UNIVERSIDADE DE SÃO PAULO
Instituto de Ciências Matemáticas e de Computação

A deep learning approach to age newborns' fingerprints minutiae maps

Luiz Fernando Puttow Southier

Monograph - MBA in Artificial Intelligence and Big Data

ICMC USP SÃO CARLOS

Luiz Fernando Puttow Southier

A deep learning approach to age newborns' fingerprints minutiae maps

Monograph presented to the Departamento de Ciências de Computação do Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo - ICMC/USP, as part of the requirements for obtaining the title of Specialist in Artificial Intelligence and Big Data.

Concentration area: Artificial Intelligence and Big Data

Advisor: Prof. Dr. Dalcimar Casanova

Original version

São Carlos

2024

Ficha catalográfica elaborada pela Biblioteca Prof. Achille Bassi
e Seção Técnica de Informática, ICMC/USP,
com os dados inseridos pelo(a) autor(a)

S727d

Southier, Luiz Fernando Puttow
A deep learning approach to age newborns? fingerprints minutiae maps / Luiz Fernando Puttow Southier; orientador Dalcimar Casanova. -- São Carlos, 2024.
47 p.

Trabalho de conclusão de curso (MBA em Inteligência Artificial e Big Data) -- Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, 2024.

1. biometrics. 2. newborn. 3. minutiae map. 4. deep learning. I. Casanova, Dalcimar, orient. II. Título.

Bibliotecários responsáveis pela estrutura de catalogação da publicação de acordo com a AACR2:
Gláucia Maria Saia Cristianini - CRB - 8/4938
Juliana de Souza Moraes - CRB - 8/6176

Luiz Fernando Puttow Southier

A deep learning approach to age newborns' fingerprints minutiae maps

Monografia apresentada ao Departamento de Ciências de Computação do Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo - ICMC/USP, como parte dos requisitos para obtenção do título de Especialista em Inteligência Artificial e Big Data.

Área de concentração: Inteligência Artificial e Big Data

Advisor: Prof. Dr. Dalcimar Casanova

Versão original

São Carlos

2024

ACKNOWLEDGEMENTS

We extend our sincere gratitude to InfantID for their invaluable operational support to our neonatal biometrics research group at UTFPR. Their partnership has been instrumental in advancing our exploration of biometric technologies in neonatal care. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001, and by the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPQ).

ABSTRACT

Southier, L. F. P. **A deep learning approach to age newborns' fingerprints minutiae maps**. 2024. 47 p. Monograph (MBA in Artificial Intelligence and Big Data) - Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, São Carlos, 2024.

Despite the recognized need for neonatal identification systems, they rely on minutiae maps that are distorted over time due to finger growth, weight loss, and skin alterations, affecting recognition rates. The literature has attempted to address the challenge of emulating changes in the minutiae map by utilizing scaling factors. However, the effectiveness of these approaches, especially in younger individuals, remains unclear. While scaling factors have been used to mitigate the aging effect, this method oversimplifies the growth pattern assumption, neglecting distinctive growth patterns observed in young children. This study aims to fill this gap by developing and evaluating an artificial intelligence solution to emulate the changes in a newborn minutiae map over time. The hypothesis is that a Deep Learning-based growth model could provide a more effective approach than traditional scaling factors, offering a more accurate and adaptable representation of newborn development and mitigating the impact of the aging effect on biometric features. However, without a high-quality longitudinal database, a differentiable loss function based on the difference between two minutiae maps, and a method for aligning these maps, implementing such an approach becomes unfeasible. Future research should address these critical aspects to advance the development of effective biometric solutions for neonatal identification.

Keywords: biometrics; newborn; minutiae map; deep learning.

RESUMO

Southier, L. F. P. **A deep learning approach to age newborns' fingerprints minutiae maps**. 2024. 47 p. Monografia (MBA em Inteligência Artificial e Big Data) - Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, São Carlos, 2024.

Apesar da necessidade reconhecida de sistemas de identificação neonatal, eles dependem de mapas de minúcias que se distorcem ao longo do tempo devido ao crescimento dos dedos, perda de peso e alterações na pele, afetando as taxas de reconhecimento. A literatura tentou abordar o desafio de emular as mudanças no mapa de minúcias utilizando fatores de escalonamento. No entanto, a eficácia dessas abordagens, especialmente em indivíduos mais jovens, permanece incerta. Embora os fatores de escalonamento tenham sido usados para mitigar o efeito da idade, esse método simplifica excessivamente a suposição de padrão de crescimento, negligenciando padrões de crescimento distintos observados em crianças pequenas. Este estudo visa preencher essa lacuna desenvolvendo e avaliando uma solução de inteligência artificial para emular as mudanças no mapa de minúcias de um recém-nascido ao longo do tempo. A hipótese é que um modelo de crescimento baseado em Deep Learning poderia fornecer uma abordagem mais eficaz do que os fatores de escalonamento tradicionais, oferecendo uma representação mais precisa e adaptável do desenvolvimento dos recém-nascidos e mitigando o impacto do efeito da idade nas características biométricas. No entanto, sem um banco de dados longitudinal de alta qualidade, uma função de perda diferenciável baseada na diferença entre dois mapas de minúcias e um método para alinhar esses mapas, a implementação de tal abordagem torna-se inviável. Pesquisas futuras devem abordar esses aspectos críticos para avançar no desenvolvimento de soluções biométricas eficazes para identificação neonatal.

Palavras-chave: biometria; recém-nascido; mapa de minúcias; aprendizado profundo.

LIST OF FIGURES

CONTENTS

1 INTRODUCTION

Every year, the global population witnesses the birth of over 134 million children, resulting in more than 656 million children under the age of five worldwide in 2022 (United Nations - Department of Economic and Social Affairs, 2023). Despite commendable efforts and advancements over the past decades, the imperative to enhance child survival persists. The year 2021 alone witnessed approximately 13,800 deaths per day of children under the age of five, an alarming statistic underscoring the pressing need for improved child care and survival strategies (UNICEF United Nations International Children's Emergency Fund, 2023b).

The lack of efficient identification, as highlighted by the World Health Organization (WHO), pose challenges in accurately tracking vaccination schedules (WHO World Health Organization, 2023a). The consequences are evident in the stagnation of global vaccination coverage, with 14.3 million children not receiving vaccinations in 2022, compared to 12.9 million in the pre-pandemic era (WHO World Health Organization, 2023b). Furthermore, 1.5 million children succumb to vaccine-preventable diseases each year (UNICEF United Nations International Children's Emergency Fund, 2023a). In regions with a high prevalence of HIV, timely diagnosis and intervention for HIV-exposed infants are crucial, and biometric solutions can aid in identifying and tracking these infants, ensuring prompt access to life-saving treatments (Sirirungsi *et al.*, 2016). In nutrition programs, the importance of neonatal identification is underscored by reports of fraud diverting food meant for needy children (WFP United Nations World Food Programme, 2023). Biometric data could offer a solution by facilitating the monitoring of nutrition programs, allowing for early interventions, and promoting healthy growth by tracking weight gain and vital metrics in children (Grantham-McGregor *et al.*, 2014). From a child protection standpoint, an efficient infant identification system would mitigate baby swaps in hospitals, locate missing or abducted children, and enhance airport security to prevent child trafficking (Sedlak, 2002). Neonatal biometrics also play a crucial role in national identity programs, ensuring secure and lifelong identities for children.

Despite the consensus on the need for neonatal identification systems, neonates' delicate skin and tiny fingers pose challenges for conventional biometric systems designed for adult use (Jain; Cao; Arora, 2014). Adding to the complexity, in the initial five days post-birth, individuals undergo a daily body weight loss of 1-2%, accompanied by a reduction in extracellular water from 40% to 30% (Lindower, 2017; Sharma; Ford; Calvert, 2011). Preterm infants may experience even more pronounced effects (Visscher *et al.*, 2015), leading to a profound shift in body composition that influences the composition and thickness of the fingers skin layers (Visscher *et al.*, 2015; Finlayson *et al.*, 2022).

Additionally, the rapid changes and growth in the children's fingers in the early stages of life impose an extra challenge referred to by the literature as the *ageing effect* (Galbally; Haraksim; Beslay, 2018a). The ageing effect refers to the fingerprint quality and recognition performance that varies as the time lapse between the reference fingerprint sample, usually acquired at early hours of life, and the probe fingerprint sample, collected at older ages, changes. This effect is primarily related to the concept of fingerprint permanence, which refers to the ability of fingerprints to remain consistent and recognizable over time (Galbally; Haraksim; Beslay, 2018b).

The biometric systems use a set of key points in the fingerprint, referred to as *minuatiae map*, to perform the recognition. To obtain a minutiae map from a fingerprint, first, the captured fingerprint needs to be segmented. The segmented image differentiates the raised lines of the finger's skin, named *ridges*, from the depressed or low areas between the ridges, named *valleys*. The minutiae map is defined by the points where ridges or valleys terminate. Each minutia m_i in the minutiae map $M = \{m_0, ..., m_i\}$ is represented by a tuple in the form of $m_i = (x_i, y_i, \theta_i)$, being x_i and y_i, and θ_i the cartesian coordinates and the angle for the i-th minutia, respectively. The changes due to the finger's growth, body weight loss, and skin composition alterations cause the minutiae map from a fingerprint to distort as time passes. These distortions, while minor for short time lapses become increasingly detrimental to the biometric system's recognition rate for more extended time lapses (Jain; Cao; Arora, 2014; Camacho *et al.*, 2017; Jain *et al.*, 2016b; Galbally; Haraksim; Beslay, 2018b; Engelsma *et al.*, 2021). Figure 1 illustrates the concepts of segmented fingerprints, ridge, valley, and minutiae maps.

Figure 1 – Comparison of original and segmented fingerprint image and its corresponding minutiae map.

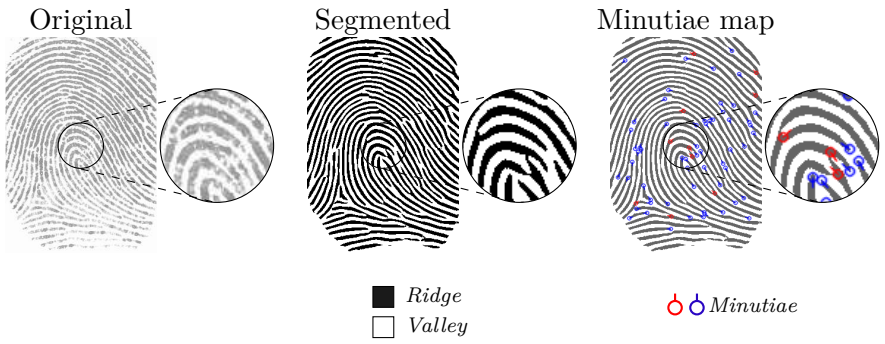

To overcome this challenge, the biometric systems would benefit from a computational solution that could emulate the minutiae map from a fingerprint after some timelapse. Specifically, given a reference minutiae map M_y collected at an age y and a time-lapse Δt, the solution f would emulate the fingerprint ageing and calculate the aged minutia map $M_{y'}$ at an age y', such that $y' = y + \Delta t$, and that $M_{y'} = f(M_y)$.

Recognizing this gap, the literature has presented approaches to emulate the changes

in the minutia map by applying scaling factors. Some studies apply a fixed scaling factor (Jain *et al.*, 2016b), which proves suboptimal compared to using a scaling factor based on age category (Camacho *et al.*, 2017). While two-factor growth models have been suggested (Haraksim; Galbally; Beslay, 2019; Markert *et al.*, 2019), they have not been evaluated on younger individuals.

While applying scaling factors has mitigated the ageing effect in the reported studies, this strategy seems to be an oversimplification because it assumes that fingerprints have a linear growth that such factors can model. However, as reported in (Schneider, 2010), young children tend to exhibit distinctive growth patterns distorting the minutia map in other ways. Additionally, these studies only use the minutiae maps as input for performing the ageing emulation. However, since young children tend to have a low-quality minutia map (Camacho *et al.*, 2017), the quantity and quality of minutiae are largely affected, and additional information, such as the ridge and valley configuration that defines the minutiae relation, would hypothetically benefit the emulation solution.

This study's primary objective is to create and evaluate an artificial intelligence solution to emulate the changes a newborn minutiae map can suffer due to natural growth over time. The investigation aims to assess the feasibility of this approach, evaluating the model's efficacy in attenuating distortions arising from ageing and the distinct physical transformations characteristic of this demographic.

The underlying hypothesis of this work is that employing artificial intelligence can provide a better approach to emulating the growth and changes in the minutiae map of newborns over time. It is believed that by developing a deep learning-based growth model that includes as input the original fingerprint, the segmented fingerprint (valley and ridge relation), and the minutiae map, it is possible to mitigate the impact of the ageing effect on biometric features, thereby providing a more accurate and adaptable representation of newborn development than simply applying scaling factors as reported by the literature. In general terms, the workflow of the proposed solution is illustrated in Figure 2, in which M_y is the minutia map at age y, Δt is the time-lapse, and $M_{y'}$ is the calculated minutia map at age y'.

Figure 2 – Workflow of the proposed solution for emulating minutia map growth.

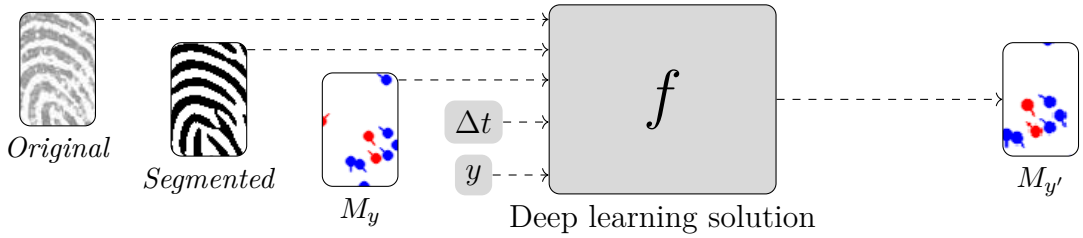

To achieve the proposed goal, the following is conducted:

- Prepare and preprocess a database regarding newborn fingerprints containing several images collected at different times for each fingerprint. This is currently being conducted in the scope of a research group[1].
- Propose and evaluate a deep learning model that has as input the original and segmented fingerprint images, the minutia map, and the information of age and timelapse.

[1] Information Processing Research Group - Federal University of Technology - Paraná (in Portuguese: *Grupo de Pesquisas em Processamento de Informação na Universidade Tecnológica Federal do Paraná - UTFPR*). The research has been approved with a Certificate of Presentation of Ethical Appreciation 73791023.7.0000.0177 at the Brazil Platform.

2 THEORETICAL BACKGROUND

This chapter introduces the theoretical background, laying the groundwork for the subsequent discussions. Section 2.1 explores biometric systems, explains the concepts of identification and verification, and shows fingerprint aspects, especially minutiae maps, and how they are used for recognition. Section 2.2 presents the fundamental concepts of neural networks and deep learning. Lastly, Section 2.3 reviews related studies, offering insights into existing research and providing a contextual comparison of the study's contributions.

2.1 Biometric systems

Biometrics refers to the automated recognition of individuals based on unique biological and behavioral characteristics, as highlighted by Jain, Hong and Pankanti (2000). Widely acknowledged for its sophistication and reliability, biometric technology finds extensive application across diverse security domains (Li, 2009). This technology leverages an array of distinctive physical or behavioral attributes, including but not limited to fingerprints, facial characteristics, iris patterns, voice-prints, and even gait, to establish and verify one's identity, as extensively discussed by (Li, 2009).

Figure 3 provides a comprehensive illustration of a biometric recognition system, delineating the process of capturing, processing, and analyzing biometric data to authenticate individuals.

Figure 3 – General schema of a biometric recognition system.

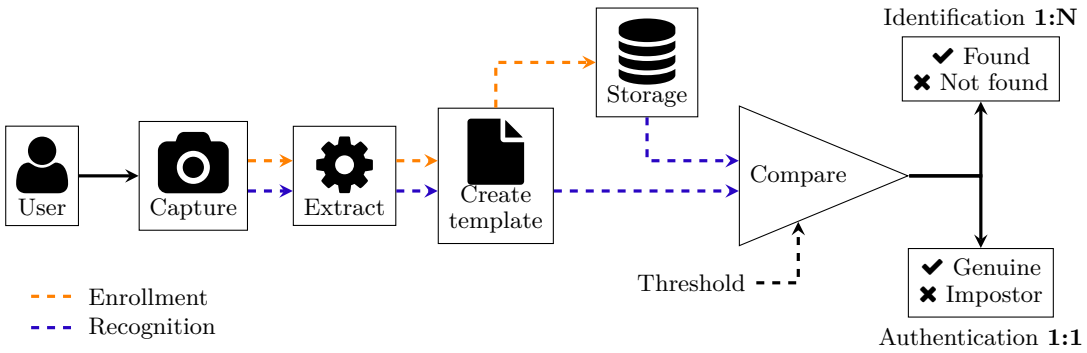

Source: Southier *et al.* (2023)

The initial stage in a biometric system involves *enrollment*, where one or multiple biometric samples are gathered, processed, and stored as *templates*. Typically, this processing phase entails the extraction of key features. Subsequently, when enrolled users seek access to secure resources, they once again provide their biometric data to the capture mechanism, a step referred to as *recognition*. Following this, the processing system

undertakes one of two actions: authentication or identification (Jain; Hong; Pankanti, 2000).

Authentication occurs when the system exclusively verifies the live biometric against stored templates from a specific user, constituting a one-to-one (1:1) match. Conversely, in identification, also known as verification, the system compares the live biometric against a database containing stored templates from multiple individuals, conducting a one-to-many (1:N) search for identification purposes. In both scenarios, if a match is detected, granting access signifies the individual as *genuine*. Conversely, if no match is found, the individual is deemed an *impostor*. Both matching cases tolerate a certain degree of error, which is defined by a predetermined threshold (Jain; Ross; Prabhakar, 2004; Phillips *et al.*, 2000).

Biometric systems encompass a plethora of features, as discussed in Jain and Kumar (2012). For instance, the human eye presents unique resources for developing recognition technologies, such as the intricate texture of the *iris* and the distinctive pattern of fundus blood vessels. Another option for human recognition lies in *palmprints*, which reveal the physical traits of the skin when pressed against a surface, including lines, points, and texture (Jain; Hong; Pankanti, 2000). Additionally, factors like *hand geometry*, encompassing finger lengths, widths, and overall hand dimensions, contribute to recognition. However, hand geometry biometrics might not be ideal for widespread personal identification due to its lack of distinctiveness on a large scale (Li, 2009). The details about fingerprints are discussed in Section 2.1.1.

Authentication aims to ascertain whether a person is who they claim to be. These systems calculate the biometric *score* by juxtaposing the stored template with the user's input (Li, 2009). This process is called *matching*. If the biometric score surpasses a predetermined *threshold*, the user is deemed *accepted*. Conversely, if the score falls below the threshold, the user is *rejected*. Equation 2.1 demonstrates the decision rule of the verification process (Jain; Bolle; Pankanti, 2011):

$$(I, x) \in \begin{Bmatrix} accepted, & if\ s \geq \eta \\ rejected, & if\ s < \eta \end{Bmatrix} \tag{2.1}$$

where I is the individual's identity; x is the set of input features; s is the matching score; and η is a threshold value.

In identification, the goal is to verify if a person is known to the system. In this scenario, the input is compared with the templates of all registered identities. The output can be the comparison with the highest similarity or a warning that the identity is not registered in the database. Equation 2.2 demonstrates the decision rule of the identification process (Jain; Bolle; Pankanti, 2011):

$$x \in \begin{cases} I_{n_0}, & \text{if } n_0 = \arg\max_n s_n \text{ and } s_{n_0} \geq \eta \\ I_{N+1}, & \text{otherwise} \end{cases} \tag{2.2}$$

where x is the set of input features; I is the set of registered identities; I_{n_0} is the identity with the highest similarity to the input; I_{N+1} is an unregistered identity; s is the matching score; and η is a threshold value.

Ideally, a flawless biometric system would accept genuine users and reject impostors unfailingly. However, real-world scenarios diverge from this ideal. As illustrated in Figure 4, contingent upon the threshold and the score obtained, certain impostors may be accepted, while some genuine users might be rejected.

Figure 4 – Distribution scores with a defined threshold

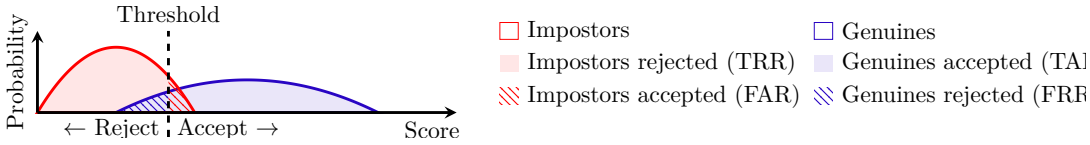

Source: Southier *et al.* (2023)

The evaluation of biometric systems involves several key metrics that provide insights into their performance and effectiveness in identifying genuine users while rejecting impostors (Jain; Bolle; Pankanti, 2011). These metrics are essential for assessing the system's reliability and suitability for specific applications. One fundamental metric is the *False Reject Rate (FRR)*, which quantifies the percentage of genuine users that are incorrectly rejected by the system. Its complement, 1-FRR, represents the *True Acceptance Rate (TAR)*, indicating the percentage of genuine users accepted by the system. On the other hand, the *True Reject Rate (TRR)* measures the share of impostors correctly rejected, while the *False Acceptance Rate (FAR)* represents the percentage of impostors incorrectly accepted by the system, with TRR=1-FAR.

These four rates are interconnected and dependent on the chosen threshold for decision-making (Jain; Bolle; Pankanti, 2011). By adjusting the threshold, it's possible to trade off between false acceptance and false rejection rates, influencing the system's overall performance. Two commonly used curves to illustrate this behavior are the *Detection Error Trade-off (DET)* curve and the *Receiver Operating Characteristics (ROC)* curve. These curves provide a graphical representation of the trade-off between false acceptance and false rejection rates across different threshold values, aiding in the optimization of system parameters. In an identification scenario, each live biometric sample is compared against all samples stored in the gallery. The resulting similarity scores are then sorted and ranked. The *True Positive Identification Rate (TPIR)* measures the probability of correctly identifying the individual within the top K ranks of the sorted list. This metric provides valuable insights into the system's ability to accurately identify individuals from

a large pool of candidates. To visualize the identification performance comprehensively, the *Cumulative Match Characteristic (CMC)* curve is often utilized. This curve illustrates the probability of correctly identifying the individual as a function of the number of samples considered, offering a clear depiction of the system's identification capabilities across different rank thresholds.

Figure 5 presents examples of DET, ROC, and CMC curves for two distinct biometric systems, providing stakeholders with visual representations of their performance characteristics and aiding in the decision-making process regarding system selection and configuration.

Figure 5 – Examples of DET (left), ROC (center), and CMC (right) curves

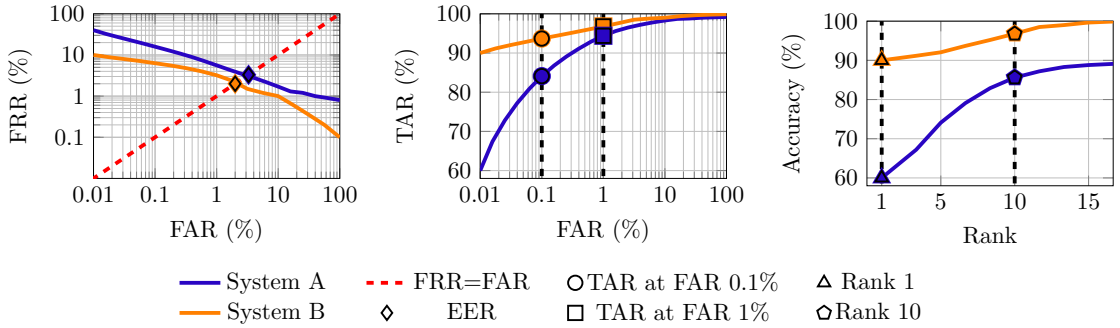

Source: Southier *et al.* (2023)

The *Detection Error Trade-off (DET)* curve and the *Receiver Operating Characteristics (ROC)* curve are essential tools for evaluating the performance of biometric systems, particularly in terms of error rates. The DET curve plots the False Acceptance Rate (FAR) and False Reject Rate (FRR) on both axes using a logarithmic scale. At a specific point on the DET curve, known as the *Equal Error Rate (EER)*, the FAR equals the FRR. This point provides a concise summary of the system's performance, indicating the threshold where the rates of false acceptance and false rejection are balanced.

On the other hand, the ROC curve illustrates the True Acceptance Rate (TAR) against various FARs. Typically, TAR values corresponding to FARs of 0.1% and 1% are commonly used for comparing different biometric systems. These values offer insights into the system's performance under different levels of security requirements and false acceptance tolerances.

2.1.1 Fingerprints

A more practical approach involves utilizing only *fingerprints* (or *digitprints*) for biometric recognition. For capturing fingerprint templates, *scanners*, also referred to as readers or sensors, are employed. Biometric features are captured as images, with each image termed an *impression* (Li, 2009). These scanners come in various forms, including contact or contactless variants, those equipped with multiple cameras, employing diverse

illumination schemes, and more. However, one of the most crucial properties of a scanner is its resolution. This aspect becomes especially pertinent when the collected zone's size in an image is reduced, such as in the case of neonatal fingerprints (Engelsma *et al.*, 2019; Jain *et al.*, 2016a). Resolution denotes the amount of information in pixels that the resulting image can encapsulate per inch (pixels per inch or ppi). Higher scanner resolutions translate to greater information extraction but also entail higher costs. Typically, standard fingerprint scanners used in adults boast a resolution of 500 ppi. Fingerprint features consist of *ridges* alternating with *valleys*, running mostly parallel but exhibiting occasional directional changes and terminations. Figure 6 shows examples of scanning and the corresponding collected fingerprint.

Figure 6 – Example of fingerprint collection.

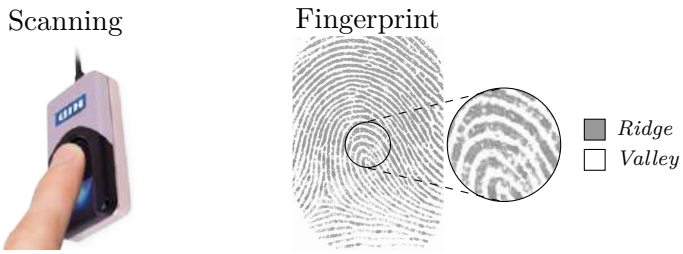

Upon capturing a fingerprint, the original image undergoes several preprocessing steps, including scaling and filtering, to derive its *segmented* representation. This segmented image enhances the fingerprint features and is normally expressed as a binarized image. Subsequently, the *directional field* is computed, depicting how the ridges and valleys traverse the print. The directional field is not normally used for matching, but it helps detect the minutiae directions in later steps. Next, the *skeletonized* version of the segmented image is calculated. Figure 7 shows an example of a fingerprint undergoing these steps.

Figure 7 – Fingerprint preprocessing pipeline

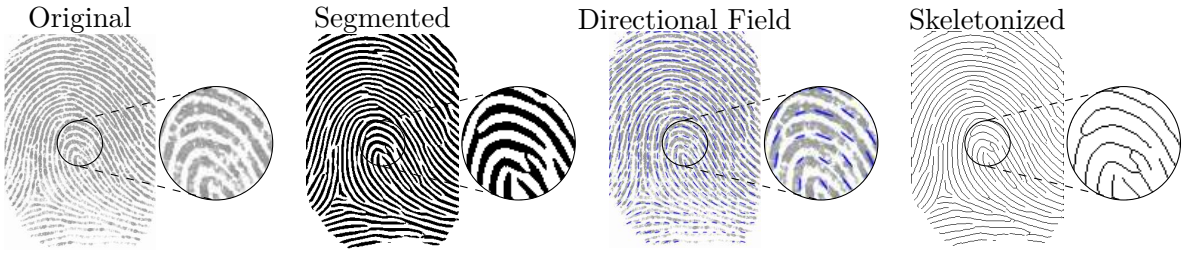

Using the skeletonized image, the extraction of the *minutia map* occurs, identifying bifurcations and terminations. The minutiae map is defined by the points where ridges or valleys terminate. Each minutia m_i in the minutiae map $M = \{m_0, ..., m_i\}$ is represented by a tuple in the form of $m_i = (x_i, y_i, \theta_i)$, being x_i and y_i, and θ_i the cartesian coordinates and the angle for the i-th minutia, respectively. Furthermore, minutiae can be classified as

termination or bifurcation. Pixels with three neighbors are identified as bifurcations, and pixels with only one neighbor are identified as terminations. Figure 8 shows an example of a minutiae map highlighting a termination and a bifurcation.

Figure 8 – Example of minutiae map.

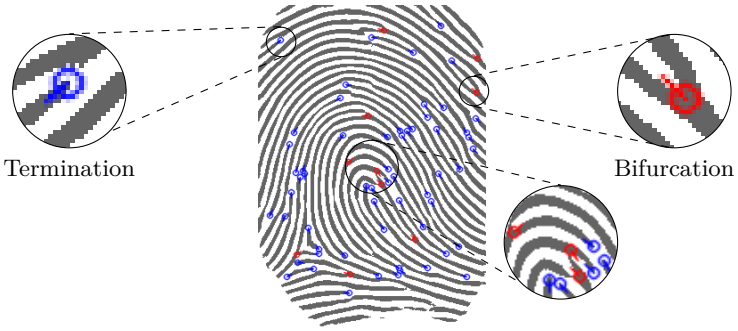

Typically, the minutia maps of two fingerprints are compared to calculate the *match score* (Jain; Kumar, 2012). First, an alignment geometric transformation is performed so that the two minutiae maps are in the same coordinate system. Next, the *matching* is executed, which forms corresponding pairs of minutiae. Finally, the match score is generated based on the matched minutiae pairs. The higher the matched pairs, the higher the match score. Figure 9 shows two examples of matching: (a) shows a genuine user with a high match score; (b) shows an impostor with a low match score.

Figure 9 – Examples of genuine and impostor matchings.

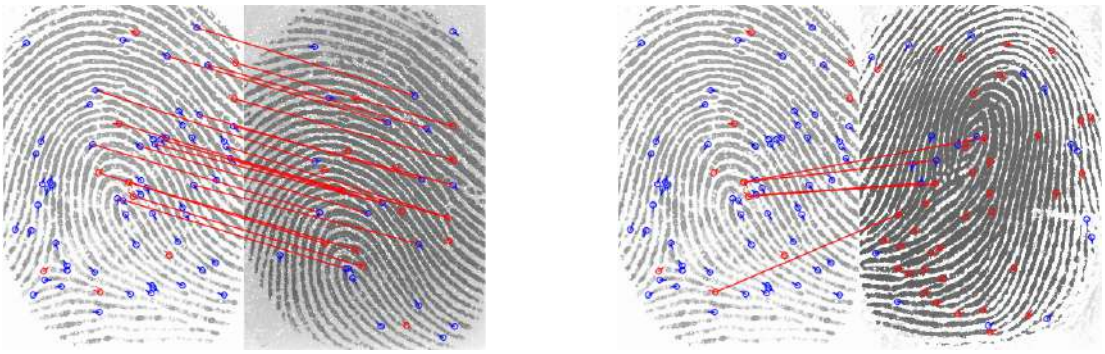

(a) Genuine – high match score

(b) Impostor – low match score

2.2 Neural Networks and Deep Learning

Based on the findings of neuroscience, in particular, the hypothesis that mental activity consists primarily of electrochemical activity in networks of brain cells called neurons, some researchers worked to create artificial neural networks (Russell; Norvig, 2010). Researchers in Artificial Intelligence and statistics became interested in the more abstract properties of neural networks, such as their ability to perform distributed computation, tolerate noisy inputs, and learn. Neural networks are composed of neurons connected by

directed links. One of the simplest neural networks, *perceptron*, contains a single input layer and an output node. Its basic architecture is shown in Figure 10. The input layer contains d nodes that transmit the d features $\overline{X} = [x_1, \ldots, x_d]$ with edges of weight $\overline{W} = [w_1, \ldots, w_d]$ to the output node. A bias input $x_0 = 1$ with an associated weight w_0 may be used. The output node first computes a weighted sum of its inputs $\overline{W} \cdot \overline{X}$, then it applies an *activation function* g to calculate the output $\hat{y} = g(\overline{W} \cdot \overline{X}) = g\left(\sum_{i=0}^{d} w_i x_i\right)$.

Figure 10 – Perceptron architecture.

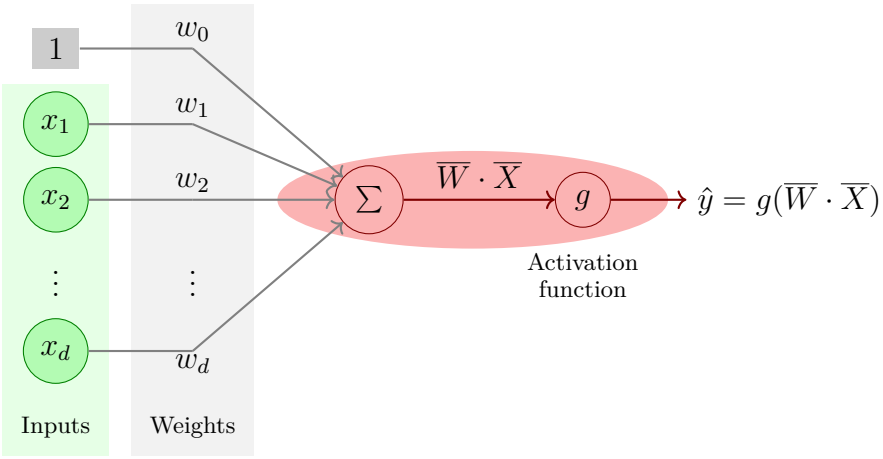

Source: Adapted from Aggarwal *et al.* (2018)

The training of a perceptron involves adjusting its weights to minimize errors in its predictions. *Gradient descent* is a popular optimization algorithm used for this purpose. Initially, the weights of the perceptron are set randomly or initialized using certain strategies. Then, for each training example, the perceptron makes a prediction $\hat{y}$ based on the input features and its current weights. The error $E(\overline{X})$ between the predicted output $\hat{y}$ and the actual output y is calculated using a *loss function*. The objective is minimize the loss function by using its gradient with respect to the weights:

$$\overline{W} \leftarrow \overline{W} + \alpha E(\overline{X})\overline{X} \tag{2.3}$$

The learning algorithm repeatedly cycles through all the training examples in random order and iteratively adjusts the weights until convergence is reached. Each such cycle is referred to as an *epoch*. The parameter α is the learning rate of the neural network (Aggarwal *et al.*, 2018).

Multilayer neural networks contain multiple computational layers; the additional intermediate layers (between input and output) are called hidden layers because the computations performed are not visible to the user. The specific architecture of multilayer neural networks is called *feed-forward* networks because successive layers feed into one another in the forward direction from input to output. The default architecture of feed-forward networks assumes that all nodes in one layer are connected to those of the next

layer. Therefore, the neural network's architecture is almost fully defined once the number of layers and the number/type of nodes in each layer have been specified (Aggarwal *et al.*, 2018). Figure 11 shows an example of a feed-forward network.

Figure 11 – Example of feed-forward network architecture

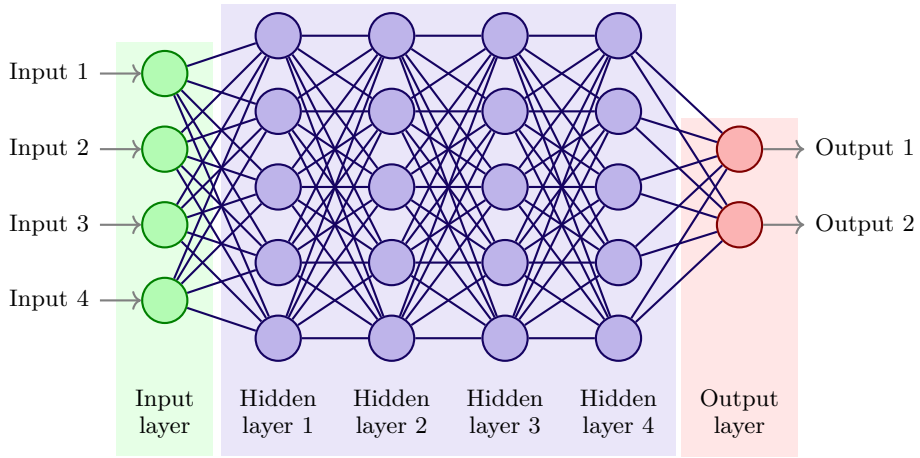

Learning in neural networks is accomplished through *backpropagation*. During the forward phase, with fixed weights, the signal is propagated through the network, layer by layer, until reaching the output. Changes only occur in the activation potentials and outputs of the neurons in the network. In the backward phase, an error signal is produced by comparing the desired output with the obtained output. This error is then retropropagated through the network, layer by layer. Adjustments are made to the synaptic weights of the network based on this error signal, allowing the network to learn from its mistakes and improve its performance over time (Bittencourt, 1998).

Previously, conventional machine-learning methods struggled to handle raw natural data effectively. Building pattern-recognition systems necessitated meticulous engineering and substantial domain expertise to construct feature extractors, converting raw data (e.g., pixel values in images) into suitable internal representations or feature vectors for subsequent processing by learning subsystems. Representation learning has emerged as a solution, enabling machines to automatically discern necessary representations from raw data for detection or classification. *Deep Learning* techniques represent a sophisticated form of representation learning, incorporating multiple layers of representation, each obtained by combining simple yet nonlinear modules. These modules progressively transform input from raw form to increasingly abstract representations, facilitating the learning of complex functions (LeCun; Bengio; Hinton, 2015).

A Convolutional Neural Network (CNN) is a type of deep, feedforward network that proved notably easier to train and exhibited superior generalization compared to networks with full connectivity between adjacent layers (LeCun; Bengio; Hinton, 2015). Its architecture, depicted in Figure 12, operates similarly to a conventional feed-forward neural network, with the key distinction that its layer operations are spatially organized and

feature sparse, deliberately designed connections between layers. The initial stages consist of two primary types of layers: convolutional layers and pooling layers. In convolutional layers, units are arranged in feature maps, where each unit is linked to local patches in the preceding layer's feature maps via a set of weights known as a *filter* bank (LeCun *et al.*, 1989).

Figure 12 – Example of convolutional neural network architecture

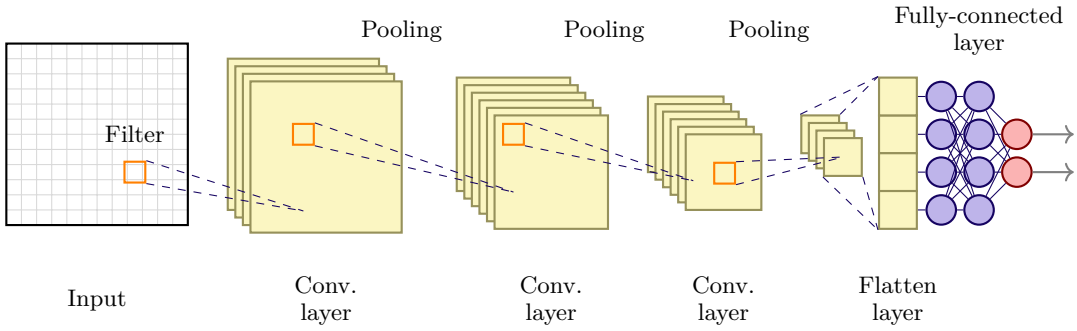

The process of arranging the units in features maps is achieved by convolution. It involves applying a filter (also known as a kernel or weight) to the input data through a sliding window, producing a feature map. Figure 13 shows an example of a input I with size 7x7, a 3x3 filter K and the corresponding convolution result $I * K$.

Figure 13 – Example of 2D convolution operation

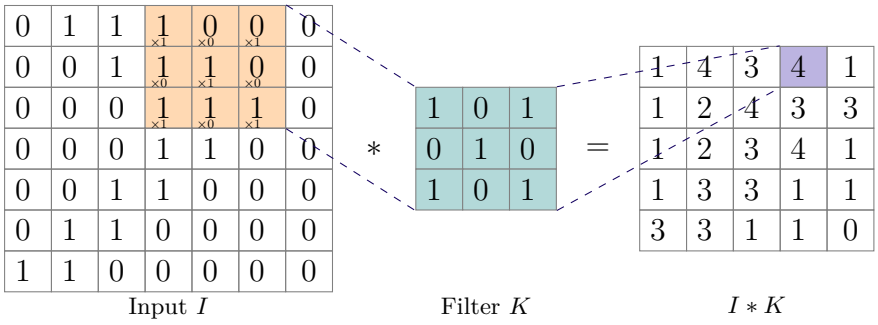

Pooling is a down-sampling technique, typically inserted between convolutional layers. It involves sliding a small window over each feature map independently, selecting either the maximum or average value within the window at each position. This process condenses the learned features effectively, making representations more robust and computationally efficient.

2.3 Related studies

During the initial stages of life, children experience rapid changes and growth in their fingers, presenting an additional challenge known in the literature as the *ageing effect*. This phenomenon describes variations in fingerprint quality and recognition performance as the time gap between the reference fingerprint sample, typically obtained in infancy,

and the probe fingerprint sample collected at later stages, increases(Galbally; Haraksim; Beslay, 2018a; Galbally; Haraksim; Beslay, 2018b). The literature presents studies related to measuring the impact of the ageing effect on authentication and identification rates.

Jain *et al.* (2016b) performed two experiments regarding measuring the ageing effect. In the first experiment, children from 0 to 5 years old (average age of two years old) were split into 3 groups: 0 to 6 months, 6 to 12 months, and greater than 12 months old. For each of these groups, three time lapses were tested between the enrollment fingerprint and the probe: six months, ten months, and one year. The study found no changes when varying the timelapse for the older group (greater than 12 months old). However, when the time lapse increases, the authentication and identification rates decrease for the two younger groups. In the second experiment, Jain *et al.* (2016b) increased the scanner resolution, and while the results were better, the overall behavior from the first experiment was observed again. Galbally, Haraksim and Beslay (2018b) conducted research with children from five to twelve years old. The authors increased the time lapse between enrollment and probe fingerprint from 0 (no time lapse) to seven years. Even though the children in the study are older than newborns and do not represent the inherent challenges of newborns' fingerprints, it is still possible to observe that the results reinforce the findings of Jain *et al.* (2016b). Finally, Engelsma *et al.* (2021) studied the biometric response of children from two to three months old being submitted to time lapses of three, nine, and twelve months. The findings of previous studies were reinforced on authentication and identification rates. Table 1 shows the numerical results from the studies, both for authentication (TAR at FAR 0.1%) and identification (Rank 1).

A current gap in research lies in the absence of techniques capable of accurately replicating the growth of fingerprints over time to mitigate the ageing effect, illustrated by the above-mentioned studies. Four studies are known in the literature for trying to narrow this gap. Jain *et al.* (2016b) illustrates that fixed scaling factors can be used to adjust the size of the fingerprint and, as a consequence, the minutiae maps considering children's age and the scanner resolution. However, this approach was shown suboptimal by Camacho *et al.* (2017) that proposed a varying scaling factor based on the children's age.

Since the above-mentioned approaches are only rescaling the images (even if using different scaling factors for each age), they have the flaw of considering that a fingerprint would grow in all directions at the same rate. Another flaw is that they consider the center of the fingerprint to be centered in the image. Markert *et al.* (2019) presents a statistical model to detect anisotropy in fingerprint growth, that is, detect the difference in proximal-distal axis (height of the fingerprint) than along the medial-lateral axis (width of the fingerprint) growth. Furthermore, Haraksim, Galbally and Beslay (2019) proposed a growth model based on the fingerprint's center and rotation, which is shown in Figure 14.

The Figure depicts the relative distances D_0 and D_1 of a specific minutiae point

Table 1 – Ageing impact on authentication and identification scores.

Study	Age	Time lapse	TAR at FAR 0.1%	Rank 1
Jain *et al.* (2016b)	0-6, 6-12, >12m	6m 10m 12m	66.7, 92.8, 100 77.3, 96.2, 100 71.1, 94.9, 100	-
Jain *et al.* (2016b)	6-12, >12m	4m 6m	100, 100 98.9, 100	100, 100 99.4, 100
Jain *et al.* (2016b)	0-6m	2m 4m 6m	31.9 18 9.8	42.2 33.6 31.1
Galbally, Haraksim and Beslay (2018b)	5-12y	0y 1y 2y 3y 4y 5y 6y 7y	96 95 94 93 90 89 86 80	-
Engelsma *et al.* (2021)	2-3m	3m 9m 1y	95 90 85	95 90 90

Note: m for months and y for years.

Figure 14 – Fingerprint Growth model proposed in Haraksim, Galbally and Beslay (2019)

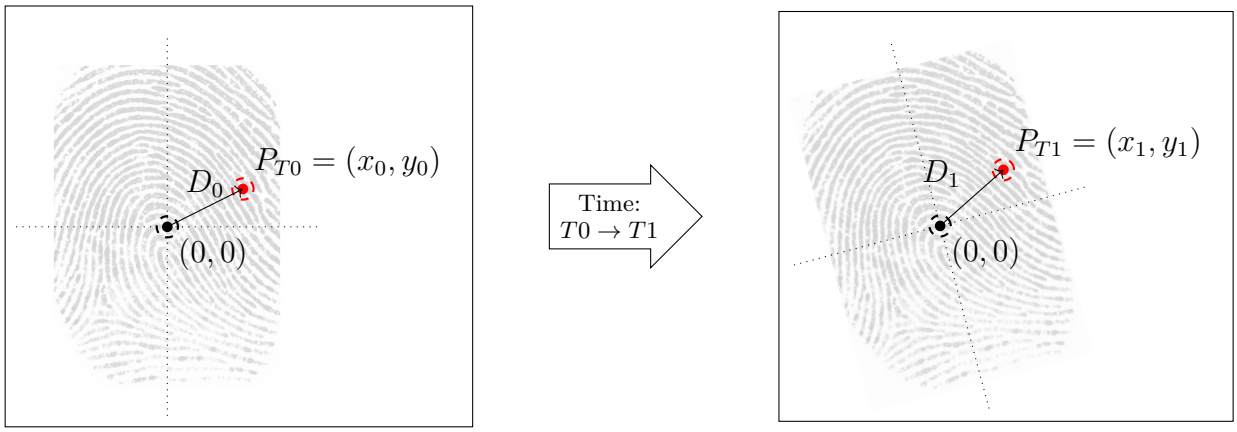

Fingerprint sample at age $T0$ Fingerprint sample at age $T1$

(red) obtained at ages T_0 and T_1, respectively, in relation to the fingerprint's center (black). Dashed circles indicate the potential location range for the centers and minutiae points due to uncertainties arising from two factors: (1) elastic distortion occurring during fingerprint acquisition via touch and (2) inaccuracies in minutiae positioning.

The proposed growth model is a mathematical function, denoted as $f(P)$, which

operates as follows: for a given minutia point $P_{T0} = (x_0, y_0)$ relative to the fingerprint's center and acquire at age $T0$, the function provides an estimate P_{T1} of the same point's location at a later age ($T1$), such that $P_{T1} = (x_1, y_1) = f(P_{T0})$. Also, $P_{T1} = f(P_{T0})$ can be expressed as $P_{T1} = (x_0 \times g_{T0}(T1), y_0 \times g_{T0}(T1))$ where $g_{T0}(T1)$ is a function that calculates the growth coefficient based on any combination of $T0$ and $T1$ (within the age limits of 5 and 16 years), enabling the computation of minutiae point displacement at any age.

While employing scaling factors has helped mitigate the effects of ageing in these related studies, this approach appears to oversimplify the issue by assuming linear fingerprint growth. Moreover, these studies solely rely on minutiae maps as input for aging emulation, overlooking the fact that young children often have low-quality minutiae maps (Camacho *et al.*, 2017), which can significantly impact the quantity and quality of minutiae. Including additional information, such as the ridge and valley configuration defining minutiae relations, could potentially enhance the emulation solution. Considering the above, the main objective of this study is to develop and evaluate an artificial intelligence solution capable of emulating the changes undergone by a newborn's minutiae map due to natural growth over time.

3 MATERIALS AND METHODS

This chapter presents the materials and methods utilized to achieve this work goal. We aimed to perform five experiments and, in each experiment, train and test a neural network model to emulate a specific growth period. The dataset used across the experiments are detailed in Section 3.1. Frame 1 shows the input and output used for the training step in each experiment.

Frame 1 – Experiments details

Experiment	Input	Output
1	Original image at 0 days + Segmented image at 0 days + Minutiae map at 0 days	Minutiae map at 7 days
2		Minutiae map at 14 days
3		Minutiae map at 1 month
4		Minutiae map at 2 months
5		Minutiae map at 3 months

For each experiment, we have as input the original fingerprint image, its segmented version, and its minutiae map. The input is the fingerprint collected on the first day of life. Respectively, in each experiment, we have trained a neural network to emulate the minutiae map at 7 days, 14 days, 1 month, 2 months, and 3 months of life. To illustrate how the fingerprints change overtime, Figure 15 shows an example of the same finger being collected in several points in time.

Figure 15 – The same fingerprint collected in different points in time

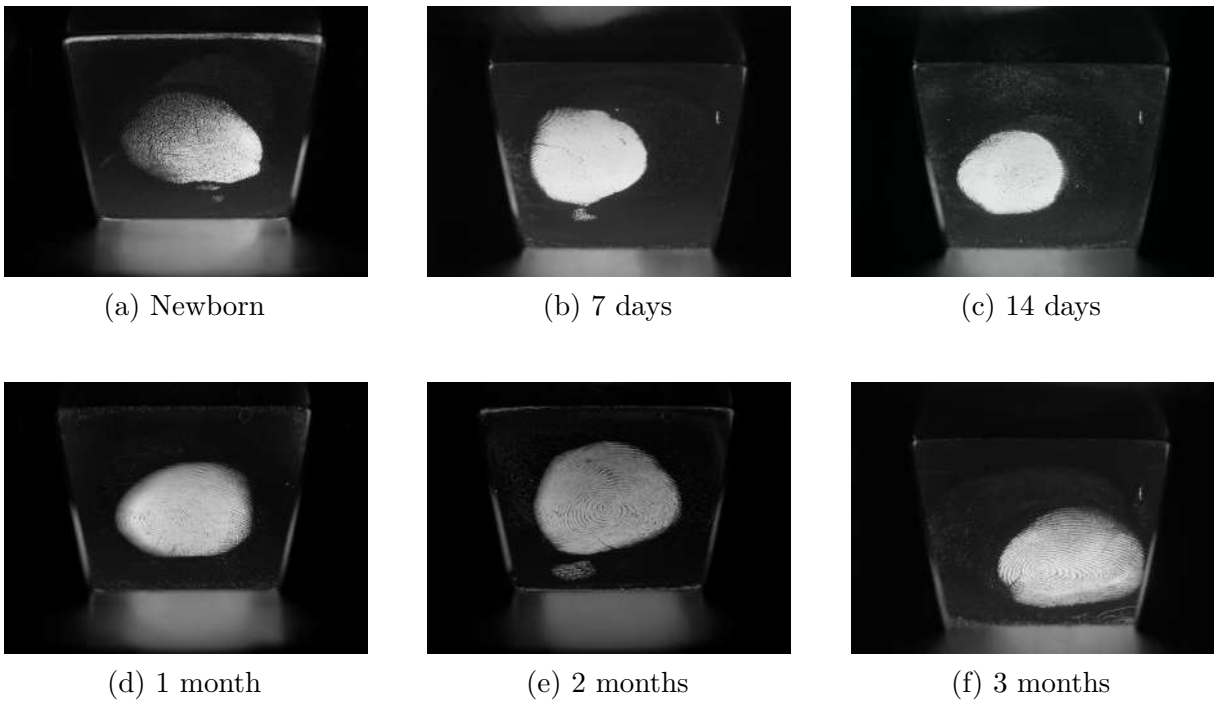

(a) Newborn (b) 7 days (c) 14 days

(d) 1 month (e) 2 months (f) 3 months

In our model, we selected the Adam optimizer due to its adaptive learning rate and momentum, which make it effective for a wide range of deep-learning tasks. For the loss function, we chose cosine similarity loss to focus on the angular distance between predicted and actual values, which is particularly useful when the direction of the target vector is more important than its magnitude. We monitored the model's performance using accuracy, a straightforward metric that provides insight into the correctness of predictions. We split our dataset into training and testing sets using an 80-20 split to ensure that our model was evaluated on unseen data, enhancing its generalizability. Additionally, we used a validation split of 20% of the training data to fine-tune the model and prevent overfitting during training. We set the batch size to 2, which, although small, allowed for more frequent updates of model parameters and reduced memory requirements. Finally, we limited training to 10 epochs, balancing the need for sufficient learning with the risk of overfitting and excessive training time.

Since our goal in each experiment was to emulate the ageing on neonatal fingerprints without other factors, we have had to hinder those factors by performing a preprocessing step referred to as *cropping*. The details of this step are detailed in Section 3.2. The inputs and outputs were submitted to the cropping step in all experiments.

For calculating the segmented version of any fingerprint, we used the segmentor provided by the partner company (InfantID, 2024). To calculate minutiae maps, we executed the MindTCT software (Ko, 2007). The model used in each experiment had its architecture detailed in Section 3.3.

3.1 Dataset

The fingerprint collection was conducted in partnership with a research compnay (InfantID, 2024). The infants had their hands sanitized and had each of their fingers pressed on a high-resolution scanner by a specialized team. Figure 16 illustrates the collection process.

The database comprises fingerprints collected from newborns from December 2023 to July 2024. The newborns had their fingerprints collected from all ten fingers on five sessions: the first day of life, 7 days, 14 days, 1 month, 2 months, and 3 months old. Thumbs and index fingers from both hands were collected twice at each session. The dataset is composed of 1,330 images of newborns, 1,330 images of 7 days babies, 1,134 images of 14 days babies, 1,134 images of 1-month-old babies, 966 images of 2-month-old babies, and 868 images of 3-months-old babies. The fingerprints that had their quality too low, i.e., no minutiae maps could be extracted in the preprocessing step, were not considered in the experiment.

Figure 16 – Newborn submitted to the fingerprint collection process

3.2 Cropping step

The preprocessing stage of cropping, to which all images were subjected, is a crucial step designed to eliminate external factors unrelated to ageing in the minutiae maps. This process aims to ensure that the subsequent analysis focuses solely on age-related changes by addressing and mitigating various discrepancies that might occur between the input and output images. These discrepancies can include minor rotations and decentralization of the fingerprints. By standardizing these aspects, the cropping preprocessing ensures a consistent and accurate basis for analyzing the ageing effects on fingerprint minutiae, thereby enhancing the reliability and validity of the results. For cropping, we used the software provided by the partner company (InfantID, 2024).

3.3 CNN architecture

Each convolutional neural network architecture consists of two main parts: the contracting path (encoder) and the expansive path (decoder), as illustrated in Figure 17. Here is a layer-by-layer description of the net used with input consisting of three grayscale images of 500x500 pixels (original fingerprint image, its segmented version, and its minutiae map) and output of a single grayscale image of 500x500 pixels (minutiae map):

- **Block 1**:
  - Two convolutional layers with 64 filters, each of size 3×3, with ReLU activation and 'same' padding.

  - Followed by a max pooling layer with a 2×2 pool size to reduce the spatial dimensions.

- **Block 2**:
  - Two convolutional layers with 128 filters, each of size 3×3, with ReLU activation and 'same' padding.
  - Followed by a max pooling layer with a 2×2 pool size to further reduce the spatial dimensions.

- **Block 3**:
  - Two convolutional layers with 256 filters, each of size 3×3, with ReLU activation and 'same' padding.
  - Followed by a max pooling layer with a 2×2 pool size.

- **Block 4**:
  - Two convolutional layers with 512 filters, each of size 3×3, with ReLU activation and 'same' padding.
  - Followed by a max pooling layer with a 2×2 pool size.

- **Block 5**:
  - Two convolutional layers with 1024 filters, each of size 3×3, with ReLU activation and 'same' padding.
  - This block represents the deepest part of the network, capturing the most abstract features.

- **Block 6**:
  - Upsampling of the feature map from Block 5 using a 2×2 upsampling layer to increase the spatial dimensions.
  - Dimensionality adjustment for the output of Block 4 using a convolutional layer with 512 filters and 2×2 kernel size.
  - Concatenation of the upsampled output from Block 5 and the adjusted output from Block 4.
  - Two convolutional layers with 512 filters, each of size 3×3, with ReLU activation and 'same' padding.

- **Block 7**:
  - Upsampling of the feature map from Block 6 using a 2×2 upsampling layer.

  - Dimensionality adjustment for the output of Block 3 using a convolutional layer with 256 filters and 2×2 kernel size.
  - Zero padding to adjust the dimensions of the upsampled output from Block 6.
  - Concatenation of the upsampled and padded output from Block 6 and the adjusted output from Block 3.
  - Two convolutional layers with 256 filters, each of size 3×3, with ReLU activation and 'same' padding.

- **Block 8**:
  - Upsampling of the feature map from Block 7 using a 2×2 upsampling layer.
  - Dimensionality adjustment for the output of Block 2 using a convolutional layer with 128 filters and 2×2 kernel size.
  - Concatenation of the upsampled output from Block 7 and the adjusted output from Block 2.
  - Two convolutional layers with 128 filters, each of size 3×3, with ReLU activation and 'same' padding.

- **Block 9**:
  - Upsampling of the feature map from Block 8 using a 2×2 upsampling layer.
  - Dimensionality adjustment for the output of Block 1 using a convolutional layer with 64 filters and 2×2 kernel size.
  - Concatenation of the upsampled output from Block 8 and the adjusted output from Block 1.
  - Two convolutional layers with 64 filters, each of size 3×3, with ReLU activation and 'same' padding.

- **Output Layer**:
  - A final convolutional layer with 1 filter of size 1×1 and a sigmoid activation function, producing the final output of the network, typically a segmentation map or binary mask.

In total, each model has 32.772.545 trainable parameters. Figure 17 shows structurally the model used. The experiments were performed in a Google Colab[1] A100 GPU environment with a system RAM of 83.5 GB, GPU RAM of 40.0GB, and a Disk capacity of 201.2 GB.

[1] https://colab.research.google.com/

Figure 17 – CNN architecture used in the experiments

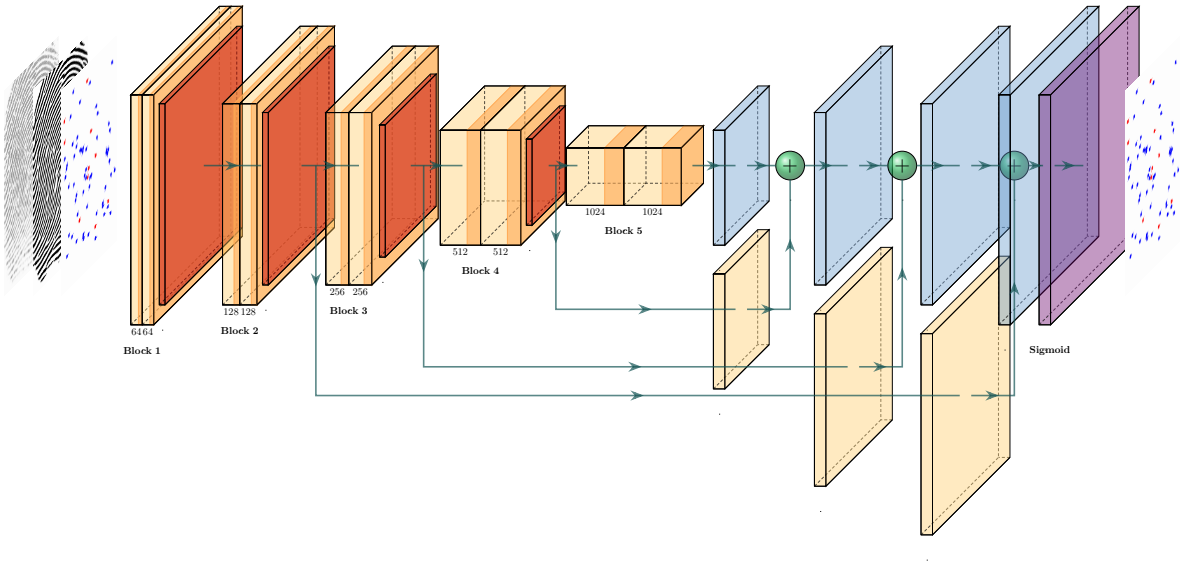

4 RESULTS AND DISCUSSION

Firstly, several images were discarded from the experiment due to their low quality, which prevented the extraction of minutiae maps. These subpar images lacked the necessary clarity and detail required for accurate minutiae detection, rendering them unsuitable for further analysis. As a result, they were excluded from the dataset to maintain the integrity and reliability of the experimental results. The final number of image pairs used in each of the five experiments is shown in Table 2.

Table 2 – Final number of pairs used in the experiment after filtering out bad quality fingerprints

Experiment	**Number of image pairs**
1	513
2	435
3	392
4	258
5	149

To illustrate the difference between high and bad-quality fingerprints, Figure 18 shows an example of each. It is possible to observe that the bad-quality segmented fingerprint has considerably less identifiable minutiae when compared with the high-quality one.

Figure 18 – Comparison of Images: Bad Quality vs. High Quality

(a) Original (Bad Quality) (b) Cropped (Bad Quality) (c) Segmented (Bad Quality)

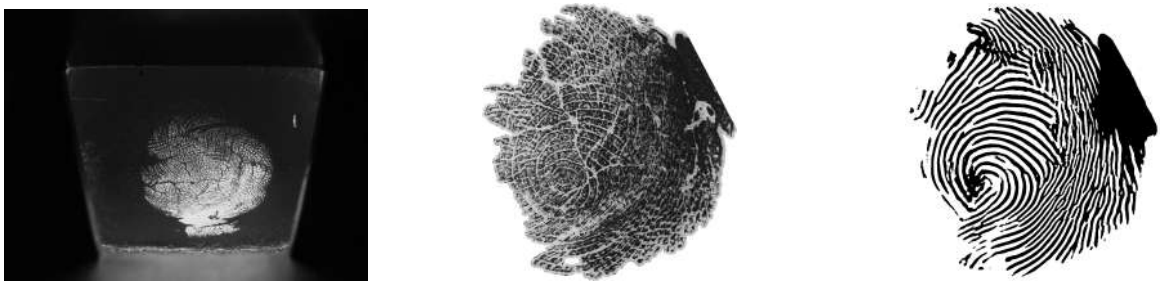

(d) Original (High Quality) (e) Cropped (High Quality) (f) Segmented (High Quality)

By performing the experiments as described in the Chapter 3, the accuracy for testing and validation shown in Figure 19 was obtained. Despite the accuracy being close to one across all experiments, these results are not appropriate or meaningful. The models are generating predominantly white images with a few random black dots representing minutiae. This superficial pattern is easy for the model to replicate, leading to artificially high accuracy scores. However, this metric fails to capture the true quality and relevance of the generated images. The high accuracy does not reflect the model's ability to produce realistic or useful representations of fingerprint growth, highlighting that accuracy alone is not a sufficient measure of success in this context.

Figure 19 – Accuracy of testing and validation

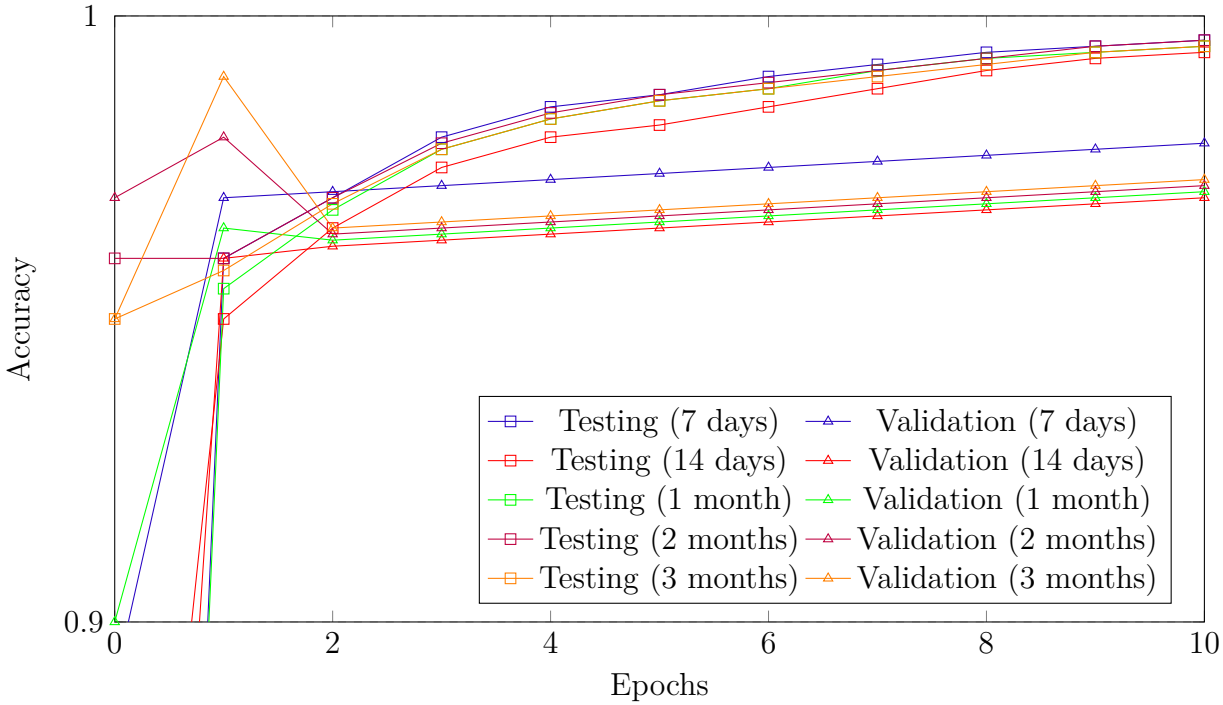

Using convolutional neural networks (CNNs) with images to emulate the growth of newborn fingerprints may not be the most effective approach due to the inherent loss of information during the conversion process. When converting minutiae maps (xyt) into images and then back from images to minutiae maps, vital data can be lost, leading to reduced accuracy and representation power. While the image dataset provides a broad and powerful representation, this may not translate effectively back into the xyt format, potentially hindering the model's ability to capture the intricate details necessary for accurate fingerprint analysis.

In this experiment, the choice of loss function is critical, as it must satisfy TensorFlow's requirements for differentiability to enable effective training. Ideally, a loss function tailored to matching techniques, such as Bozorth3 or similar methods, would be employed to better align with the nature of fingerprint matching. These techniques are specifically designed to compare and score the similarity between fingerprint minutiae sets, making

them more suitable for this task. However, integrating such matching-based loss functions while ensuring they remain differentiable presents a significant challenge that needs to be addressed for improved performance. To illustrate this, Figure 20 shows three examples of predicted and expected minutiae maps. When comparing examples 1, 2 and 3 individually by using the cosine similarity between pixels, since the predicted and expected images in each example have most pixels white, the similarity is around 99%. That could explain why all models have reached around 99% accuracy but the results are not meaningful.

Figure 20 – Comparison of Images: predicted and expected minutiae maps

(a) Expected output 1

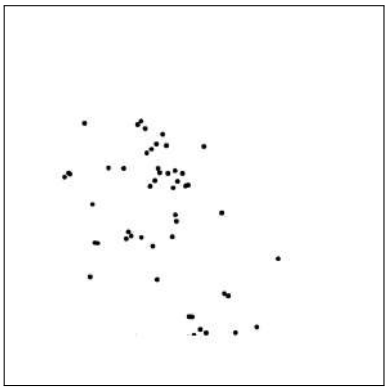

(b) Expected output 2

(c) Expected output 3

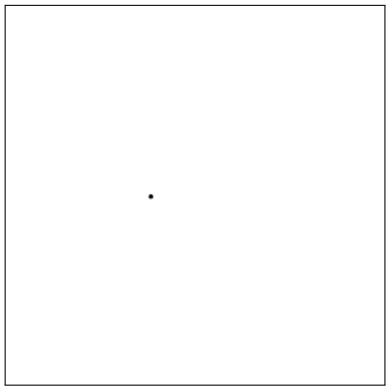

(d) Predicted output 1

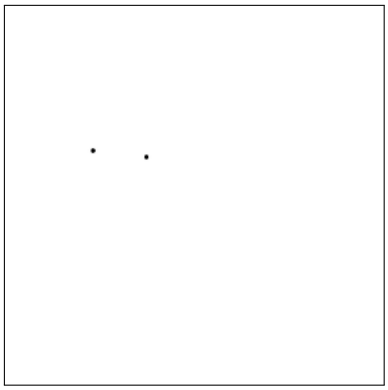

(e) Predicted output 2

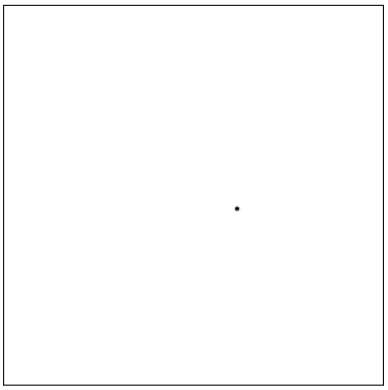

(f) Predicted output 3

Furthermore, the preprocessing techniques employed, such as cropping and rotating fingerprints, may not adequately isolate the aging behavior of the fingerprints. To ensure that the same region of the finger is consistently presented in pairs X and Y, an alignment technique would be more appropriate. This would help control for variations that are unrelated to aging and focus the model's learning on the actual growth and changes in the fingerprints over time. Without precise alignment, the model may be inadvertently learning from irrelevant differences, thus diluting the effectiveness of the experiment.

The quantity of samples used for training and testing in this experiment is another critical factor. While obtaining fingerprint data from newborns is understandably challenging, increasing the sample size could significantly enhance the experiment's outcomes. A

larger dataset would provide more varied examples, potentially leading to more robust and generalized model performance. The limited sample size, as shown in Table 2, currently restricts the model's ability to learn diverse patterns of fingerprint growth, which may result in overfitting or an inability to generalize to new data.

The quality of the collected fingerprints also plays a pivotal role in the success of this experiment. Despite having a large number of fingerprints initially, a substantial portion was discarded during preprocessing due to quality issues. This not only reduced the available data but also potentially excluded important variations that could have contributed to the model's learning. Ensuring higher quality in the collected fingerprints, possibly through improved scanning techniques or stricter quality controls, would lead to a more reliable dataset and better experimental results. The impact of quality can be seen on Figure 18.

In summary, while the experiment aims to emulate the growth of newborn fingerprints using CNNs, all these critical challenges must be addressed.

5 CONCLUSION

Each year, millions of children are born worldwide, leading to a significant number of young children in need of improved survival strategies. Despite substantial progress in recent decades, enhancing child survival remains a critical challenge. Effective interventions like vaccinations, proper nutrition, and protection programs are crucial for preserving young lives. However, their success depends heavily on robust identification systems. The lack of efficient monitoring and identification systems impacts vaccination coverage and contributes to preventable deaths. Biometric systems for neonatal identification could address fraud in nutrition programs, child protection, and national identity management. Conventional biometric methods struggle with the unique characteristics of neonatal skin and fingers. Rapid growth and changes in body composition affect fingerprint quality and stability, making recognition difficult. This study proposes developing an artificial intelligence model to simulate changes in newborn fingerprints over time, aiming to provide a more accurate representation of child development by addressing the limitations of previous methods that relied on scaling factors.

In conclusion, the experiment highlights several challenges in using CNNs to emulate the growth of newborn fingerprints, particularly the loss of information during the conversion between minutiae maps and images. This loss undermines the effectiveness of the image-based approach, suggesting that direct use of xyt data might yield more accurate results. Additionally, the choice of loss function is critical, requiring a balance between alignment with fingerprint matching techniques and TensorFlow's differentiability constraints. The current preprocessing methods, which include cropping and rotating, may not effectively isolate fingerprint aging behavior, necessitating more precise alignment techniques. The experiment's outcomes are also impacted by the limited sample size and the quality of collected fingerprints, both of which restrict the model's ability to generalize and learn diverse growth patterns. Future work should focus on increasing the dataset size, enhancing fingerprint quality, and exploring alternative methods that utilize xyt data directly to improve the accuracy and reliability of fingerprint growth emulation.

REFERENCES

AGGARWAL, C. C. *et al.* **Neural networks and deep learning**. [*S.l.: s.n.*]: Springer, 2018. v. 10.

BITTENCOURT, G. **Inteligência artificial: ferramentas e teorias**. [*S.l.: s.n.*], 1998.

CAMACHO, V. *et al.* Recognizing infants and toddlers over an on-production fingerprint database. *In*: **2017 International Conference of the Biometrics Special Interest Group (BIOSIG)**. Darmstadt, Germany: IEEE, 2017. p. 1–5.

ENGELSMA, J. J. *et al.* Infant-id: Fingerprints for global good. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, IEEE, v. 44, n. 7, p. 3543–3559, 2021.

ENGELSMA, J. J. *et al.* Infant-prints: Fingerprints for reducing infant mortality. **CoRR**, abs/1904.01091, p. 67–74, 2019.

FINLAYSON, L. *et al.* Depth penetration of light into skin as a function of wavelength from 200 to 1000 nm. **Photochemistry and Photobiology**, Wiley Online Library, v. 98, n. 4, p. 974–981, 2022.

GALBALLY, J.; HARAKSIM, R.; BESLAY, L. Fingerprint quality: A lifetime story. *In*: **2018 International Conference of the Biometrics Special Interest Group (BIOSIG)**. Darmstadt, Germany: IEEE, 2018. p. 1–5.

GALBALLY, J.; HARAKSIM, R.; BESLAY, L. A study of age and ageing in fingerprint biometrics. **IEEE Transactions on Information Forensics and Security**, IEEE, v. 14, n. 5, p. 1351–1365, 2018.

GRANTHAM-MCGREGOR, S. *et al.* Effects of integrated child development and nutrition interventions on child development and nutritional status (vol 1308, pg 11, 2014). **ANNALS REPORTS**, BLACKWELL SCIENCE PUBL OSNEY MEAD, OXFORD OX2 0EL, ENGLAND, v. 1309, p. 63–63, 2014.

HARAKSIM, R.; GALBALLY, J.; BESLAY, L. Fingerprint growth model for mitigating the ageing effect on children's fingerprints matching. **Pattern Recognition**, Elsevier, v. 88, p. 614–628, 2019.

INFANTID. 2024. Https://natosafe.com.br/, Access: 22.04.2024.

JAIN, A.; BOLLE, R.; PANKANTI, S. **Introduction to biometrics**. [*S.l.: s.n.*]: Springer, 2011.

JAIN, A.; HONG, L.; PANKANTI, S. Biometric identification. **Communications of the ACM**, ACM New York, NY, USA, v. 43, n. 2, p. 90–98, 2000.

JAIN, A. K. *et al.* Giving infants an identity: Fingerprint sensing and recognition. *In*: **Proceedings of the Eighth International Conference on Information and Communication Technologies and Development**. New York, NY, United States: Association for Computing Machinery, 2016. p. 1–4.

JAIN, A. K. *et al.* Fingerprint recognition of young children. **IEEE Transactions on Information Forensics and Security**, IEEE, v. 12, n. 7, p. 1501–1514, 2016.

JAIN, A. K.; CAO, K.; ARORA, S. S. Recognizing infants and toddlers using fingerprints: Increasing the vaccination coverage. *In*: **IEEE international joint conference on biometrics**. Clearwater, FL, USA: IEEE, 2014. p. 1–8.

JAIN, A. K.; KUMAR, A. **Biometric recognition: an overview**. Berlin, Germany: Springer, 2012. 49–79 p.

JAIN, A. K.; ROSS, A.; PRABHAKAR, S. An introduction to biometric recognition. **IEEE Transactions on circuits and systems for video technology**, IEEE, v. 14, n. 1, p. 4–20, 2004.

KO, K. **User's Guide to NIST Biometric Image Software (NBIS)**. [*S.l.: s.n.*]: NIST Interagency/Internal Report (NISTIR), National Institute of Standards and Technology, Gaithersburg, MD, 2007.

LECUN, Y.; BENGIO, Y.; HINTON, G. Deep learning. **nature**, Nature Publishing Group UK London, v. 521, n. 7553, p. 436–444, 2015.

LECUN, Y. *et al.* Handwritten digit recognition with a back-propagation network. **Advances in neural information processing systems**, v. 2, 1989.

LI, S. Z. **Encyclopedia of Biometrics: I-Z.** Berlin, Germany: Springer Science & Business Media, 2009. v. 1.

LINDOWER, J. B. Water balance in the fetus and neonate. *In*: **Seminars in Fetal and Neonatal Medicine**. Amsterdam, Netherlands: Elsevier, 2017. v. 22, p. 71–75.

MARKERT, K. *et al.* Detecting anisotropy in fingerprint growth. **Journal of the Royal Statistical Society Series C: Applied Statistics**, Oxford University Press, v. 68, n. 4, p. 1007–1027, 2019.

PHILLIPS, P. J. *et al.* An introduction evaluating biometric systems. **Computer**, IEEE, v. 33, n. 2, p. 56–63, 2000.

RUSSELL, S. J.; NORVIG, P. **Artificial intelligence a modern approach**. [*S.l.: s.n.*]: London, 2010.

SCHNEIDER, J. Quantifying the dermatoglyphic growth patterns in children through adolescence. **National Institute of Justice December**, v. 1, p. 1, 2010.

SEDLAK, A. **National estimates of missing children: An overview**. [*S.l.: s.n.*]: US Department of Justice, Office of Justice Programs, Office of Juvenile . . . , 2002.

SHARMA, A.; FORD, S.; CALVERT, J. Adaptation for life: a review of neonatal physiology. **Anaesthesia & Intensive Care Medicine**, Elsevier, v. 12, p. 85–90, 2011.

SIRIRUNGSI, W. *et al.* Early infant hiv diagnosis (eid) and entry to hiv care cascade: seven-year experience in thailand. **The lancet. HIV**, Europe PMC Funders, v. 3, n. 6, p. e259, 2016.

SOUTHIER, L. F. P. *et al.* **Systematic Literature Review on Neonatal Fingerprint Recognition**. 2023. Preprint.

UNICEF United Nations International Children's Emergency Fund. **Ending preventable child deaths: how britain can lead the way**. 2023. Available at: https://www.unicef.org.uk/wp-content/uploads/2020/01/Unicef-UK-Ending-Preventable-Child-Deaths_Report-2020.pdf.

UNICEF United Nations International Children's Emergency Fund. **Under-five mortality**. 2023. Available at: https://data.unicef.org/topic/child-survival/under-five-mortality/.

United Nations - Department of Economic and Social Affairs. **World Population Prospects 2022**. 2023. Available at: https://population.un.org/wpp/.

VISSCHER, M. O. *et al.* Newborn infant skin: physiology, development, and care. **Clinics in dermatology**, Elsevier, v. 33, n. 3, p. 271–280, 2015.

WFP United Nations World Food Programme. **WFP demands action after uncovering misuse of food relief intended for hungry people in Yemen**. 2023. Available at: https://www.wfp.org/news/wfp-demands-action-after-uncovering-misuse-food-relief-intended-hungry-people-yemen.

WHO World Health Organization. **Immunization Agenda 2030: A Global Strategy To Leave No One Behind**. 2023. Available at: https://www.who.int/publications/m/item/immunization-agenda-2030-a-global-strategy-to-leave-no-one-behind.

WHO World Health Organization. **Immunization coverage**. 2023. Available at: https://www.who.int/news-room/fact-sheets/detail/immunization-coverage.