to perceptual quality metrics such as the SSIM.

The SR models that were optimized using the customized loss function (FaceLoss) presented in general lower image quality metrics than their pairs that used their own objective function. They also presented a longer time for inference, which may be an indication that fewer weights were zero since more computation was measured. Despite that, every evaluated

architecture was able to reach real-time inference in the small GPU used for training and testing the models. The interpolation method (Bicubic) was still around three times faster than the fastest SR model. This can indicate that, for situations where processing time weights more than accuracy, interpolation methods can still be a considered direction.

5.2.2 Task 2 - Watch-List ICB-RW (1x5 Problem)

The results for the task of recognizing which suspect was correctly identified by the surveillance camera can be seen in Table 3. Fine-tuning was not performed. Therefore, the recognition pipeline did not have direct samples that resembled the gallery/probe data of this experiment.

The best obtained results came from SRGAN and its variants. In the setting where the face image was not resized previously, only these algorithms were able to overcome the baseline (Bicubic). When the cropped face had its resolution reduced for simulating lower resolution scenarios, they still were the best performing group, but other networks were also able to beat the interpolation method.

Table 3 – Rank-1 Accuracy (in %) for Face Recognition task (1xN) on 90 subjects from the Quis-Campi dataset (ICB-RW)

Results Quis-Campi / ICB-RW (N = 5 suspects).			
	No Resize No Margin	Size 40 No Margin	Size 40 Margin 1.3
SRCNN	74.89	61.78	64.67
SRCNN Coord	78.22	62.89	67.11
SubCNN	64.67	55.33	57.56
SubCNN Coord	67.33	58.44	62.22
FSRCNN	72.00	60.22	65.56
FSRCNN Coord	78.22	64.00	69.11
FSRCNN Coord Face Loss	78.44	63.33	67.33
SRGAN	85.78	77.33	72.00
SRGAN Coord	83.80	78.40	68.90
SRGAN FaceLoss	85.11	78.22	71.78
SRGAN Coord FaceLoss	84.89	76.00	71.11
Bicubic	83.11	62.00	62.44

→ Red color highlights the architectures proposed in this thesis.

As the resolution of the faces on the probe data for the ICB-RW had naturally almost the same resolution as the ones for the gallery (around 200x200), the 1.3 margin did not have much effect on the accuracy results, which was also noticed in the work of [Aghdam et al. \(2019\)](#). As in this experiment the sizes of the gallery and probe data were always matched before upsampling, the 40x40 resizing may have caused a drastic loss of high-frequency details and discriminative features, which can be noticed by the decrease in accuracy ratings.

Every accuracy result on this task, except for the SubCNN and its variant, overcame the results of Ghaleb et al. (2018), the best performing system in the ICB-RW challenge at that time. He was able to achieve a Rank-1 IR rate of 71.7%, which differs from the proposed SRGAN and SRGAN FaceLoss by a margin of 17.08% and 16.41%, respectively. These two architectures would also beat the results of Aghdam et al. (2019), who was able to achieve a Rank-1 rate of 84.22%, the highest registered so far.

5.2.3 Task 3 - Attendance Evaluation (1xN Problem)

The results for the task of evaluating which students are present in each of the classrooms can be seen in Tables 4, 5, and 6. Similarly to the other task, neither training nor fine-tuning was performed using gallery/probe data for this experiment.

For Classroom 1, the best performing algorithm for all settings was the SRGAN and its variants. The FSRCNN model presented the highest obtained accuracy, but it did not keep a performance consistency. For the setting where no resize was employed, all architectures were able to beat the baseline. Yet, when the margin was applied to increase the amount of information within the image, bicubic interpolation overcame even the SRGAN and two of its variants.

Table 4 – Accuracy (in %) for Face Recognition task in Classroom 1 (1xN)

Results UFS Classroom 1 (N = 15 students)			
	No Resize No Margin	Size 40 No Margin	Size 40 Margin 1.3
SRCNN	62.50	66.67	83.33
SRCNN Coord	64.58	70.83	85.42
SubCNN	68.75	64.58	66.67
SubCNN Coord	70.83	66.67	68.75
FSRCNN	66.67	72.92	83.33
FSRCNN Coord	66.67	66.67	89.58
FSRCNN Coord Face Loss	64.58	68.75	89.58
SRGAN	85.42	81.25	75.00
SRGAN Coord	81.25	79.17	81.25
SRGAN FaceLoss	85.42	83.33	83.33
SRGAN Coord FaceLoss	81.25	81.25	79.17
Bicubic	58.33	66.67	83.33

→ **Red color** highlights the architectures proposed in this thesis.

For Classroom 2, it was possible to conclude that the images had a high degree of degradation since the highest accuracy was around 73%, and the setting without resizing and margin adjustment had around 55%. The best obtained results came from the SRGAN with coordinate convolution in the setting where the margin was adjusted. All the other SRGAN related models happened to hit the same accuracy, which gives a hint that they might have similar weights.

For Classroom 3, probe images might have had a higher resolution than in previous classroom experiments since the results for the setting without margin adjustment presented the highest rate, similar to the results on the simulated ICB-RW benchmark. The most consistent models were the proposed architectures based on SRGAN with results around 80%, which overcame the baseline in every possible setting. However, when the size and margin were adjusted, FSRCNN presented the best results and the outcome for the other models became similar.

Table 5 – Accuracy (in %) for Face Recognition task in Classroom 2 (1xN)

Results UFS Classroom 2 (N = 16 students)			
	No Resize No Margin	Size 40 No Margin	Size 40 Margin 1.3
SRCNN	26.56	23.99	58.06
SRCNN Coord	26.56	31.32	50.92
SubCNN	23.99	28.94	41.39
SubCNN Coord	21.61	16.85	43.77
FSRCNN	33.88	26.56	51.10
FSRCNN Coord	31.50	26.56	55.86
FSRCNN Coord Face Loss	33.88	26.56	55.86
SRGAN	39.01	45.97	67.95
SRGAN Coord	38.83	36.45	72.71
SRGAN FaceLoss	48.72	48.72	67.95
SRGAN Coord FaceLoss	53.66	38.64	67.95
Bicubic	26.74	26.56	55.49

→ **Red color** highlights the architectures proposed in this thesis.

Table 6 – Accuracy (in %) for Face Recognition task in Classroom 3 (1xN)

Results UFS Classroom 3 (N = 12 students)			
	No Resize No Margin	Size 40 No Margin	Size 40 Margin 1.3
SRCNN	40.00	40.00	66.67
SRCNN Coord	36.67	56.67	56.67
SubCNN	40.00	46.67	60.00
SubCNN Coord	40.00	43.33	60.00
FSRCNN	46.67	43.33	80.00
FSRCNN Coord	46.67	43.33	73.33
FSRCNN Coord Face Loss	40.00	33.33	80.00
SRGAN	70.00	63.33	70.00
SRGAN Coord	76.67	76.67	73.33
SRGAN FaceLoss	83.33	76.67	70.00
SRGAN Coord FaceLoss	86.67	80.00	70.00
Bicubic	46.67	40.00	66.67

→ **Red color** highlights the architectures proposed in this thesis.

It was possible to check that, for all experiments in this task, increasing the amount of information within the image with a 1.3 margin resulted in an increase in accuracy for most

algorithms. This increase did help the SR network to provide more discriminative face images for the feature extractor since the average accuracy results were higher in general for such setting.




Figure 25 – Performance results for accuracy on ICB-RW and UFS Classroom 1 data. Proposed architectures have an asterisk in their names.(Source: Author's own)




Figure 26 – Performance results for accuracy on UFS Classroom 2 and UFS Classroom 3 data. Proposed architectures have an asterisk in their names. (Source: Author's own)

5.2.4 Hypotheses Discussion

For checking if there is correlation between image quality metrics and accuracy performance, the application of a correlation test was necessary. Since it was not possible to confirm if the original image data distribution approached normality, the Spearman Correlation Coefficient was calculated since it is specific for nonparametric data. Its results can be seen in Table 7.

Table 7 – Spearman Correlation PSNR/SSIM vs. Accuracy

	PSNR	SSIM
Spearman Correlation Coefficient	-0.3625	0.1159
p-value	5.671 e-07	0.121

As can be seen by Table 7, the null hypothesis (there is no dependency) can be rejected in the case of SSIM, yet it was not the case for the PSNR since its p-value was less than 0.05. However, this result implies that there is a negative correlation involving PSNR. This outcome can be explained by the fact that the best performing models for accuracy (SRGAN and its

variants) had the worst PSNR results when compared to the other models, which occasionally performed poorly in the face verification step.

Regarding accuracy, all architectures that took advantage of coordinate convolutions kept or increased its accuracy performance 72% of the time for all experiments. However, even with the positive results, after applying the Wilcoxon Signed-Rank Test to check if the hypothesis of having CoordConvs brought substantial gains, the p-value value was equal to 0.083. This p-value was not sufficient to reject the null hypothesis, which may indicate that either the results data distributions were the same or there was not sufficient data to point their difference. Since the evaluation of different architectures was meant to be applied to real-world LR data, which in this case is limited, this CoordConv operator still needs to be further explored to have a real measure of its potential for the SR task. Even so, it has shown already promising results for general SR in face biometrics and can be employed for different architectures performing similar tasks.

Architectures that were optimized with the specially designed for feature extraction “FaceLoss” kept or increased its accuracy performance 77% of the time for all experiments. Also, they presented the highest average accuracy across all experiments and most of the best results, as it is shown in Figures 25 and 26. The Wilcoxon Signed-Rank Test presented a p-value of 0.03, which means that the null hypothesis (results have same distribution) can be rejected. Therefore, architectures with this loss had results from a different data distribution when compared to the same models trained on their general standard losses.

According to Figures 25 and 26, the verification accuracy increased when the 1.3 margin was applied to naturally LR images (UFS Classroom data) and decreased when probe images were in HR and already matched the gallery resolution (ICB-RW data). This suggests that a simple system that monitors what to apply to detected faces depending on their size could be elaborated to take advantage of this characteristic and provide significant improvements for a LR face recognition pipeline.

In general, deeper architectures (from literature and adapted) had better performance over the shallower ones and the baseline for the executed experiments. Even the simplest SR models (SRCNN, SubCNN, FSRCNN and their variants) were able to beat bicubic interpolation by at least a tiny margin, as shown in Figures 25 and 26. Nevertheless, this small margin may not justify the use of a simple DL model for upscaling images prior to recognition in real-world situations since the loss of FPS is still substantial and needs to be taken into consideration.

5.3 Final Considerations

In this chapter, the experiments elaborated for validating the initial hypotheses were presented, and the results regarding the proposed architectures were discussed. The use of deep SR models for enhancing image features before verification proved to be a beneficial step in the low-resolution face recognition pipeline. The models that made use of adversarial

training (GANs) with different loss functions presented not only the best results regarding visual perception, as seen in Appendix [A](#), but also recognition performance since they were able to produce a clearer image for feature extraction.

The final chapter presents the final considerations taking into account the whole thesis and the conclusions.

6

Conclusions

Super-resolution has shown vastly on recent works, and also reaffirmed in this thesis, both its ability to enhance the clarity and visual aspects of images and its potential to improve the accuracy performance of face recognition systems.

In this work, several state-of-the-art deep learning architectures for SR were implemented, modified and evaluated with the objective of enhancing face recognition performance in two naturally LR datasets. The application of SR models with the proposed change that addressed the use of an operator that takes into account the position information resulted in an equal or better accuracy 72% of the time over the use of the same architectures without adaptation. Meanwhile, applications of SR architectures that were optimized with the loss that prioritized better feature extraction obtained a comparable or improved accuracy 77% of the time.

The deeper networks (SRGAN and its variants) presented results that were both perceptually good for the human eye and the evaluation of the recognition criteria. The Inception ResNet V1 feature extractor with the proposed SRGAN FaceLoss architecture, the best performing SR model, had an average accuracy in all experiments of 73.54%, which overcame the bicubic interpolation baseline by 17%. Also, this same setup was able to achieve 85.11% in the simulated ICB-RW dataset, which is an indication that such strategy might be able to defeat the SOTA model for the original benchmark since its accuracy was around 84%.

The study performed in this thesis also confirmed that most of the other recently proposed SR deep learning architectures, even the not so deep ones, could be effective for recovering discriminant features of LR face images in real-world settings. In addition, the deeper network presented real-time capabilities when using a small GPU, which can facilitate their implementation for real-world surveillance systems.

Regarding the specific objectives elaborated for this thesis, some points are important to be highlighted:

- Even though SR algorithms are usually optimized for upscaling an image and obtaining good PSNR metrics, not always this super-resolved image is going to present the most discriminative features for face recognition.
- Coordinated Convolutions presented gains for both image quality metrics and verification accuracy in the pipeline when applied to SR network architectures. However, they still need further studies since the data distribution for the results of networks with and without them presented high similarities, as confirmed by the Wilcoxon Signed-Rank Test.
- The use of a custom loss function to enhance the discriminative face features in images allows solid gains to the accuracy performance of a LR face recognition pipeline.

As future work, a simple system can be proposed for applying the best margin to the crop size of a detected face to ensure there is enough information available for feature extraction. Notwithstanding, a comparative study should be necessary for evaluating in which size the recognition accuracy would start to drop.

Also, different strategies may be employed to tackle the deficiencies presented in an SR pipeline for LR face recognition. One of the problems is the need for a fixed upscale factor when training the SR network. This challenge may be tackled by the use of a meta-upscaling strategy based on the recent work of [Hu et al. \(2019\)](#) where, for example, the weights of the upscaling network may be predicted based on the knowledge acquired by meta-features extracted from similar face datasets.

Another challenge to be solved is that some deep generative SR architectures may not be able to achieve real-time performance in CPU or mobile devices due to higher computational complexity. To overcome that, the use of distillation methods to prune these networks may improve time for inference to the cost of losing some performance ([ZHANG et al., 2018](#)).

Bibliography

- ABELLO, A. A.; JR., R. H. Optimizing super resolution for face recognition. In: SBC. *SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*. [S.l.], 2019. Cited 3 times in page(s) [17](#), [35](#), and [42](#).
- AGHDAM, O. A. et al. Exploring factors for improving low resolution face recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. [S.l.: s.n.], 2019. p. 0–0. Cited 6 times in page(s) [34](#), [36](#), [40](#), [41](#), [49](#), and [50](#).
- AGUSTSSON, E.; TIMOFTE, R. Ntire 2017 challenge on single image super-resolution: Dataset and study. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. [S.l.: s.n.], 2017. Cited in page [27](#).
- ATAER-CANSIZOGLU, E. et al. Verification of very low-resolution faces using an identity-preserving deep face super-resolution network. *arXiv preprint arXiv:1903.10974*, 2019. Cited in page [34](#).
- BAKER, S.; KANADE, T. Hallucinating faces. *fg*, Citeseer, v. 2000, p. 83–88, 2000. Cited in page [31](#).
- BAO, J. et al. Cvae-gan: fine-grained image generation through asymmetric training. In: *Proceedings of the IEEE International Conference on Computer Vision*. [S.l.: s.n.], 2017. p. 2745–2754. Cited in page [23](#).
- BEGIN, I.; FERRIE, F. P. Comparison of super-resolution algorithms using image quality measures. In: IEEE. *The 3rd Canadian Conference on Computer and Robot Vision (CRV'06)*. [S.l.], 2006. p. 72–72. Cited in page [27](#).
- BERGER, G.; PEYRARD, C.; BACCOUCHE, M. Boosting face recognition via neural super-resolution. In: *ESANN*. [S.l.: s.n.], 2016. Cited in page [33](#).
- CAO, Q. et al. Vggface2: A dataset for recognising faces across pose and age. In: IEEE. *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. [S.l.], 2018. p. 67–74. Cited 3 times in page(s) [8](#), [29](#), and [36](#).
- CHANG, C.-H. et al. Memory and perception-based facial image reconstruction. *Scientific reports*, Nature Publishing Group, v. 7, n. 1, p. 6499, 2017. Cited in page [15](#).
- CHEN, Y. et al. Deep feature extraction and classification of hyperspectral images based on convolutional neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, IEEE, v. 54, n. 10, p. 6232–6251, 2016. Cited in page [29](#).
- CHEN, Y. et al. Fsrnet: End-to-end learning face super-resolution with facial priors. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, p. 2492–2501, 2018. Cited 4 times in page(s) [26](#), [27](#), [32](#), and [37](#).
- COLE, F. et al. Synthesizing normalized faces from facial identity features. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2017. p. 3703–3712. Cited in page [42](#).

CROSSWHITE, N. et al. Template adaptation for face verification and identification. *Image and Vision Computing*, Elsevier, v. 79, p. 35–48, 2018. Cited in page 29.

DESHPANDE, A. *A Beginner's Guide To Understanding Convolutional Neural Networks*. 2017. <<https://adeshpande3.github.io/A-Beginner's-Guide-To-Understanding-Convolutional-Neural-Networks/>>. Accessed: 2019-11-25. Cited 2 times in page(s) 8 and 20.

DONAHUE, C.; MCAULEY, J.; PUCKETTE, M. Synthesizing audio with generative adversarial networks. *arXiv preprint arXiv:1802.04208*, 2018. Cited in page 23.

DONG, C. et al. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, IEEE, v. 38, n. 2, p. 295–307, 2015. Cited 3 times in page(s) 26, 31, and 43.

DONG, C.; LOY, C. C.; TANG, X. Accelerating the super-resolution convolutional neural network. In: SPRINGER. *European conference on computer vision*. [S.l.], 2016. p. 391–407. Cited 4 times in page(s) 8, 32, 43, and 44.

ELSAYED, A. et al. Unsupervised face recognition in the wild using high-dimensional features under super-resolution and 3d alignment effect. *Signal, Image and Video Processing*, Springer, v. 12, n. 7, p. 1353–1360, 2018. Cited in page 34.

FACELI, K. et al. Inteligência artificial: Uma abordagem de aprendizado de máquina. 2011. Cited in page 29.

FELDSTEIN, S. The global expansion of ai surveillance. *Carnegie Endowment*. <https://carnegieendowment.org/2019/09/17/global-expansion-of-ai-surveillance-pub-79847>, 2019. Cited in page 16.

GECER, B. et al. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2019. p. 1155–1164. Cited in page 42.

GERCHBERG, R. Super-resolution through error energy reduction. *Optica Acta: International Journal of Optics*, Taylor & Francis, v. 21, n. 9, p. 709–720, 1974. Cited in page 16.

GHALEB, E. et al. Deep representation and score normalization for face recognition under mismatched conditions. *Ieee Intelligent Systems*, IEEE, v. 33, n. 3, p. 43–46, 2018. Cited in page 50.

GOBBINI, M. I.; HAXBY, J. V. Neural systems for recognition of familiar faces. *Neuropsychologia*, Elsevier, v. 45, n. 1, p. 32–41, 2007. Cited in page 15.

GOODFELLOW, I. et al. Generative adversarial nets. In: *Advances in neural information processing systems*. [S.l.: s.n.], 2014. p. 2672–2680. Cited in page 22.

HARIS, M.; SHAKHNAROVICH, G.; UKITA, N. Task-driven super resolution: Object detection in low-resolution images. *arXiv preprint arXiv:1803.11316*, 2018. Cited 3 times in page(s) 16, 32, and 42.

HAXBY, J. V.; HOFFMAN, E. A.; GOBBINI, M. I. The distributed human neural system for face perception. *Trends in cognitive sciences*, Elsevier, v. 4, n. 6, p. 223–233, 2000. Cited in page 15.

- HE, K. et al. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2016. p. 770–778. Cited in page 21.
- HENNINGS-YEOMANS, P. H.; BAKER, S.; KUMAR, B. V. Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. In: IEEE. *2008 IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.], 2008. p. 1–8. Cited in page 33.
- HOWARD, J. et al. *fastai*. [S.l.]: GitHub, 2018. <<https://github.com/fastai/fastai>>. Cited 2 times in page(s) 39 and 47.
- HU, X. et al. Meta-sr: A magnification-arbitrary network for super-resolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2019. p. 1575–1584. Cited in page 56.
- HUANG, D.; LIU, H. A short survey of image super resolution algorithms. *Journal of Computer Science Technology Updates*, v. 2, n. 2, p. 19–29, 2015. Cited 3 times in page(s) 8, 24, and 25.
- HUANG, G. B. et al. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In: . [S.l.: s.n.], 2008. Cited in page 33.
- JIAO, L.; ZHAO, J. A survey on the new generation of deep learning in image processing. *IEEE Access*, IEEE, 2019. Cited 2 times in page(s) 8 and 44.
- JOHNSON, J.; ALAHI, A.; FEI-FEI, L. Perceptual losses for real-time style transfer and super-resolution. In: SPRINGER. *European conference on computer vision*. [S.l.], 2016. p. 694–711. Cited in page 23.
- JOHNSON, M. H. et al. Newborns’ preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, Elsevier, v. 40, n. 1-2, p. 1–19, 1991. Cited in page 15.
- KIM, D. et al. Progressive face super-resolution via attention to facial landmark. *arXiv preprint arXiv:1908.08239*, 2019. Cited 2 times in page(s) 16 and 37.
- KRIZHEVSKY, A.; SUTSKEVER, I.; HINTON, G. E. Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. [S.l.: s.n.], 2012. p. 1097–1105. Cited 2 times in page(s) 16 and 20.
- LECUN, Y.; BENGIO, Y.; HINTON, G. Deep learning. *nature*, Nature Publishing Group, v. 521, n. 7553, p. 436, 2015. Cited 5 times in page(s) 16, 19, 20, 21, and 29.
- LEDIG, C. et al. Photo-realistic single image super-resolution using a generative adversarial network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2017. p. 4681–4690. Cited 4 times in page(s) 23, 26, 32, and 43.
- LI, J.; FENG, J.; KUO, C.-C. J. Deep convolutional neural network for latent fingerprint enhancement. *Signal Processing: Image Communication*, Elsevier, v. 60, p. 52–63, 2018. Cited in page 16.
- LI, P. et al. On low-resolution face recognition in the wild: Comparisons and new techniques. *IEEE Transactions on Information Forensics and Security*, IEEE, v. 14, n. 8, p. 2000–2012, 2019. Cited in page 34.

- LIU, R. et al. An intriguing failing of convolutional neural networks and the coordconv solution. In: *Advances in Neural Information Processing Systems*. [S.l.: s.n.], 2018. p. 9605–9616. Cited 4 times in page(s) 8, 16, 23, and 24.
- LIU, Z. et al. Deep learning face attributes in the wild. In: *Proceedings of International Conference on Computer Vision (ICCV)*. [S.l.: s.n.], 2015. Cited 4 times in page(s) 8, 27, 36, and 37.
- MA, S. et al. Da-gan: Instance-level image translation by deep attention generative adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2018. p. 5657–5666. Cited in page 23.
- MUTTU, Y.; VIRANI, H. Effective face detection, feature extraction & neural network based approaches for facial expression recognition. In: IEEE. *2015 International Conference on Information Processing (ICIP)*. [S.l.], 2015. p. 102–107. Cited in page 28.
- NEVES, J.; MORENO, J.; PROENÇA, H. Quis-campi: an annotated multi-biometrics data feed from surveillance scenarios. *IET Biometrics*, IET, v. 7, n. 4, p. 371–379, 2017. Cited 4 times in page(s) 8, 37, 38, and 40.
- NEVES, J.; PROENÇA, H. Icb-rw 2016: International challenge on biometric recognition in the wild. In: IEEE. *2016 International Conference on Biometrics (ICB)*. [S.l.], 2016. p. 1–6. Cited 4 times in page(s) 8, 37, 41, and 46.
- NGUYEN, H. V.; BAI, L. Cosine similarity metric learning for face verification. In: SPRINGER. *Asian conference on computer vision*. [S.l.], 2010. p. 709–720. Cited in page 29.
- NGUYEN, K. et al. Super-resolution for biometrics: A comprehensive survey. *Pattern Recognition*, Elsevier, v. 78, p. 23–42, 2018. Cited 3 times in page(s) 15, 16, and 25.
- OUYANG, N. et al. Deep joint super-resolution and feature mapping for low resolution face recognition. In: IEEE. *2018 IEEE International Conference of Safety Produce Informatization (IICSPI)*. [S.l.], 2018. p. 849–852. Cited in page 17.
- PEYRARD, C.; MAMALET, F.; GARCIA, C. A comparison between multi-layer perceptrons and convolutional neural networks for text image super-resolution. In: *VISAPP (I)*. [S.l.: s.n.], 2015. p. 84–91. Cited in page 33.
- PURKAIT, P.; PAL, N. R.; CHANDA, B. A fuzzy-rule-based approach for single frame super resolution. *IEEE Transactions on Image processing*, IEEE, v. 23, n. 5, p. 2277–2290, 2014. Cited in page 25.
- RASTI, P. et al. Convolutional neural network super resolution for face recognition in surveillance monitoring. In: SPRINGER. *International conference on articulated motion and deformable objects*. [S.l.], 2016. p. 175–184. Cited 3 times in page(s) 15, 16, and 33.
- REIBMAN, A. R.; BELL, R. M.; GRAY, S. Quality assessment for super-resolution image enhancement. In: IEEE. *2006 International Conference on Image Processing*. [S.l.], 2006. p. 2017–2020. Cited in page 27.
- RIBEIRO, E.; UHL, A. Exploring texture transfer learning via convolutional neural networks for iris super resolution. In: IEEE. *2017 International Conference of the Biometrics Special Interest Group (BIOSIG)*. [S.l.], 2017. p. 1–5. Cited in page 16.

SÁ, J. M. D. d. C. *Registro de Classe Automatizado Utilizando Reconhecimento Facial*. 74 p. Bachelor's Thesis — Universidade Federal de Sergipe, 2019. Cited 4 times in page(s) 8, 38, 39, and 40.

SCHROFF, F.; KALENICHENKO, D.; PHILBIN, J. Facenet: A unified embedding for face recognition and clustering. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2015. p. 815–823. Cited 3 times in page(s) 29, 35, and 40.

SHI, W. et al. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2016. p. 1874–1883. Cited 5 times in page(s) 8, 19, 32, 43, and 44.

SZEGEDY, C. et al. Inception-v4, inception-resnet and the impact of residual connections on learning. In: *Thirty-First AAAI Conference on Artificial Intelligence*. [S.l.: s.n.], 2017. Cited 2 times in page(s) 21 and 40.

SZEGEDY, C. et al. Going deeper with convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. [S.l.: s.n.], 2015. p. 1–9. Cited 3 times in page(s) 8, 20, and 21.

TIAN, J.; MA, K.-K. A survey on super-resolution imaging. *Signal, Image and Video Processing*, Springer, v. 5, n. 3, p. 329–342, 2011. Cited in page 16.

TIAN, L.; SUZUKI, A.; KOIKE, H. Task-oriented evaluation of super-resolution techniques. In: IEEE. *2010 20th International Conference on Pattern Recognition*. [S.l.], 2010. p. 493–498. Cited 2 times in page(s) 26 and 31.

TIMOFTE, R. et al. Ntire 2018 challenge on single image super-resolution: Methods and results. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. [S.l.: s.n.], 2018. Cited in page 27.

UPADHYAY, U.; SINGHAL, B.; SINGH, M. Spinal stenosis detection in mri using modular coordinate convolutional attention networks. In: IEEE. *2019 International Joint Conference on Neural Networks (IJCNN)*. [S.l.], 2019. p. 1–8. Cited in page 24.

VEDADI, F.; SHIRANI, S. A map-based image interpolation method via viterbi decoding of markov chains of interpolation functions. *IEEE Transactions on Image Processing*, IEEE, v. 23, n. 1, p. 424–438, 2014. Cited in page 25.

VEZHNEVETS, V. Face and facial feature tracking for natural human-computer interface. In: . [S.l.: s.n.], 2002. Cited in page 28.

WANG, M.; DENG, W. Deep face recognition: A survey. *arXiv preprint arXiv:1804.06655*, 2018. Cited in page 29.

WANG, Z. et al. Studying very low resolution recognition using deep networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2016. p. 4792–4800. Cited in page 34.

WANG, Z.; CHEN, J.; HOI, S. C. Deep learning for image super-resolution: A survey. *arXiv preprint arXiv:1902.06068*, 2019. Cited 4 times in page(s) 8, 16, 26, and 27.

WANG, Z.; SHE, Q.; WARD, T. E. Generative adversarial networks: A survey and taxonomy. *arXiv preprint arXiv:1906.01529*, 2019. Cited 3 times in page(s) 16, 22, and 23.

XU, X.; CHEN, Y.-C.; JIA, J. View independent generative adversarial network for novel view synthesis. In: *Proceedings of the IEEE International Conference on Computer Vision*. [S.l.: s.n.], 2019. p. 7791–7800. Cited in page 24.

YANG, Z. et al. Semi-supervised qa with generative domain-adaptive nets. *arXiv preprint arXiv:1702.02206*, 2017. Cited in page 23.

YU, J. et al. Generative image inpainting with contextual attention. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [S.l.: s.n.], 2018. p. 5505–5514. Cited in page 23.

YU, X. et al. Face super-resolution guided by facial component heatmaps. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. [S.l.: s.n.], 2018. p. 217–233. Cited in page 37.

ZAFEIRIOU, S.; ZHANG, C.; ZHANG, Z. A survey on face detection in the wild: past, present and future. *Computer Vision and Image Understanding*, Elsevier, v. 138, p. 1–24, 2015. Cited in page 28.

ZAFEIROULI, K. et al. Efficient, lightweight, coordinate-based network for image super resolution. In: IEEE. *2019 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*. [S.l.], 2019. p. 1–9. Cited 2 times in page(s) 24 and 33.

ZHANG, C.; ZHANG, Z. A survey of recent advances in face detection. 2010. Cited in page 29.


ZHANG, K. et al. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, IEEE, v. 23, n. 10, p. 1499–1503, 2016. Cited 2 times in page(s) 28 and 38.


ZHANG, L. et al. Adaptive importance learning for improving lightweight image super-resolution network. *arXiv preprint arXiv:1806.01576*, 2018. Cited in page 56.


Appendix


APPENDIX A – Perceptual Results of SR Algorithms for 4x Upscaling


In this section it is possible to check the perceptual results for the SR experiment described in Section 5.1.1. All the images were upscaled from a grid of 40x40 to a size of 160x160.











APPENDIX B – Average Training Losses for the SR Algorithms

This section presents the behavior for the training losses of each network obtained during the experiment described in Section 5.1.1.

