Enhancing Healthcare Data Privacy with Homomorphic Encryption

By

Mugo, John Mwangi

153032

Submitted in Partial fulfilment of the Requirements for the Degree of Master of Science in Information Systems Security at Strathmore University

School of Computing and Engineering Sciences

Strathmore University

Nairobi, Kenya

June 2025

Declaration and Approval

Declaration

I declare that this work has not been previously submitted and approved for the award of a degree by this or any other University. To the best of my knowledge and belief, the dissertation contains no material previously published or written by another person except where due reference is made in the dissertation itself.

Name: **Mugo, John Mwangi**

Signature: .....

Date: 28/03/2025

Approval

The dissertation of Mugo, John Mwangi was reviewed and approved for examination by the following:

Dr. Joseph Sevilla

Supervisor,

School of Computing and Engineering Sciences,

Director of @iLabAfrica and @iBizAfrica,

Strathmore University.

Signature: .....

Date: 29/03/2025

Abstract

Artificial intelligence in healthcare has birthed advances in fields such as medical diagnostics and personalised treatments. However, the use of patient information raises pertinent privacy issues, creating the challenge of protecting patient data privacy and confidentiality. This dissertation focused on creating a privacy-preserving artificial intelligence application for melanoma diagnosis by employing Fully Homomorphic Encryption (FHE) to bolster data confidentiality. The main objectives included examining the challenges of managing health data for AI, exploring FHE applications, developing an application that integrates FHE for confidential processing of health data, and verifying its accuracy, efficiency and privacy assurances. The methodology included conducting a literature review, requirement analysis, designing an FHE-enabled architecture using Concrete ML, system development, testing and validation. The application was developed using the Streamlit framework and TensorFlow's ResNet50 model combined with clinically pre-trained weights calibrated for melanoma detection. A minimal custom model designed for FHE compatibility was then implemented through a knowledge distillation process from the ResNet50 teacher model. This was to address the inherent computational constraints of homomorphic encryption while preserving diagnostic capabilities. Thorough testing and validation was carried out to ensure the application functions as per the requirements and achieved the set goals. A verified security expert attempted to access unencrypted patient data during FHE operations and confirmed that the data remained securely encrypted throughout the processing pipeline. The application offered accurate melanoma detection while ensuring verifiable end-to-end encryption of patient data throughout the process in compliance with healthcare privacy laws. This demonstrates FHE's potential in facilitating the adoption of privacy-preserving AI in healthcare and other data-sensitive fields.

KEY WORDS: Artificial Intelligence, Concrete ML, Confidentiality, Fully Homomorphic Encryption, Healthcare, Machine Learning, Melanoma, Privacy, ResNet50, Streamlit.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AI - Artificial Intelligence

ANN - Artificial Neural Network

CEH - Certified Ethical Hacker

CNN - Convolutional Neural Networks

CPU - Central Processing Unit

DOFHE - Delay Optimised Fully Homomorphic Encryption

DPA - Data Protection Act

ECC - Elliptic Curve Cryptography

FHE – Fully Homomorphic Encryption

FPGAs – Field Programmable Gate Arrays

FV HE – Fan-Vercauteren (FV) Homomorphic Encryption

GDPR - General Data Protection Regulation

GM - Goldwasser-Micali

GPU - Graphics Processing Unit

HE - Homomorphic Encryption

HELM - Hardware Emulation Language for Modelling

HDL - Hardware Description Language

HIPAA - Health Insurance Portability and Accountability Act

IV - Initialisation Vector

LA - Lion Algorithm

LWE - Learning with Errors

ML – Machine Learning

NLP - Natural Language Processing

PDF - Portable Document Format

PHI – Protected Health Information

PPRL – Privacy Preserving Record Linkage

RAD – Rapid Application Development

RCSD - Robust Cramer Shoup Decryption

RCSDOFH – Robust Cramer Shoup Delay Optimised Fully Homomorphic

RL – Reinforcement Learning

RLWE - Ring Learning With Errors

RPA - Robotic Process Automation

RSA - Rivest, Shamir, Adleman

STHE - Statistical Transformation with Homomorphic Encryption

WSL - Windows Subsystem for Linux

Definition of Terms

Concrete ML	This is an open-source library that allows the creation of Fully Homomorphic Machine Learning models. It facilitates the conversion of ML models into their FHE equivalents. (Zama.ai, 2024)
Encryption	The process of transforming data into an unintelligible format that prevents unauthorised parties from comprehending it even if they gain access to it, without the appropriate decryption key. (Menezes et al., 2018).
Fully Homomorphic Encryption (FHE)	FHE is an encryption technique that enables computations to be carried out on data that is encrypted, ensuring that the results remain accurate when decrypted. It enables the processing of encrypted data without necessitating decryption. (Gentry, 2009).
Machine Learning (ML)	Machine learning is a field within computer science that entails developing models and algorithms that help computers improve their performance on tasks by learning from experience. These algorithms analyse sample data, called training data, to create models that can make predictions without needing explicit instructions for each task. (Jordan and Mitchell, 2015).
Melanoma	A type of skin cancer that starts in and affects the melanoctyes. These cells produce melanin, the pigment that gives skin its colour. (Schadendorf et al., 2018).
ResNet50	A deep convolutional neural network (CNN) with 50 layers. It is designed and suited for image recognition and classification tasks. (tensorflow.org, 2024).

Acknowledgements

First and foremost, I am deeply grateful to God for the gift of life and his abundant grace. It is by his hand that I have made it this far in my academic journey.

I am grateful to my beloved parents for their unwavering love, prayers, and sacrifices. My dear siblings and friends, thank you for always being there for me and providing a constant support system.

Special appreciation to my supervisor, Dr Joseph Sevilla, for his invaluable guidance and feedback. Your expertise and dedication have been instrumental in shaping this research.

Chapter 1
Introduction

1.1 Introduction

The rise and widespread adoption of Artificial Intelligence (AI) in the healthcare sector is reshaping how medical data is handled, diagnoses are made, treatment options are recommended, and patient care is provided. According to Davenport and Kalakota (2019), AI systems powered by machine learning algorithms can analyse large volumes of complex health information to extract valuable insights for clinical purposes. This has the potential to greatly benefit healthcare professionals by aiding them in making precise diagnostics, predicting future health issues and customising medical plans tailored to individual patients' needs (Shaheen, 2021). However, AI's success in healthcare relies heavily on access to data like medical images, genetic information, lab test results, medical histories and treatment outcomes. This reliance on health information poses challenges in ensuring robust privacy measures and maintaining patient trust (Kaissis et al., 2020).

Health-related data is highly confidential and sensitive. The personal details stored in these records require stringent security safeguards to prevent unauthorised access and misuse (Cohen and Mello, 2018). Furthermore, laws such as the Data Protection Act (DPA) in Kenya, the General Data Protection Regulation (GDPR) in Europe and the Health Insurance Portability and Accountability Act (HIPAA) in the USA, regulate how medical data should be handled ethically and place restrictions on its usage (Bajwa et al., 2021). DPA covers various areas of data protection to enforce the secure handling of personal data, including medical data (Kenyalaw, 2019). HIPAA focuses on Protected Health Information (PHI) overseen by entities like insurers and healthcare providers in the US while GDPR has a reach covering all data of EU citizens with special emphasis on data that is sensitive for instance, health information (Cohen and Mello, 2018; Voigt and Von dem Bussche, 2017).

Although DPA, HIPAA and GDPR aim to protect patient privacy, they introduce complexities when it comes to using data for healthcare-focused AI applications. Different entities need to navigate these regulations by obtaining explicit consent, ensuring data minimisation and

purpose limitation, and implementing robust security measures (Voigt and Von dem Bussche, 2017). Non-compliance may lead to reputation damage and even fines. The privacy challenges associated with these applications relying on health data create a hurdle for their adoption. Patients might hesitate to divulge their information if they don't trust how it will be handled. The potential consequences of data breaches and misuse are severe. Unauthorised disclosure of diagnoses, such as cancer, could result in stigma and discrimination. Therefore it is crucial to develop methods that allow medical AI systems to utilise health data while safeguarding end-to-end privacy (Bajwa et al., 2021).

Homomorphic encryption (HE) has risen to become a promising technology that enables processing of encrypted data without necessitating its decryption (Gentry, 2009). This facilitates obtaining insights from data while upholding confidentiality (Acar et al., 2018). Integrating encryption into healthcare AI processes could enable systems to analyse personal health information while protecting patient privacy (Alowais et al., 2023). However further research is necessary to determine the integration strategies and evaluate real-world feasibility and performance. This dissertation aims to investigate the potential of combining AI with privacy-preserving encryption in the field of melanoma diagnostics.

1.2 Problem Statement

The utilisation of AI technology in the healthcare sector presents a dilemma: how to fully leverage the advantages offered by AI in terms of data analysis, improved diagnostics, personalised insights and better care, all while safeguarding the vast amount of highly sensitive patient information from privacy breaches, unauthorised access, cyber threats and data leaks. Maintaining the confidentiality and security of healthcare information is essential and of utmost importance. However, integrating AI deeply while maintaining security measures poses technological and regulatory challenges. Moreover, medical data is governed by laws like DPA, HIPAA, and GDPR that restrict its usage creating obstacles for AI applications that rely on this data (Cohen and Mello, 2018; Kenyalaw, 2019; Voigt and Von dem Bussche, 2017). Balancing the utilisation of AI's potential in healthcare with protecting privacy and data confidentiality is a task that requires attention. Hence innovative solutions are essential to enable AI to handle medical information while respecting patient privacy.

1.3 Research Hypothesis

The successful development of an AI melanoma diagnostic system utilising homomorphic encryption will maintain the confidentiality of sensitive patient health data while providing accurate detection. The system will guarantee the security of encrypted patient data throughout the diagnostic process by allowing its analysis without decrypting it.

1.4 General Objective

To develop an artificial intelligence system for melanoma diagnostics that uses homomorphic encryption to ensure the confidentiality of patient health data.

1.5 Research Objectives

1. To analyse the challenges in handling and processing sensitive patient data required for AI applications in healthcare.
2. To explore existing research on applying homomorphic encryption in artificial intelligence.
3. To design, develop and test an AI system employing homomorphic encryption to uphold the confidentiality of patient health data.
4. To validate the effectiveness of the developed solution in upholding the confidentiality of sensitive patient health information.

1.6 Research Questions

1. What are the key challenges in handling and processing sensitive patient data needed in healthcare-focused AI applications?
2. What current research efforts and solutions utilise homomorphic encryption in AI?
3. How can an effective artificial intelligence healthcare system employing homomorphic encryption be designed, developed, and tested to maintain the confidentiality of patient data?
4. To what extent can the developed AI system uphold the confidentiality of sensitive patient health information?

1.7 Scope

The research focuses on developing an artificial intelligence system for melanoma diagnostics using homomorphic encryption techniques. The evaluation will assess the solution's confidentiality, data integrity, and performance.

1.8 Limitations

Potential reluctance by healthcare providers to share real patient data due to concerns about data privacy, security risks, and liability despite assurances of maintaining confidentiality and adhering to governing regulations. This data is integral for system development and testing.

1.9 Justification

This research investigates the potential of combining artificial intelligence and homomorphic encryption techniques to develop a melanoma diagnostic system that maintains patient data privacy while still providing accurate results. Successfully demonstrating such a system could help address ethical AI concerns in healthcare by enabling providers to fully utilise AI's benefits without compromising confidentiality obligations. This would offer regulators the confidence to develop policies that support AI integration. Patients would benefit from privacy assurances, which would help them trust in the security of their information. In summary, this study aims to enhance technologies that protect privacy in order to maximise the adoption and use of AI in healthcare while ensuring data confidentiality.

Chapter 2
Literature Review

2.1 Introduction

This chapter covers past literature on enhancing data privacy using homomorphic encryption in different settings. The chapter involves a discussion of different theoretical frameworks and perspectives that are relevant to this area of knowledge, the empirical studies that try to establish the existing association between the study variables, and the research gaps. Through reviewing a wide range of literature, this chapter seeks to give a thorough overview of the topic and to help contextualise new research within the existing body of knowledge.

2.2 Theoretical Framework

This study's theoretical framework focuses on various homomorphic encryption (HE) systems that have piqued the interest of researchers in the cryptography field. It begins with a review of the Fully Homomorphic Encryption (FHE) Model by Gentry (2009). The study also delves into the system introduced by Goldwasser (1984) and further explores other systems, like The HEAWS Model, Reinforcement Learning (RL) Model, Robust Cramer Shoup Delay Optimised Fully Homomorphic (RCSDOFH) and the Ring Learning With Errors (RLWE).

2.2.1 Fully Homomorphic Encryption (FHE) Model

Gentry (2009) developed the Fully Homomorphic Encryption (FHE) Model. The core concept of FHE enables arbitrary computations to be carried out on data that is encrypted, ensuring that the results remain accurate when decrypted (Menon and Wu, 2022). This principle enables data to be processed and analysed without revealing the initial content, thereby facilitating complex data processing tasks while maintaining privacy (Bonte et al., 2022).

The theory of fully homomorphic encryption (FHE) proposes that by using advanced encryption schemes that support both multiplication and addition operations on data that is encrypted, data privacy can be enhanced even further. FHE enables processing data that is encrypted without compromising its confidentiality, increasing the level of security for sensitive information. In a mathematical context, if a public key, pk, is used to encrypt a plain text message, x, to create a

ciphertext c, then:

$$c = \text{Enc}_{pk}(x) \tag{2.1}$$

For instance, in the realm of large-scale dialogue models, the goal might be to calculate language model scores for a user without disclosing their input.

Considering the user's input as x and the language model parameters as f, the goal is to calculate the score $f(x)$. Using Fully Homomorphic Encryption (FHE), the input, x, can be encrypted on the user's device to generate $c = \text{Enc}_{pk}(x)$. The computation under encryption is then performed on the server to derive the encrypted score of $\text{Enc}_{pk}(f(x))$, which is later transmitted back to the user's device for decryption.

Figure 2.1 depicts Fully Homomorphic Encryption.

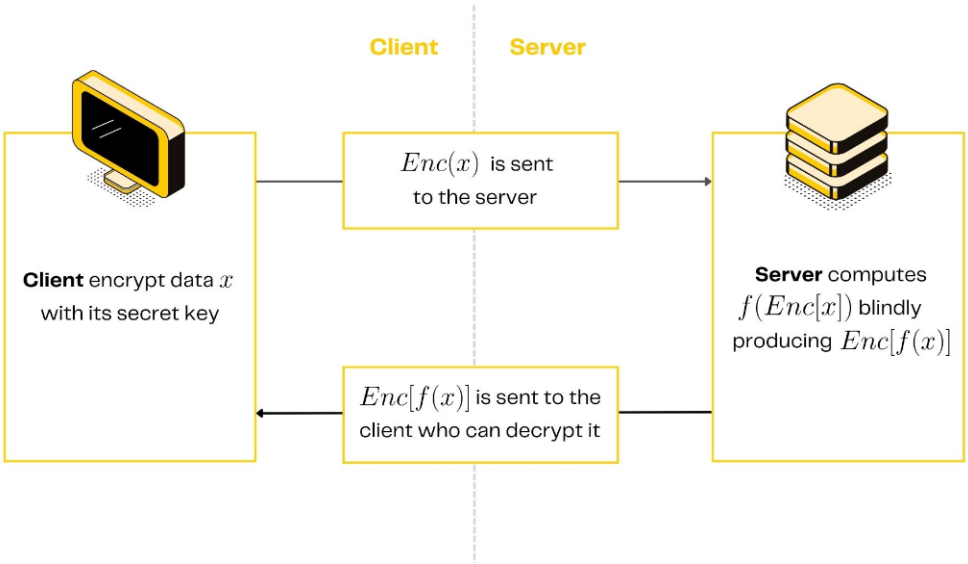

Figure 2.1: Fully Homomorphic Encryption (FHE) Depiction.

Source: Zama.ai (2024)

Despite its potential, Fully Homomorphic Encryption (FHE) faces several challenges, including significant computational and memory overhead that makes it slow for real-time or resource-intensive applications. Its complexity and large ciphertext and key sizes introduce inefficiencies in storage, transmission, and key management. Additionally, FHE's broad support for computa-

tions is countered by inefficiencies in specific operations, and the lack of established standards complicates its implementation and integration (Cao and Liu, 2015).

However, this technology presents a promising opportunity for enhancing privacy. Despite current challenges in computational efficiency, it is anticipated that with advancements in algorithms and hardware, FHE will see widespread adoption. This model represents a significant advancement in the integration of privacy-preserving technologies within dialogue systems, offering enhanced technical assurance for user data security (Zhang et al., 2023). In this study, FHE will be used to enhance healthcare data privacy by allowing secure processing and analysis of sensitive medical information without decryption.

2.2.2 Goldwasser-Micali Scheme

This encryption scheme created by Goldwasser and Micali in 1982 is of significance since it has influenced many subsequent approaches to homomorphic encryption. Similar to the RSA algorithm, this technique involves the computations: $n = p \cdot q$, where n is the product of two large primes, p and q (Rivest et al., 1977). Encryption under this scheme is relatively simple requiring a product and a square operation while decryption involves exponentiation and is computationally demanding. The decryption process is intricate with the time complexity of $O(k \cdot l(p)^2)$ where $l(p)$ denotes the bit length of p (Goldwasser, 1984).

However, one drawback of this scheme is its limitation to encrypt a bit at a time which can lead to inefficiencies when encrypting multiple bits. Encrypting k bits would thus incur a cost of $O(k \cdot l(p)^2)$ which may not be very efficient despite its practicality. Furthermore, encrypting a single bit results in the ciphertext becoming significantly larger causing complications with the process (Sen, 2013). The Goldwasser-Micali (GM) scheme is also criticised by Maimuţ and Teşeleanu (2020) because it is computationally intensive resulting in slow encryption and decryption processes. The model also produces large ciphertexts leading to inefficiencies in storage and transmission. The model also lacks inherent semantic security, meaning it does not guarantee that ciphertexts reveal no additional information about plaintexts beyond their length.

Nevertheless, this encryption scheme is critical in ensuring the security and confidentiality of data while enabling various operations on data that is encrypted without necessitating decryption. Therefore, it can be used as a tool for enhancing data privacy in relation to health information.

2.2.3 The HEAWS Model

Turan et al. (2020) designed a novel methodology for accelerating functions that are homomorphic estimation on data that is encrypted using a domain-centric coprocessor system known as HEAWS. They also developed an efficient coprocessor system for the FV HE technique, leveraging the large-scale capabilities of Amazon FPGAs. Furthermore, they introduced an artificial neural network (ANN) for more rapid energy consumption forecasting in a smart grid context, achieving a fivefold increase in speed.

The HEAWS model is criticised for it incurs significant computational overhead due to the complex operations involved in homomorphic encryption and attribute-based access control. The integration of homomorphic encryption with attribute-based access requires sophisticated key management strategies. Homomorphic encryption typically results in larger ciphertexts compared to plaintexts which can lead to increased storage and transmission requirements, which may not be ideal for resource-constrained environments. It is also criticised for lacking widespread practical implementations and standardised frameworks. The HEAWS model is useful because it can enhance healthcare data privacy by integrating homomorphic encryption with advanced access control mechanisms (Rao et al., 2023).

2.2.4 Reinforcement Learning (RL) Model

The reinforcement learning (RL) model was brought forth by Park et al. (2020) with the aim of improving cloud systems privacy. They also introduced in the cloud systems a method for performing arithmetic operations without requiring the decryption of ciphertexts. To access RL-based applications, users were only allowed to submit ciphertexts to the cloud computing system using the Homomorphic Encryption (HE) scheme. Performance evaluation and analysis of the anticipated Privacy Preserving Record Linkage (PPRL) architecture were conducted using various intelligent service scenarios on cloud computing-based systems.

The Reinforcement Learning (RL) Model is criticised for requiring extensive computational resources and time to train due to its large number of iterations needed to explore and exploit the environment effectively. The model is also limited for it needs a vast number of interactions with the environment to learn effectively and is therefore impractical in scenarios where data collection is costly or time-consuming (Wang and Hong, 2020). However, the Reinforcement Learning (RL) model is applicable to this study for providing valuable contributions to optimising privacy-preserving systems. RL model is also useful because it can enhance the adaptability of privacy-preserving systems by enabling them to adjust encryption parameters and access controls in response to evolving threats and changing data privacy requirements

2.2.5 Robust Cramer Shoup Delay Optimised Fully Homomorphic (RCSDOFH)

Li et al. (2020) introduced a system called Robust Cramer Shoup Delay Optimised Homomorphic (RCSDOFH) to enhance privacy protection. The system involved three phases. The Robust Cramer Shoup Decryption (RCSD) method was designed to reduce communication processing time and overhead. Next, a Delay Optimised Fully Homomorphic Encryption (DOFHE) technique was introduced to minimise data latency and network delays. Additionally, this approach calculated the delivery delay between the signal transmitted by an IoT device and the base station.

The RCSDOFH scheme is limited in that it may produce larger ciphertexts compared to plaintexts like many FHE schemes leading to increased storage requirements and higher costs for data transmission. The model is also criticised for implementation complexity. Implementation of RCSDOFH can be complex due to the intricacies involved in its cryptographic components and optimisations. This complexity can lead to difficulties in ensuring correct and secure implementations, potentially introducing vulnerabilities (Li et al., 2020). The model is however applicable to this study for its optimisation for delay reduction improves the efficiency of homomorphic encryption processes, addressing one of the major concerns of FHE regarding performance overhead.

2.2.6 Ring Learning With Errors (RLWE)

Su et al. (2020) presented a Fully Homomorphic Encryption (FHE) model based on the Ring Learning with Errors (RLWE) problem. The researchers showcased a rapid implementation of the levelled FHE scheme and the design of a highly parallel architecture on an FPGA to enhance the efficiency of FHE schemes. The clock frequency was amplified through circuit and block-level pipeline strategies, thereby boosting homomorphic evaluation functions and the processing speed of polynomial multipliers to reduce computation latency and improve overall performance.

The model is limited for its complexity in implementation due to the mathematical intricacies involved. Correct and secure implementation requires careful attention to detail and expertise in lattice-based cryptography. The model is also criticised because the ciphertexts generated using RLWE can be quite large due to the nature of polynomial arithmetic and the need to include error terms, resulting in higher storage and transmission costs. The model is however relevant to this topic for its ability to support practical and secure homomorphic encryption which makes it a valuable component in developing effective privacy-enhancing technologies for healthcare (Bootland et al., 2020).

2.2.7 FHE and Data Security in Machine Learning

Concrete-ML, developed by Machine Learning experts at Zama, is an open-source library that allows the creation of Fully Homomorphic Machine Learning models (Zama.ai, 2024). This open-source library is unique in its ability to convert ML models into their FHE equivalents. This provides significant benefits for individuals and businesses looking to utilise ML capabilities while keeping their data secure. FHE operates without requiring decryption during computation, making it a promising option for application development. Concrete-ML effectively utilises this to improve its functionality and utility (Shollo et al., 2022).

The Concrete-ML base model is trained initially with unencrypted data and then transformed into a Concrete-Numpy program that can process encrypted inputs for secure and confidential computations (Iatropoulou et al., 2023). This advancement in technology, illustrated by Concrete-ML, allows for the utilisation of machine learning capabilities while safeguarding data privacy.

It signifies a significant step towards a future where the benefits of ML can be harnessed without compromising on data security (Bharati and Podder, 2022). As the use of AI and machine learning continues to grow across diverse industries, technologies like Fully Homomorphic Encryption (FHE) and tools such as Concrete-ML will play a crucial part in ensuring data security, consumer trust, and progress towards a secure digital environment. As the potential of ML continues to expand, encryption technologies will be essential in maintaining privacy and security in the era of advanced digital technologies (Bharati and Podder, 2022).

Security and privacy are closely linked concepts. Security is focused on ensuring information is available, intact, and confidential, while privacy specifically concerns the rights individuals have regarding their personal information (Bertino and Ferrari, 2017). When it comes to handling personal data, privacy is prioritised, but information security is crucial for protecting against unauthorised access to information assets. Personal data encompasses any information related to an individual, such as their name, identification details, address, social security number, and bank account information among others (Maple, 2017).

Several strategies have been proposed to address the balance between security and privacy in machine learning. The most commonly utilised technologies for privacy in ML include homomorphic encryption, secure multiparty computing, trusted execution and differential privacy (Bharati and Podder, 2022). Differential privacy is utilised to prevent the attacker from discerning which instances were utilised in constructing the target model. Secure multiparty computing and homomorphic encryption safeguard the testing and training data. Trusted execution environments employ isolation and hardware-based security to protect sensitive data and training code. However, these methods significantly increase computational demands and require customisation for different types of neural networks (Bharati and Podder, 2022).

2.3 Artificial Intelligence and Healthcare

Artificial Intelligence (AI) involves machines, especially computer systems, simulating human intelligence processes, and includes applications such as natural language processing (NLP), expert systems, machine vision and speech recognition (Zhang and Lu, 2021). Davenport and

Kalakota (2019) highlight that the growing volume of data and complexity in healthcare are driving increased use of AI in the field.

Artificial Intelligence encompasses various technologies with significant relevance to healthcare. Machine learning comprises neural networks and deep learning and is pivotal for precision medicine and image analysis, such as detecting cancerous lesions in radiology images. Natural Language Processing (NLP) aids in understanding and managing clinical documentation and patient interactions. Expert systems are used for clinical decision support based on predefined rules. Physical robots are used particularly in surgical settings to enhance precision and minimally invasive procedures. Additionally, robotic process automation (RPA) streamlines administrative tasks by mimicking human actions in information systems. Each type of AI contributes uniquely to improving diagnostic accuracy, drug discovery, patient management, and operational efficiency in healthcare (Hussain et al., 2014; Lee et al., 2018).

2.3.1 AI in Healthcare data management

Valuable information can often be overlooked amid vast quantities of data, and the difficulty in linking critical data points can impede advancements in drug development, preventive medicine, and accurate diagnosis. AI addresses these challenges by efficiently managing and integrating large volumes of data, bridging information gaps in minutes. This capability streamlines healthcare administrative processes, potentially reducing time and costs while enhancing daily operations and patient experiences (Dash et al., 2019).

2.3.2 AI in Medical Diagnosis

AI has the potential to predict and diagnose diseases more quickly than most medical professionals, as it is not affected by incomplete medical histories or large caseloads that can lead to human errors. This capability to enhance the diagnostic process is among the most promising applications of AI in healthcare (Kaur et al., 2020).

2.3.3 AI in Drug Discovery

The drug development industry faces escalating costs and extensive research hours, with each drug trial costing billions and only a small fraction reaching the market. However, advancements in technology are transforming this process. AI is accelerating drug discovery by aiding in drug

design, predicting potential side effects, and identifying the best candidates for clinical trials (Blanco-Gonzalez et al., 2023).

2.3.4 AI in patient experience

AI enhances digital communication by providing personalised health tips, appointment reminders and recommendations for next steps. Through assisting with health diagnoses, AI accelerates the efficiency and accuracy of patient visits, resulting to quicker and more customised care. This streamlined approach improves patient experiences and enables hospitals, clinics, and physicians to manage a higher volume of patients each day (Meyer et al., 2020).

2.3.5 AI in Robotic Surgery

Hospitals are increasingly utilising AI and robotic systems for a range of procedures, from minimally invasive operations to complex open-heart surgeries. Surgeons operate robotic arms from a computer console, which provides a detailed, three-dimensional view of the surgical site. This advanced perspective enables the surgeon to guide the surgical team through the operation with precision. Robot-assisted surgeries have been shown to reduce complications, minimise pain, and accelerate recovery times (Bhandari et al., 2020).

2.4 Empirical Review

2.4.1 Integration of Privacy-Aware AI with HE

Rao et al. (2023) explored the relationship between privacy-aware AI with HE using Machine Learning (ML). The study emphasised that as the use of ML applications grows, the need for more data to make accurate predictions also increases and that is where big data comes in. However, the study highlighted that the limitations of big data, including privacy concerns, can be mitigated through the use of homomorphic encryption (HE) in ML processes. Through their research, Rao et al. (2023) proposed a multi-party privacy-protected ML approach that enables users to engage in AI tasks without having to reveal their private data. They utilised the Artificial Neural Network (ANN) ML technique to demonstrate that sensitive data can be kept secure while training common models. By leveraging homomorphic procedures within the main computing system, the researchers showed that gradient data can be securely transmitted and

combined among all parties. Moreover, the study emphasised the use of the Lion Algorithm (LA) to select the optimal key for the homomorphic encryption process. By modifying the learning model based on new gradient data, the researchers found that the privacy concerns associated with big data could be effectively addressed. However, it is paramount to note that this study did not specifically examine issues related to health data privacy.

2.4.2 HE Schemes for Data Protection in Cloud Computing

Ameur et al. (2023) explored the utilisation of homomorphic encryption in a multi-cloud setting to address security concerns. The researchers noted that outsourcing data to cloud platforms has become common as cloud computing technology offers high performance and data processing capabilities. However, data security and privacy remain significant challenges. To tackle these issues, they put forward a multi-cloud solution that enhances data availability and privacy by integrating private, public, and managed clouds with a unified user interface. The distribution of data across different data centres in this environment is based on data sensitivity and cloud reliability. Although the effectiveness of current encryption algorithms is established, they are resource-intensive and time-consuming. Moreover, these encryption methods necessitate the decryption of data before processing, making them inefficient. In contrast, homomorphic encryption allows data processing and manipulation while encrypted, enabling users or third parties, such as cloud providers, to perform functions on encrypted data without revealing its values.

Oh et al. (2022) evaluated the methods of maintaining data confidentiality through an Enhanced Homomorphic Encryption Scheme in mobile cloud computing. The researchers identified a need for enhanced privacy measures due to the efficiency and speed of data collection in mobile cloud environments, which can lead to vulnerabilities in data confidentiality and user privacy. To address this issue, they developed a privacy-preserving data collection technique that enables users to respond to personal information requests securely and confidentially. By utilising a Homomorphic Cryptosystem, the researchers were able to process data without decrypting it, ensuring the confidentiality of the data. Their evaluation of the scheme demonstrated a significant increase in efficiency compared to existing methods, with up to 170% improvement

in delay efficiency. This research highlights the importance of implementing robust data privacy measures in mobile cloud computing environments.

Omollo et al. (2017) carried out a research on improving the cloud computing's security of data by utilising an addition-composition fully HE scheme. The research focused on the application of HE schemes to address data security concerns in cloud computing. The study began by recognising Gentry (2009)'s work on partial HE schemes, specifically his development of a somewhat HE scheme based on ideal lattices incorporating both additive and multiplicative homomorphisms. The researchers also acknowledged advancements made by Dijkwherehe, who utilised integers instead of ideal lattices in the construction of the encryption scheme. Additionally, the study highlighted contributions by Braskerski in addressing encryption challenges through Learning with Errors (LWE) and Ring Learning with Errors (Ring LWE) problems. Despite these advancements, existing schemes still impose significant strain on computing resources. The researchers identified the need for an encryption scheme that could alleviate computational burdens on computing assets. They proposed a FHE scheme that leveraged both addition and composition operations to address this issue. The theoretical framework was translated into a JAVA algorithm and tested on hardware with minimal specifications. Results demonstrated that the new scheme facilitated faster encryption processes supported larger ciphertext sizes, and offered greater versatility. The addition-composition fully homomorphic encryption scheme was found to enhance data security in cloud computing, thereby fostering greater consumer confidence in cloud services and applications.

In a comprehensive study conducted by Lv (2021), data privacy protection was examined through the lens of homomorphic encryption. The research addressed the growing concern that data security and privacy are becoming major obstacles in the advancement of cloud computing. To enhance current homomorphic encryption technology, an ECC homomorphic encryption algorithm was introduced and implemented. This algorithm was utilised to establish a data privacy protection aggregation model, and its effectiveness in securing data operations was assessed. The results indicated that the ECC homomorphic encryption algorithm outperformed the traditional RSA encryption algorithm in terms of security. Specifically, ECC showed superior security with key lengths exceeding 300 bits compared to RSA with key lengths exceeding 2000

bits. Additionally, the ECC homomorphic encryption algorithm demonstrated faster encryption times and smaller cypher sizes compared to RSA. Overall, ECC exhibited lower computational complexity and better security performance in comparison to RSA. The study however did not focus on enhancing healthcare data.

Agarwal and Shrivastava (2021) explored methods for safeguarding data privacy in the cloud through the use of homomorphic encryption. The research was motivated by the significance of cloud computing in the field of information technology, emphasising its role in data storage and cost reduction for businesses. Despite these benefits, data security remains a major concern in cloud computing. Various techniques have been proposed to enhance secure data storage and retrieval, but many of them have drawbacks that diminish the practicality of cloud computing. The study focused on Partial HE, which enables operations to be performed on the data that is encrypted without the encryption being compromised. Unlike other security techniques that only support single-round encryption, this study aimed to secure medical data by implementing two rounds of encryption. This approach enables users to generate their random key and encrypt sensitive medical records each time they are shared with a third party, such as a foreign doctor or between a patient and a doctor. The research findings suggest that medical data can be shared securely over the internet in any file format, ensuring confidentiality and integrity. The encrypted data format prevents unauthorised access and modifications, making it difficult to search for or alter medical information without decryption. The study also highlights the difficulty of modifying medical data using this approach, as any changes can be easily detected.

2.4.3 Medical Data Sharing Using Multiparty Homomorphic Encryption

Scheibner et al. (2021) explored innovative ways to improve medical data sharing through advanced privacy-enhancing technologies. The researchers were inspired by the importance of multisite medical data sharing in modern clinical practice and research while recognising the need to protect individual privacy and maintain data utility. The study introduced the concept of multiparty homomorphic encryption, which combines homomorphic encryption and secure multiparty computation to ensure privacy and efficiency in data sharing. This approach not only meets legal requirements, such as those outlined in the European Union's General Data Protection Regulation but also streamlines the process by reducing the need for customised contractual

measures between institutions. By implementing multiparty homomorphic encryption, the researchers believe that medical research can be expedited and institutions can be encouraged to adopt common data interoperability standards.

2.4.4 Homomorphic Encryption and Machine Learning

Matias et al. (2023) explored the effects of homomorphic encryption on the efficiency of Machine Learning algorithms. The motivation for their research stemmed from the growing necessity for data sharing due to advancements in internet technology and the widespread availability of high-speed connections, particularly within business collaborations. With the increased importance of secure data transmission, processing, and storage, new encryption methods like homomorphic encryption have been introduced to facilitate confidential data sharing and analysis. The study investigated the utilisation of machine learning (ML) algorithms on data encrypted through homomorphic encryption, comparing four ML algorithms and evaluating the impact of encryption on performance metrics such as accuracy, precision, recall, and F1-Score, as well as processing time using a health dataset from the Kaggle platform. The findings of the study indicate that employing ML on homomorphically encrypted data can be done without significant performance decline, although potentially longer processing times associated with ML-based approaches working on encrypted data should be taken into consideration.

2.4.5 Data Privacy in Neural Network Training Using HE

Amorim et al. (2023) conducted a study on data privacy in neural network training and inference using homomorphic encryption. Their research delved into the various strategies and techniques employed to enhance data security and privacy in this context. The study examined the current state-of-the-art in homomorphic encryption for neural networks, identifying challenges and limitations that must be overcome to ensure its reliability and efficiency in preserving privacy. The researchers empirically explored different categories of homomorphic encryption schemes and their suitability for neural networks, as well as techniques for optimising the accuracy and efficiency of encrypted models. While the study demonstrated the potential of homomorphic encryption to provide strong data privacy guarantees for neural networks, it also highlighted obstacles such as limited support for advanced neural network operations, scalability issues, and performance trade-offs that need to be addressed.

2.4.6 Privacy Enhancement through HE and Statistical Transformation

Kumar et al. (2023) explored methods to improve user privacy by utilising privacy-chain-based homomorphic encryption schemes combined with statistical techniques. The research introduced a statistical transformation alongside a homomorphic encryption algorithm designed to modify both categorical and numerical data from the lung cancer, bank marketing, and adult income datasets, all while maintaining data utility. In the STHE approach, quasi-identifiers are first altered using an initialisation vector (IV), followed by the application of privacy-chain-based homomorphic encryption with RSA to secure data privacy. The effectiveness of the STHE algorithm was evaluated against leading algorithms using classifier models such as decision trees, random forests, extreme gradient boosting, and support vector machines. The results indicate that the proposed STHE algorithm surpasses existing methods in terms of performance, accuracy, data retrieval and transformation, and the balance between privacy preservation and data utility.

Zhang et al. (2023) reviewed sparse attention and homomorphic encryption as methods for improving privacy in large language models. They introduced a new framework for securing dialogue-based large models by integrating fully homomorphic encryption (FHE) and mechanisms to address privacy challenges in the information age. The framework aims to improve the accuracy and responsiveness of dialogue systems while safeguarding user privacy. One key aspect of the framework is its ability to process user data without decryption, thus preventing exposure of sensitive information even in case of data leakage whether during transmission or processing. Additionally, the framework reduces unnecessary computations by including a module that optimises processing efficiency by calculating attention scores only between key elements that significantly impact the output. Comparative analysis of various models revealed that the proposed method provides security similar to FHE while maintaining a minimal latency of 0.15 ms, highlighting its efficient encryption technology.

2.4.7 Challenges in Implementing HE

Gouert and Tsoutsos (2024) discussed the challenges of implementing homomorphic encryption in applications. They highlighted the difficulties in configuring homomorphic schemes and selecting appropriate parameters for different applications, even for experienced users. The

researchers proposed two approaches to address these issues: Walrus and HELM. Walrus is designed for arithmetic-intensive applications with shallow depths and high throughput requirements. It offers a user-friendly programming interface and automates parameter selection by analysing the application and considering factors such as homomorphic noise growth. The study demonstrated the effectiveness of Walrus through a neural network inference example and provided guidelines for its efficient usage. HELM, on the other hand, focuses on converting existing HDL designs to the encrypted domain for secure outsourcing on cloud servers. Unlike Walrus, HELM supports fully homomorphic encryption backends and is suitable for complex applications. This approach utilises both CPU and GPU acceleration to exploit the parallelism inherent in Boolean circuits when performing logic gate operations on encrypted individual bits.

2.5 Research Gaps

Despite the considerable progress in AI and its applications in healthcare, the privacy of these solutions remains a significant concern. The integration of privacy-aware AI with homomorphic encryption (HE) offers a promising approach to addressing this issue. This is demonstrated by Rao et al. (2023) who posit that HE can effectively protect sensitive data during machine learning (ML) processes by enabling secure data transmission and model training without exposing private information. However, their study did not specifically tackle health data privacy, which is critical given the unique sensitivity and regulatory requirements of medical information. To advance privacy protection in AI healthcare solutions, applying fully homomorphic encryption could offer enhanced security by enabling the processing of encrypted data while protecting the confidentiality of sensitive health records.

Additionally, the need for better privacy solutions in healthcare data sharing has been emphasised by recent studies such as Scheibner et al. (2021) who introduced multiparty homomorphic encryption to enhance medical data sharing while ensuring privacy and compliance with regulations. This approach, coupled with advancements in statistical transformation and HE by Kumar et al. (2023), suggests a viable path for improving privacy in healthcare data management. However, integrating these privacy-enhancing technologies with existing healthcare systems and ensuring their practical effectiveness remains challenging. FHE offers a comprehensive solution by allowing secure data processing and analysis without revealing sensitive information. To

fully realise the potential of FHE in healthcare, further research should focus on optimising its implementation for healthcare-specific scenarios, ensuring both security and usability.

2.6 Conceptual Framework

The solution consists of an architecture designed to ensure that sensitive patient data is processed by the machine learning algorithms while encrypted.

Once the user uploads the unencrypted data, an encryption key is generated and used to encrypt the uploaded data. This data is then sent to the machine learning algorithm which performs the required inference operations on the encrypted data using a secure method and then returns the encrypted inference results. The model never has access to the unencrypted data or the unencrypted inference results, ensuring that sensitive information is never exposed during processing. After this process, the encrypted inference results are returned and decrypted using the same encryption key. This process ensures that data confidentiality is maintained during inference.

Figure 2.2 show the conceptual digram.

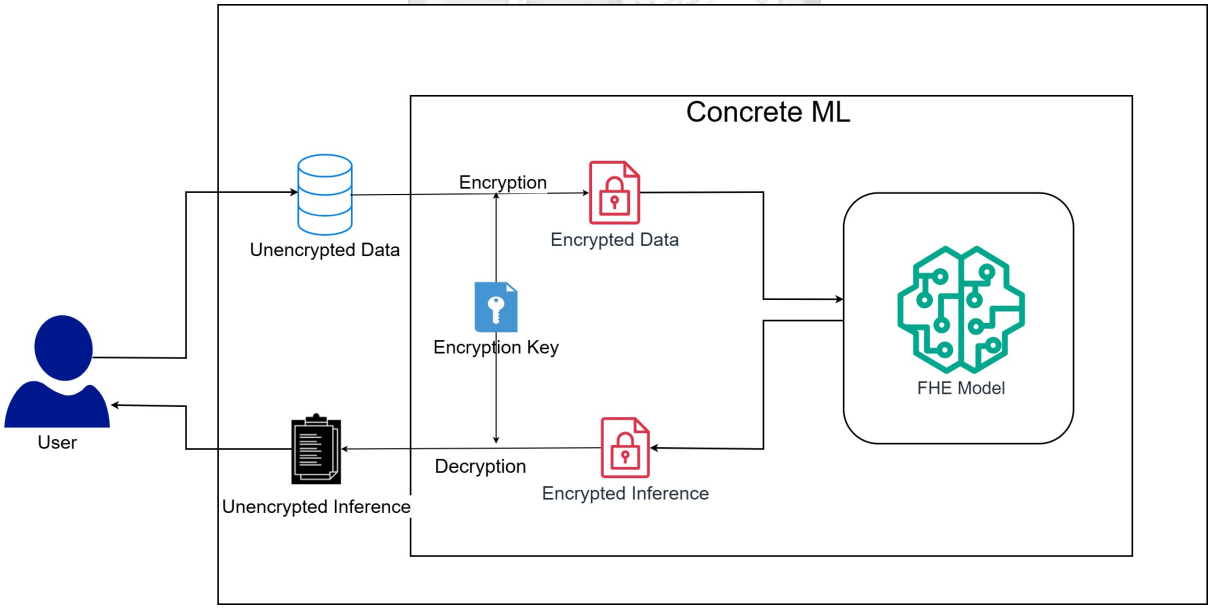

Figure 2.2: Conceptual Diagram.

Chapter 3
Methodology

3.1 Introduction

This chapter details the phases the application underwent during development, outlining the objectives and outcomes for each phase. The methodology was selected after thoroughly reviewing similar AI-based and cryptographic systems. The first objective focused on identifying challenges in managing and processing patient data for AI applications in healthcare. This was partially addressed in the literature review and was further explored during the Requirements Planning phase of the Rapid Application Development (RAD) methodology. The literature review also provided insights into applying homomorphic encryption in AI, presenting its operational challenges. The third objective aimed at developing an AI application integrating fully homomorphic encryption to protect health information privacy, was addressed during the User Design and Construction phases. Finally, the fourth objective, assessing the application's effectiveness in maintaining patient data privacy, was handled during RAD's Cutover and system validation stages.

3.2 Software Methodology

Rapid Application Development (RAD) was used to develop the AI application for melanoma detection with fully homomorphic encryption and Concrete ML (Zama.ai, 2024). RAD is a methodology emphasising a concise development cycle, using modern tools to deliver high-quality systems quickly (Sommerville, 2011). It has four phases. These are Requirements Planning, User Design, Rapid Construction, and Cutover. RAD's iterative approach helps manage costs and risks effectively while ensuring continuous refinement based on feedback. This methodology was chosen for its ability to deliver faster development and superior quality results compared to traditional life cycles, making it ideal for the melanoma detection AI application (Sommerville, 2011).

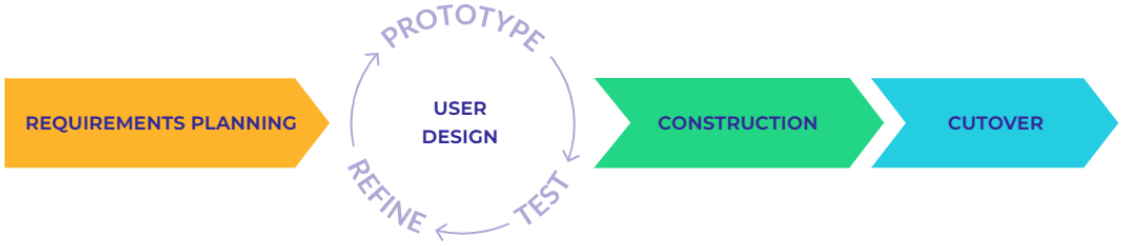

Figure 3.1: Rapid Application Development (RAD) Methodology

Source: jmix.io (2024)

3.3 Requirement Planning

The Requirements Planning phase was crucial for defining the scope and requirements of the melanoma detection AI application. This phase focused on determining the project's boundaries, such as the types of skin lesions to be detected, and the intended usage environments (Dorj et al., 2018; Shen et al., 2022). Data requirements were assessed by reviewing existing datasets like HAM10000 and ISIC 2017 to determine if additional data collection was necessary and the conclusion was that no additional data was needed (Codella et al., 2018). Technical considerations, including the selection of deep learning architectures and privacy-preserving technologies, led to an evaluation and selection of the ResNet50 model and the Concrete ML framework (tensorflow.org, 2024; Xiao et al., 2023; Zama.ai, 2024). Key performance measures such as privacy-preserving capabilities, usability, and accuracy were defined to guide the development (Brinker et al., 2018). A comprehensive schedule and project timeline, which can be found in Appendix D of this document, was established to ensure efficient progress and timely completion.

3.4 User Design

This phase aimed at creating a user-friendly interface for the AI application, ensuring alignment with user needs and workflows (Shen et al., 2022). The interface layout was designed to be clear and logical, using Streamlit to build a responsive and attractive user interface (Brinker et al., 2018; streamlit.io, 2024). Streamlit's pre-built widgets allow users to interact with the

AI application easily. Best practices for UI design were followed, optimising the layout for both desktop and mobile devices (streamlit.io, 2024). Contextual help features, like tooltips and modal windows, were used to provide users with relevant information. Iterative prototyping and user testing was conducted to refine the design based on user feedback (Gottesdiener, 1995). This phase ensured a user-friendly application was developed that meets users' needs, enhancing adoption and satisfaction.

3.5 Construction

The Construction phase implemented the application based on the prior requirements and design specifications. Infrastructure and tools were set up, including TensorFlow for implementing the ResNet50 model and Concrete Ml for FHE compilation (Brinker et al., 2018; Zama.ai, 2024). The ResNet50 model was selected because of its exceptional performance dermatological image analysis features and the availability of clinically pre-trained weights specifically calibrated for melanoma detection (tensorflow.org, 2024). These weights were obtained from an established medical imaging repository, enabling the model to distinguish between benign and melanoma skin lesions.

A comprehensive knowledge distillation methodology was implemented to transfer diagnostic expertise from the ResNet50 teacher model to a minimal custom model designed for FHE. This was necessary to address the inherent computational constraints of homomorphic encryption while preserving essential diagnostic capabilities. The minimal model utilised reduced filter quantities and a streamlined architecture to enhance FHE compatibility. (Zama.ai, 2024).

Best practices for development and code management were followed, using Git for version control (Daud et al., 2010; git scm.com, 2024). Iterative development and testing was critical to ensure that the final application met quality and performance standards (Martin, 1991). The result was an efficient application that accurately detects melanoma while protecting patient privacy, with a user-friendly interface built using Streamlit (streamlit.io, 2024).

3.6 Cutover

The Cutover phase transitioned the application from testing to a usable application, focussing on functionality, performance, and usability (Sommerville, 2011). Functional testing verified that the application met design requirements, using test cases developed during the Requirements Planning phase. Usability testing evaluated the application interface and ease of use and navigation was measured Brinker et al. (2018). Performance testing assessed the application's responsiveness and scalability under varying loads. Identified bottlenecks that could be addressed were optimised to ensure that the application worked effectively.

3.7 System Validation

A rigorous validation process confirmed that the AI application effectively addresses the intended problem (Sommerville, 2011). Privacy protection measures implemented using FHE and Concrete ML were validated and proven to maintain confidentiality by ensuring patient data is encrypted during inference (Xiao et al., 2023). Per Brinker et al. (2018), the performance of the application in the detection of melanoma was also measured.

3.8 Deployment

The application was then deployed to Hugging Face (huggingface.co, 2025). Hugging Face Spaces offers an ideal environment to host AI and ML applications as it provides a free hosting service, and resources can be upgraded to scale with the needs of the application. This is particularly important for this application, since FHE applications are resource-intensive.

Chapter 4
System Design and Architecture

4.1 Introduction

This chapter delves into the architecture and design of the FHE melanoma detection application which aligns with the requirements previously identified during the User Design and Requirement Planning phases. The application was implemented following RAD methodology.

4.2 System Design

The design and architecture chosen was based on findings from the Requirement planning phase and Literature Review stage which highlighted the need for a simplified user interface in line with other AI systems developed. This allows users to have an optimum experience while getting the full functionality of the application. This led to the identification of two modules, the Educational Content Module and the Diagnosis Module.

4.2.1 Educational Content Module

This module enables users who access the application to get valuable information about how the application works and about melanoma prevention and detection. It explains how FHE enhances privacy by enabling inference on encrypted data. The module highlights how this technology can be utilised to facilitate secure data analysis enabling healthcare providers to process patient data without compromising confidentiality and therefore maintain compliance with data regulations such as HIPAA. The workflow of the application is explained from the first step of uploading an image, encryption, and inference, to the final result of getting the decrypted inference results.

It also highlights the importance of the early detection of melanoma, stressing that it is treatable if detected early. The content encourages regular self-examination and annual dermatologist visits. The module also shows a video about diagnosing melanoma.

4.2.2 Diagnosis Module

This module carries the core functionality of the application. It enables users to upload images and get melanoma inference results. Users select the image they want to run the inference on, the image is then checked to ensure that it is a valid skin lesion. If it is not, an error message is

displayed. Once a valid skin lesion is uploaded, it is first preprocessed to ensure it meets the requirements of the FHE model. It then gets into the FHE pipeline where it is first encrypted, then inference is run on the encrypted image. The encrypted inference produced is decrypted and displayed to the user with recommendations about viable steps.

4.3 System Design Tools

This section discusses various system design tools which showcase the system architecture. The tools used are a use case diagram, a sequence diagram and a collaboration diagram.

4.3.1 Use Case Diagram

The use case diagram, Figure 4.1, shows the users' interaction with the application and the various use cases that they are involved in.

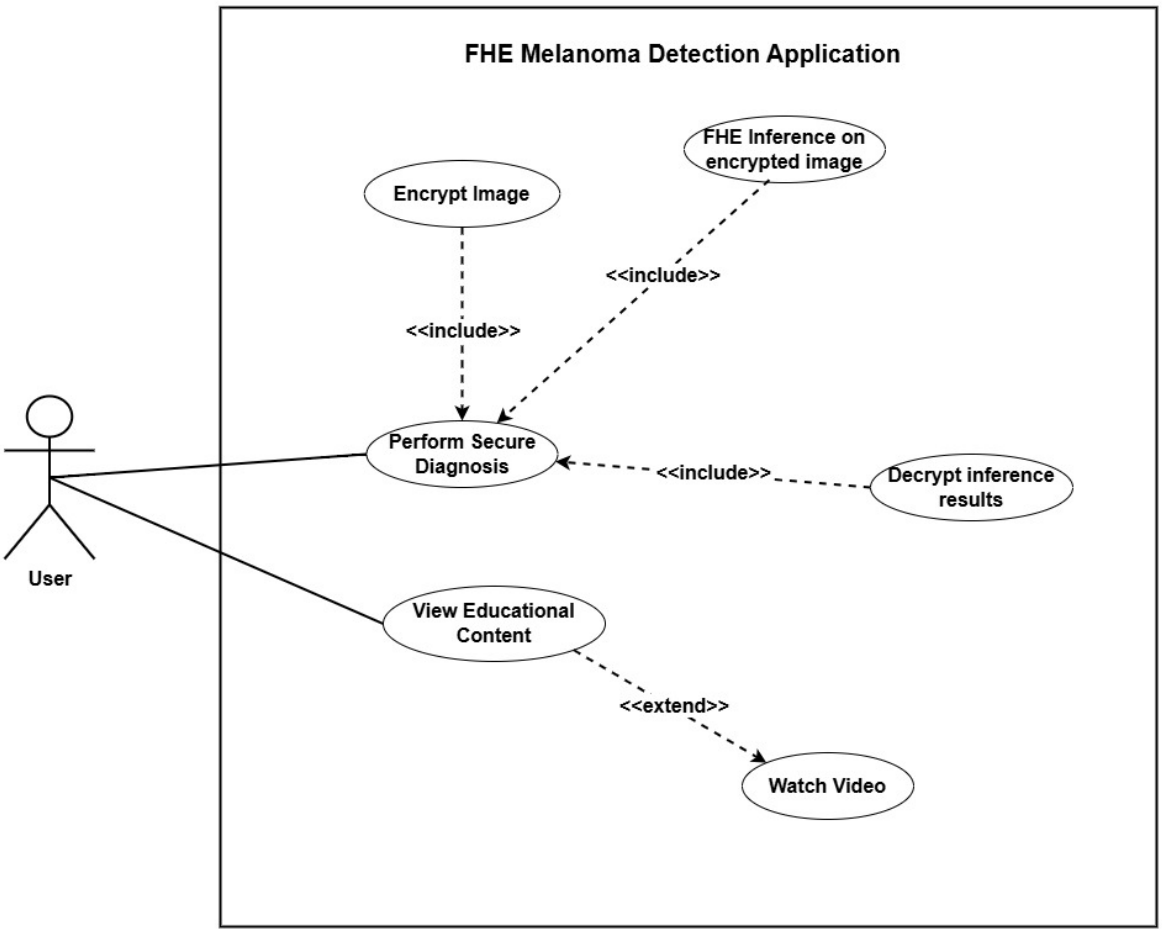

Figure 4.1: Use Case Diagram

4.3.2 View Educational Content Use Case

This use case outlines the requisite steps that users should follow to access the educational content. Table 4.1 shows these steps.

Table 4.1: View Educational Content Use Case

Use Case Name:	View Educational Content
Description:	This use case allows users to access relevant information about the application, the technology it is using and how it works. They also access information about melanoma prevention, mitigation and diagnosis.
Primary Actor:	User
Secondary Actor:	None
Preconditions:	None
Postconditions:	1. The system displays the educational content.
Main Flow:	1. The user checks the Educational Content radio button on the navigation bar. 2. The system loads the educational content page. 3. The use case ends.

4.3.3 Watch Video Use Case

This use case outlines the requisite steps that users should follow to watch the video about melanoma diagnosis. Table 4.2 shows these steps.

Table 4.2: Watch Video Use Case

Use Case Name:	Watch Video

Description:	This use case allows users to watch an informative video about melanoma diagnosis
Primary Actor:	User
Secondary Actor:	None
Preconditions:	None
Postconditions:	1. The system plays the video.
Main Flow:	1. The user checks the Educational Content radio button on the navigation bar. 2. The system loads the educational content page. 3. The user scrolls down to the video player. 4. The user clicks on the play button and the video loads. 5. The use case ends.

4.3.4 Perform Secure Diagnosis Use Case

This use case outlines the requisite steps that users should follow to upload an image and perform a melanoma diagnosis. Table 4.3 shows these steps.

Table 4.3: Perform Secure Diagnosis Use Case

Use Case Name:	Perform Secure Diagnosis
Description:	This use case allows users to upload an image of a skin lesion and perform a secure melanoma diagnosis taking advantage of FHE to preserve their privacy.
Primary Actor:	User

Secondary Actor:	None
Preconditions:	None
Postconditions:	1. The system displays the inference results and gives some recommendations.
Main Flow:	1. The system loads the diagnosis page by default. If the user is on another page, the user checks the Diagnosis radio button on the navigation bar. 2. The user clicks on the Browse files button. 3. The user selects the image they want to upload. 4. The file is uploaded and the system encrypts the image. 5. The system performs melanoma inference on the image and returns the encrypted inference. 6. The system decrypts the inference. 7. The system displays the results and gives some recommendations. 8. The use case ends.
Alternative Flows:	**3a. User uploads an invalid skin lesion** 1. The system displays an error message notifying the user that the image is not a valid skin lesion. 2. Use case resumes at main flow step 1.

4.3.5 Sequence Diagram

Figure 4.2 is a sequence diagram depicting the interaction between a user and the system when running inference on an image.

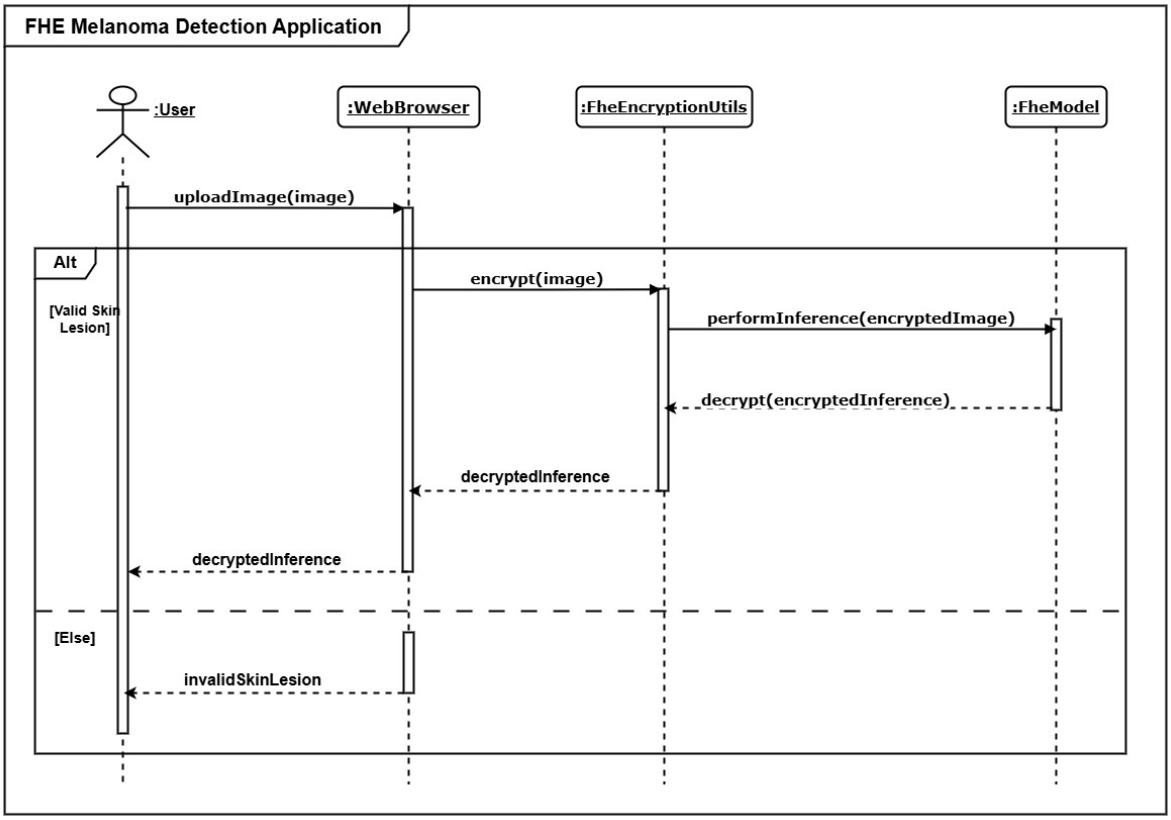

Figure 4.2: Sequence Diagram

4.3.6 Collaboration Diagram

Figure 4.3 is a collaboration diagram depicting the interaction between a user and the system when running inference on an image.

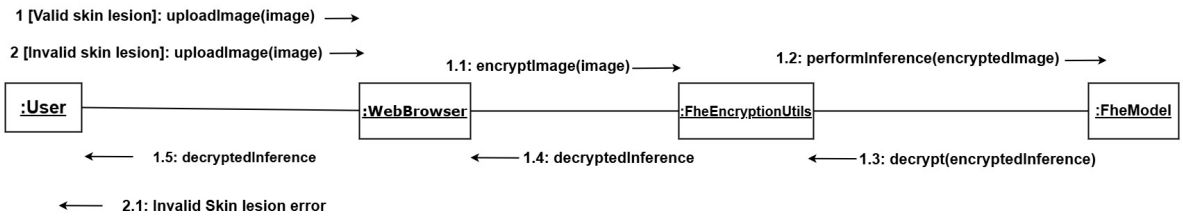

Figure 4.3: Collaboration Diagram

4.3.7 Educational Content Page Wireframe

Figure 4.4 shows a wireframe diagram of the educational content page. The page is divided into sections with the navigation on the left and the sections that shows the information to the right. At the bottom, there is a section that shows the educational video.

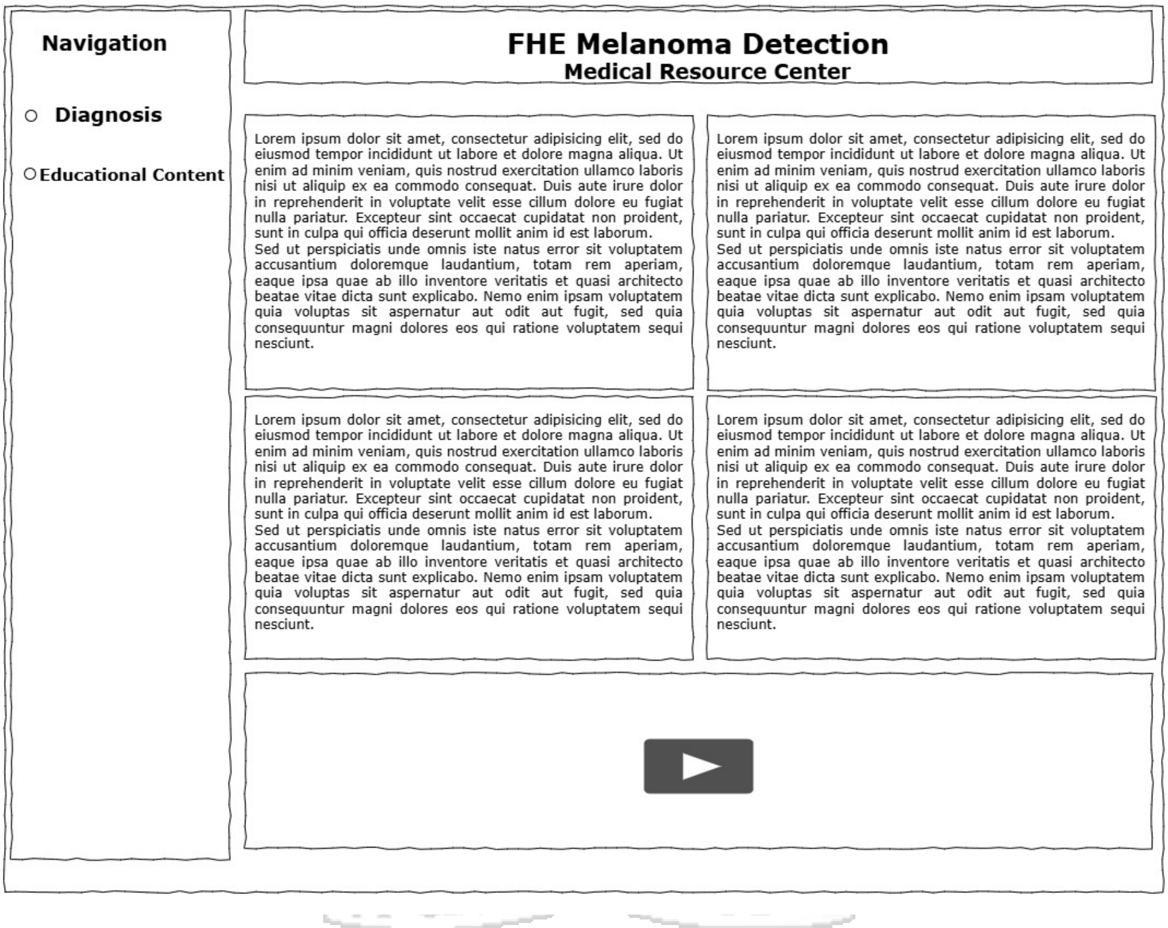

Figure 4.4: Educational Content Page Wireframe

4.3.8 Diagnosis Page Wireframe

Figure 4.5 is the wireframe for the diagnosis page. It has a very simple layout with the navigation to the left and the upload section to the right. The upload section has some text that explains the image upload process.

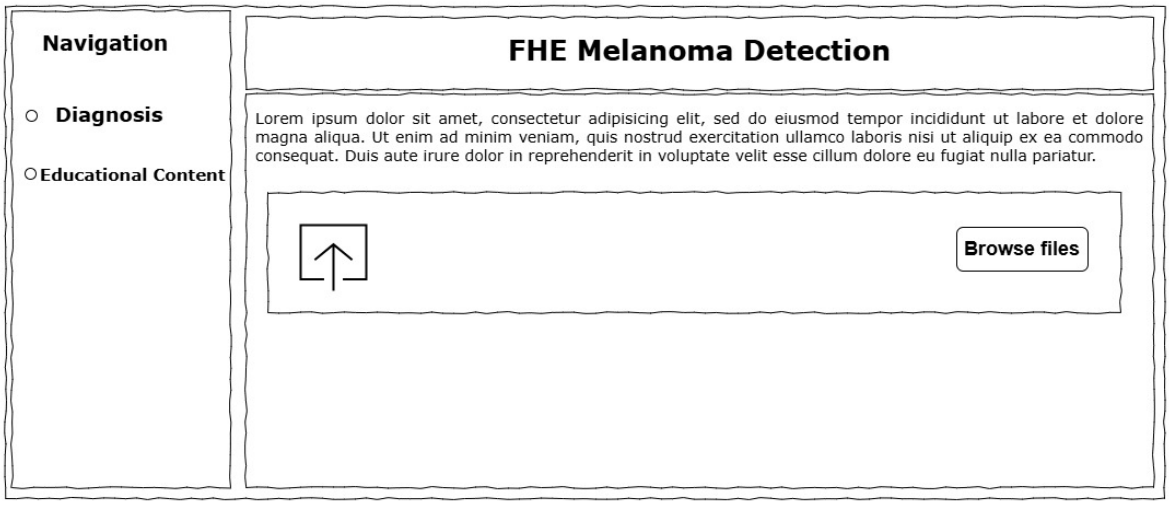

Figure 4.5: Diagnosis Page Wireframe

Chapter 5
System Implementation and Testing

5.1 Introduction

This chapter covers the implementation of the FHE Melanoma Detection Application. It highlights the significant features and includes screenshots of the important user interfaces. Various tests carried out on the system are also discussed.

5.2 Implementation Environment

This application was developed on the Streamlit framework. It is written in Python and requires various Python packages which are necessary for its operation. These packages include TensorFlow for implementing the ResNet50 model and handling neural network operations, Concrete ML for FHE compilation and execution, and, NumPy for numerical computations and various image processing libraries such as OpenCv (streamlit.io, 2024; tensorflow.org, 2024; Zama.ai, 2024).

Locally, the application was developed on Ubuntu WSL, (Windows Subsystem for Linux) and then deployed on Hugging Face's servers. These two environments have the dependencies mentioned above installed.

5.3 System Functionality

The following screenshots show important user interfaces which enable the application's functionality.

5.3.1 Educational Content

Users have access to educational content through the application. This content details how the application works and its benefits. They also have access to information about melanoma, what it is and the benefits of early detection. Users navigate to this page by checking the "Educational Content" radio button on the navigation bar.

Figure 5.1 shows the educational content page.

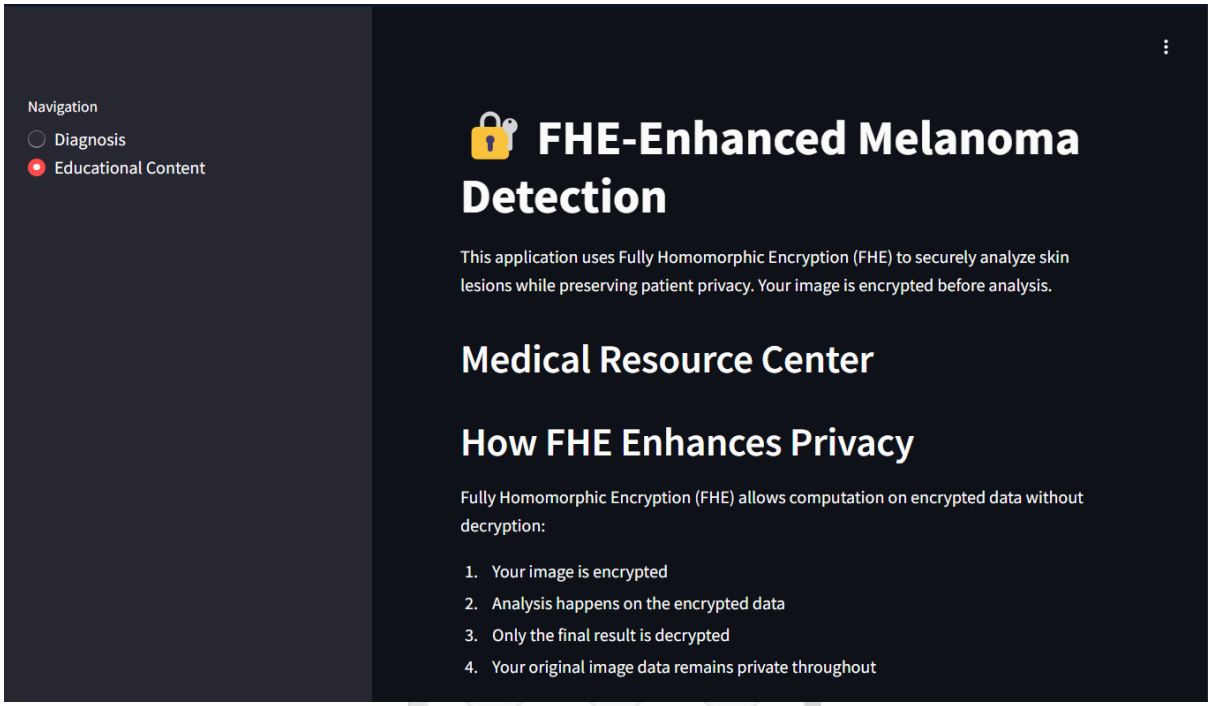

Figure 5.1: Educational Content Page

Figure 5.2 shows the educational content page with the video.

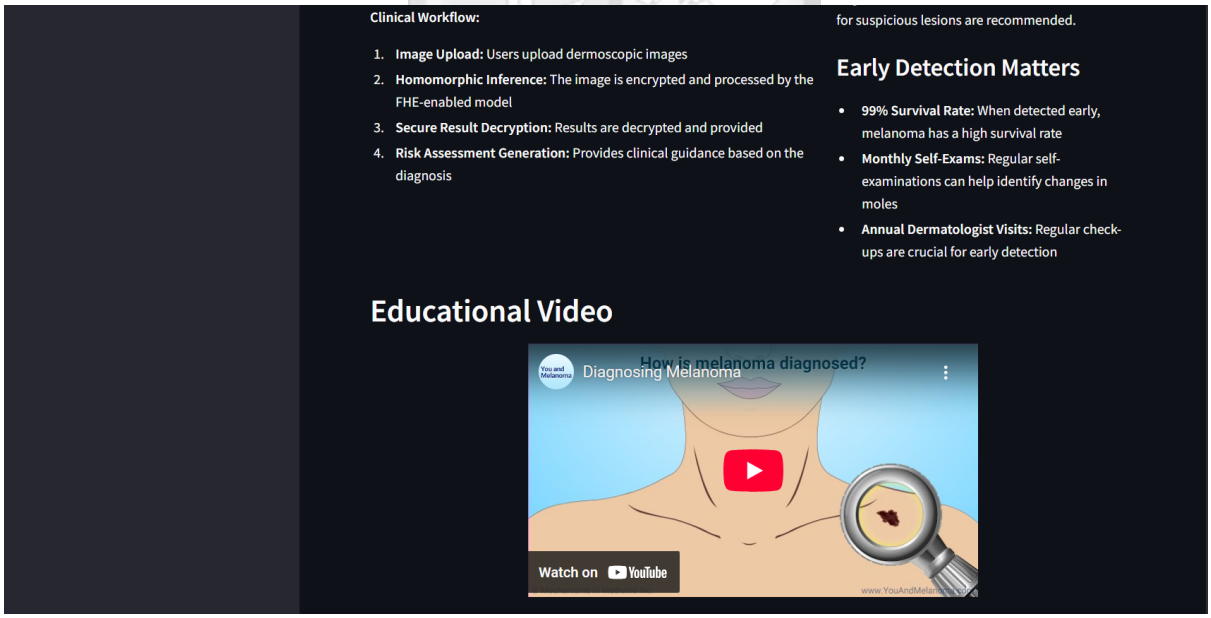

Figure 5.2: Educational Video

5.3.2 Diagnosis

This is the core functionality of the application. It enables users to upload a skin lesion image and perform melanoma diagnosis while preserving their confidentiality. This is the default page but users can navigate to this page by checking the "Diagnosis" radio button on the navigation bar.

Figure 5.3 shows the diagnosis page.

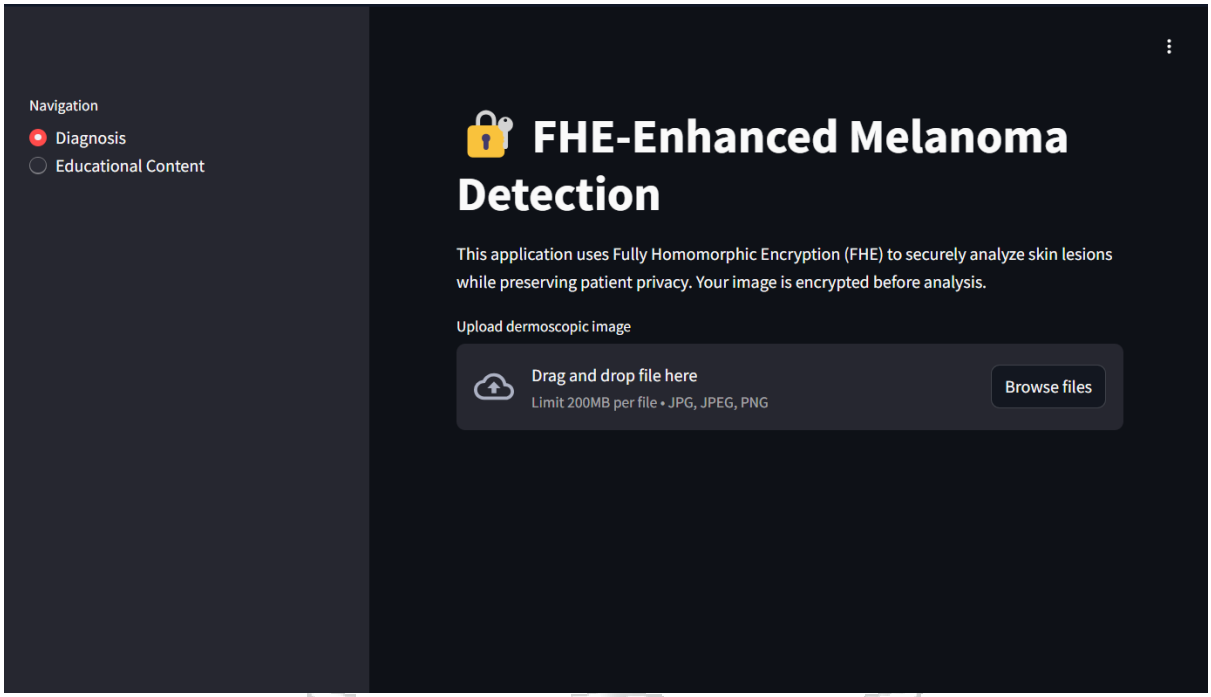

Figure 5.3: Diagnosis Page

Once a valid skin lesion is uploaded it is first preprocessed to ensure it meets the requirements of the FHE model. The preprocessing involves resizing the image, normalizing pixel values and enhancing contrast to optimize feature detection. Encryption keys are then generated specifically for this session and used to encrypt the image. Inference is then run on the encrypted image through the minimal FHE model. This computation occurs entirely on the encrypted data ensuring the actual image is never exposed. The encrypted inference produced is then decrypted and the results are displayed to the user with recommendations such as consulting a dermatologist if the melanoma probability exceeds the threshold.

5.4 Testing

Tests were conducted to verify whether the developed application functions as per the requirements. Compatibility and functional tests were carried out.

5.4.1 Compatibility Testing

The application was tested on different screen sizes to ensure it was responsive. Streamlit applications are easily customisable ensuring they automatically adjust their layouts and user interface for different screen sizes. This adaptability is achieved through Streamlit's built-in layout components such as columns and containers which dynamically reorganise content based on available screen size. This responsive behaviour enhances usability across different devices, ranging from desktop computers to tablets and mobile phones. Moreover, the navigation bar component's collapsible nature contributes to the application's adaptability as it is configured to appear on larger screens while remaining hidden on mobile devices until it is toggled. The application passed the compatibility tests.

Figure 5.4 shows the diagnosis page on a phone screen.

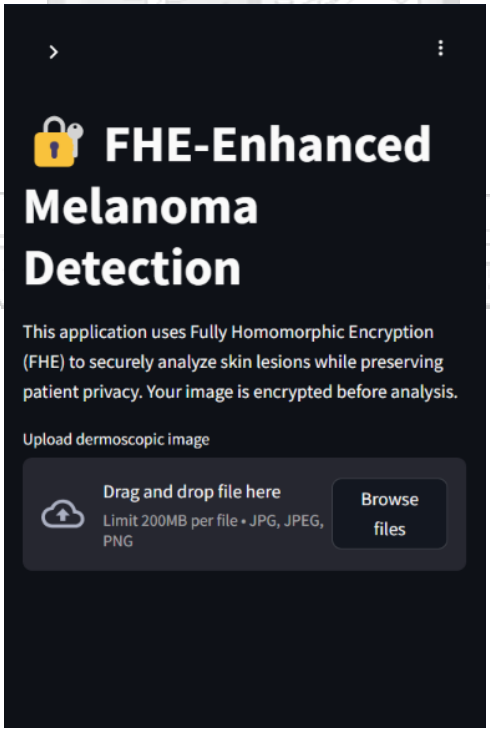

Figure 5.4: Compatibility Test

5.4.2 Functional Testing

This process involved testing the upload of various categories of invalid inputs to ensure accurate error handling. Non-image files including text documents and PDFs as well as non-dermatological content such as general objects were tested. These uploads failed and displayed an error message. This prevents unnecessary computational overheads and errors that might arise from processing invalid inputs.

Figure 5.5 shows the error message displayed when an invalid content is uploaded.

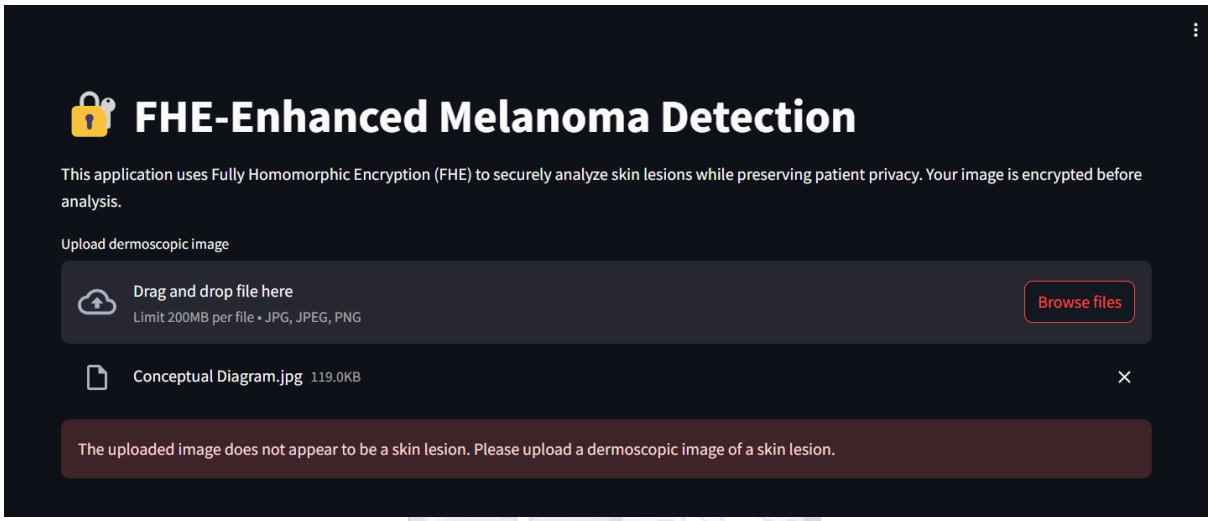

Figure 5.5: Functional Test

5.5 System Validation

The application underwent security validation conducted by a certified security expert with Certified Ethical Hacker (CEH) credentials. The expert attempted to access unencrypted patient data using various attack mechanisms such as traffic interception and memory inspection during the FHE operations. Despite the attempts, user data remained securely encrypted throughout the processing pipeline confirming the system's fundamental privacy guarantees. The expert also examined captured packets containing encrypted images and results and confirmed that all these data remained encrypted with only encrypted ciphertexts visible.

The application also underwent validation for clinical accuracy. A set of 50 dermatological images with confirmed diagnoses, 25 melanoma and 25 benign, was passed through the FHE-

compatible distilled model. The results showed that the FHE application maintained about 85% accuracy demonstrating minimal loss of diagnostic capability.

Figure 5.6 shows the output logs of a successful inference.

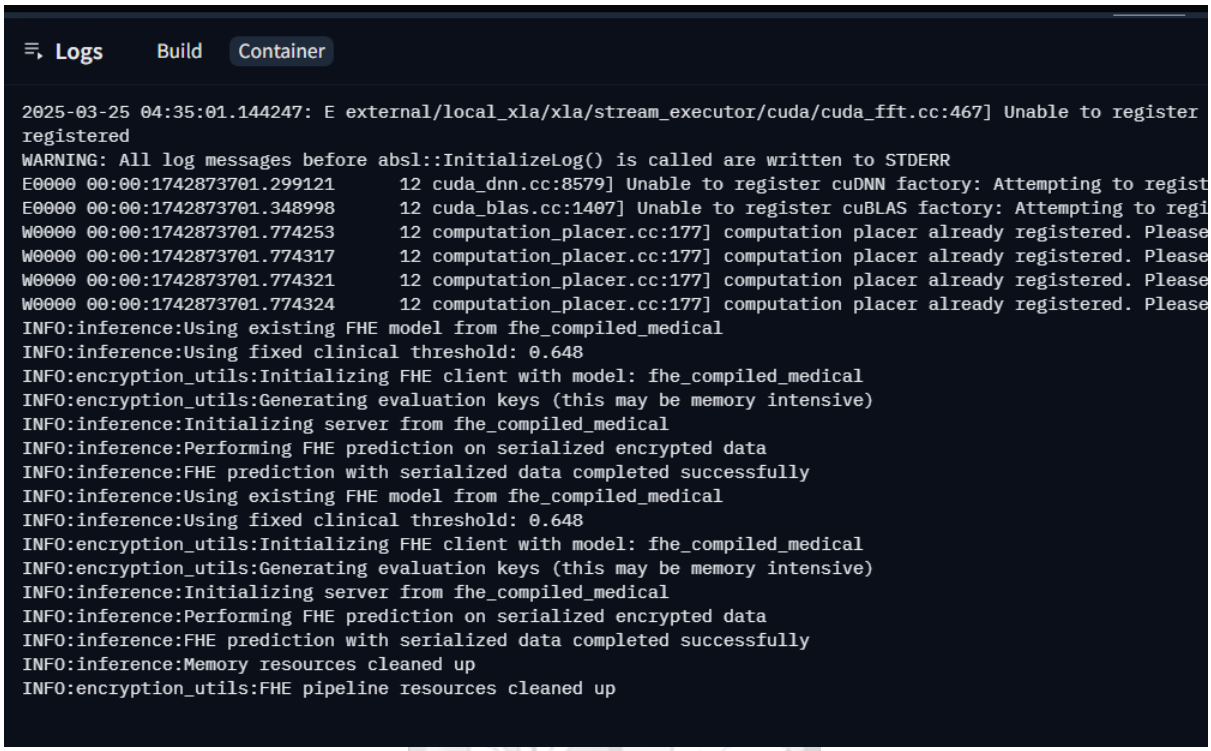

Figure 5.6: Sample Validation Output Logs

In summary, the validation confirmed that the FHE system successfully protected patient privacy throughout the diagnostic process while maintaining clinically relevant accuracy. The findings demonstrated that FHE provides a viable approach for privacy-preserving medical diagnostics, though with notable computational trade-offs that should be considered in deployment planning. These validation results provide strong evidence for the system's effectiveness in enhancing healthcare data privacy while maintaining diagnostic capability.

Chapter 6
Discussion of Results

6.1 Introduction

This chapter compares the results vis-a-vis the research objectives and the extent to which they concur with the literature review. The research hypothesis is also validated.

6.2 Objective One

The first objective, which was to analyse the challenges in handling and processing sensitive patient data required for AI applications in healthcare, was adequately addressed in the introduction and literature review. The introduction highlighted significant privacy concerns when working with healthcare data, emphasising the need to comply with regulations such as HIPAA, GDPR, and the Data Protection Act. As discussed in the literature, healthcare data contains highly sensitive personal and medical information that requires robust protection against unauthorised access and data breaches. The literature review explored how AI systems require vast amounts of data to function effectively, creating a unique hurdle between data accessibility and privacy protection. This challenge is particularly acute in healthcare, where AI offers a lot of benefits including improved diagnostic accuracy, accelerated drug discovery, enhanced patient management, and streamlined administrative processes. These challenges collectively create a complex environment for AI integration in healthcare, where privacy preservation must be carefully balanced with the potential for improvements in healthcare.

6.3 Objective Two

The second objective, exploring existing research on applying homomorphic encryption in artificial intelligence, was comprehensively covered in the literature review through an examination of various homomorphic encryption models and their applications. The review explored several theoretical frameworks including the Fully Homomorphic Encryption (FHE) Model by Gentry, the Goldwasser-Micali Scheme, the HEAWS Model, Reinforcement Learning (RL) Model, Robust Cramer Shoup Delay Optimised Fully Homomorphic (RCSDOFH), and the Ring Learning With Errors (RLWE). Each of these models offers unique approaches to enabling com-

putation on encrypted data without requiring decryption. The literature review highlighted that Concrete-ML provides an open-source library that enables the creation of Fully Homomorphic Machine Learning models. This technology allows for the conversion of standard ML models into their FHE equivalents, enabling secure processing of sensitive data. The review concluded that homomorphic encryption, particularly through implementations like Concrete-ML, offers a promising solution for maintaining privacy in AI applications by allowing data to be analysed without sacrificing confidentiality. This informed the decision to choose Concrete ML in the implementation of the FHE melanoma detection application..

6.4 Objective Three

The third objective was successfully achieved through the design, development and testing of the FHE Melanoma detection application. The implementation employed an architecture that ensured only the encrypted data is processed. The system's development process involved using a ResNet50 model as a teacher, leveraging its pre-trained weights specifically calibrated for melanoma detection, and then creating a minimal custom CNN architecture optimised for FHE compatibility through knowledge distillation. This approach successfully addressed the computational constraints of homomorphic encryption while preserving essential diagnostic capabilities.

The application was developed using Python on the Streamlit framework, which provided an intuitive and responsive user interface accessible across different devices. The responsive design ensured that users could access the system from various platforms, including mobile devices, enhancing its practical utility in clinical settings. Testing confirmed that the implemented system could effectively detect melanoma in dermatological images while keeping the data encrypted throughout the analysis process. Functional testing validated that the application properly handled input validation, correctly identifying and rejecting non-dermatological images, thereby ensuring that only valid skin lesion images entered the FHE processing pipeline.

6.5 Objective Four

The fourth objective was achieved through comprehensive validation of the application's effectiveness in maintaining user data confidentiality. A certified security expert conducted different

tests on the application attempting to intercept and access unencrypted data during the diagnostic process. Diverse testing approaches were employed to thoroughly assess the system's security characteristics. These tests confirmed that patient data remained securely encrypted throughout the processing pipeline, with all sensitive information protected during transmission and analysis.

The validation also addressed the computational correctness of the FHE operations, verifying that operations performed on encrypted data produced the same results as identical operations on unencrypted data. This confirmed the system's computational integrity while maintaining privacy. While the FHE approach successfully protected sensitive data, it introduced significant computational overhead compared to unencrypted processing. This finding represents an expected trade-off between privacy preservation and computational efficiency that must be considered in deployment planning.

6.6 Research Hypothesis

The research hypothesis that an AI melanoma diagnostic system utilising homomorphic encryption would maintain the confidentiality of sensitive patient health data while providing accurate detection has been validated through both security testing and clinical accuracy assessment. This has been covered in the discussion of results for objective three and four above. The system successfully implemented an architecture where sensitive patient information remained secure throughout the entire diagnostic workflow with all inference computations performed on encrypted data.

The validation demonstrated that the FHE-compatible model maintained clinically relevant diagnostic capabilities despite the architectural simplifications required for FHE compatibility. The application's ability to maintain a clinically relevant diagnostic threshold while operating entirely on encrypted data confirms that homomorphic encryption offers a promising solution to protect sensitive information while still enabling valuable medical insights to be extracted. These results validate the hypothesis, demonstrating that advanced medical diagnostics can be performed while maintaining patient data confidentiality through homomorphic encryption.

Chapter 7
Conclusions, Recommendations and Future Work

7.1 Conclusions

This dissertation focused on developing a privacy-preserving AI application for melanoma detection using Fully Homomorphic Encryption (FHE). The aim was to protect sensitive patient health data while enabling accurate AI-based diagnostics. The application was successfully developed and fulfilled this requirement. Through the implementation of a knowledge distillation approach, a complex ResNet50 model was transformed into a minimal CNN architecture compatible with FHE operations. The distillation process effectively transferred diagnostic expertise from the teacher model to the FHE-compatible student model. The security validation confirmed that user data remained protected throughout the analysis process, with all operations performed on encrypted data. The privacy-preserving application retained clinically relevant diagnostic capabilities despite architectural simplifications necessary for FHE implementation.

7.2 Recommendations

The application is best suited for healthcare environments where both diagnostic accuracy and patient privacy are critical concerns. This application should be deployed in clinical settings with proper hardware resources to accommodate the computational demands of homomorphic encryption. Healthcare organisations implementing this application should ensure servers have sufficient memory and multi-core processing capabilities to handle the resource-intensive FHE operations efficiently.

For optimal performance, the application should be deployed on dedicated hardware resources rather than shared infrastructure. The significant computational overhead of FHE operations requires specialised allocation of processing power to maintain reasonable response times in clinical settings. Healthcare providers should be trained on the privacy benefits of FHE to properly communicate these advantages to patients, especially for dermatological images which patients often consider highly personal.

Healthcare organisations, where the application is adopted, should establish clear protocols for handling cases where the FHE analysis results are inconclusive or require further clinical evaluation. This ensures appropriate follow-up care while preserving patient data privacy. Additionally, implementation should include robust monitoring systems to track performance metrics and identify areas for optimisation as technology advances.

7.3 Future Work

Future work may cover areas such as performance optimisation, integration with other systems and enhanced user experience.

For performance optimisation, the current implementation has significant computational overhead compared to unencrypted processing. This could be reduced through specialised hardware acceleration or further algorithmic improvements. The FHE operations could be optimised to better utilise parallel processing capabilities, potentially reducing inference time with appropriate hardware. Additionally, exploring alternative FHE schemes that offer better computational efficiency for convolutional operations could yield significant performance improvements.

Integration could be carried out to allow the application to work directly with hospital information systems and existing dermatology workflow software. This would improve adoption rates and practical utility in clinical settings. The user interface could be enhanced with additional features such as case management and automated follow-up scheduling based on risk assessment.

Finally, mobile applications could be developed for both iOS and Android platforms to facilitate remote consultations while maintaining the privacy-preserving nature of the application, which would be particularly valuable in telehealth scenarios where dermatologist access is limited but privacy concerns remain paramount.

References

Acar, A., Aksu, H., Uluagac, A. S., and Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (Csur)*, 51(4):1–35.

Agarwal, P. and Shrivastava, P. (2021). Enhancing data security in cloud computing through homomorphic encryption. *Computology: Journal of Applied Computer Science and Intelligent Technologies*, 1(1):32–39.

Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A. I., Almohareb, S. N., Aldairem, A., Alrashed, M., Bin Saleh, K., Badreldin, H. A., et al. (2023). Revolutionizing healthcare: the role of artificial intelligence in clinical practice. *BMC medical education*, 23(1):689.

Ameur, Y., Bouzefrane, S., et al. (2023). Handling security issues by using homomorphic encryption in multi-cloud environment. *Procedia Computer Science*, 220:390–397.

Amorim, I., Maia, E., Barbosa, P., and Praça, I. (2023). Data privacy with homomorphic encryption in neural networks training and inference. In *International Symposium on Distributed Computing and Artificial Intelligence*, pages 365–374. Springer.

Bajwa, J., Munir, U., Nori, A., and Williams, B. (2021). Artificial intelligence in healthcare: transforming the practice of medicine. *Future healthcare journal*, 8(2):e188.

Bertino, E. and Ferrari, E. (2017). Big data security and privacy. In *A comprehensive guide through the Italian database research over the last 25 years*, pages 425–439. Springer.

Bhandari, M., Zeffiro, T., and Reddiboina, M. (2020). Artificial intelligence and robotic surgery: current perspective and future directions. *Current opinion in urology*, 30(1):48–54.

Bharati, S. and Podder, P. (2022). Machine and deep learning for iot security and privacy: applications, challenges, and future directions. *Security and communication networks*, 2022:1–41.

Blanco-Gonzalez, A., Cabezon, A., Seco-Gonzalez, A., Conde-Torres, D., Antelo-Riveiro, P., Pineiro, A., and Garcia-Fandino, R. (2023). The role of ai in drug discovery: challenges, opportunities, and strategies. *Pharmaceuticals*, 16(6):891.

Bonte, C., Iliashenko, I., Park, J., Pereira, H. V., and Smart, N. P. (2022). Final: faster fhe instantiated with ntru and lwe. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 188–215. Springer.

Bootland, C., Castryck, W., and Vercauteren, F. (2020). On the security of the multivariate ring learning with errors problem. *Open Book Series*, 4(1):57–71.

Brinker, T. J., Hekler, A., Utikal, J. S., Grabe, N., Schadendorf, D., Klode, J., Berking, C., Steeb, T., Enk, A. H., and Von Kalle, C. (2018). Skin cancer classification using convolutional neural networks: systematic review. *Journal of medical Internet research*, 20(10):e11936.

Cao, Z. and Liu, L. (2015). On the weakness of fully homomorphic encryption. *arXiv preprint arXiv:1511.05341*.

Codella, N. C., Gutman, D., Celebi, M. E., Helba, B., Marchetti, M. A., Dusza, S. W., Kalloo, A., Liopyris, K., Mishra, N., Kittler, H., et al. (2018). Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, pages 168–172. IEEE.

Cohen, I. G. and Mello, M. M. (2018). Hipaa and protecting health information in the 21st century. *Jama*, 320(3):231–232.

Dash, S., Shakyawar, S. K., Sharma, M., and Kaushik, S. (2019). Big data in healthcare: management, analysis and future prospects. *Journal of big data*, 6(1):1–25.

Daud, N. M. N., Bakar, N. A. A. A., and Rusli, H. M. (2010). Implementing rapid application development (rad) methodology in developing practical training application system. In *2010 International Symposium on Information Technology*, volume 3, pages 1664–1667. IEEE.

Davenport, T. and Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future healthcare journal*, 6(2):94.

Dorj, U.-O., Lee, K.-K., Choi, J.-Y., and Lee, M. (2018). The skin cancer classification using deep convolutional neural network. *Multimedia Tools and Applications*, 77:9909–9924.

Gentry, C. (2009). *A fully homomorphic encryption scheme*. Stanford university.

git scm.com (2024). Git - Documentation — git-scm.com. https://git-scm.com/doc. [Accessed 09-04-2024].

Goldwasser, S. (1984). *Probabilistic encryption: Theory and applications (partial information, factoring, pseudo random bit generation)*. University of California, Berkeley.

Gottesdiener, E. (1995). Rad realities: Beyond the hype to how rad really works. *Application Development Trends*, 2(8):28–38.

Gouert, C. and Tsoutsos, N. G. (2024). Data privacy made easy: Enhancing applications with homomorphic encryption. *Cryptology ePrint Archive*.

huggingface.co (2025). Hugging Face – The AI community building the future. — huggingface.co. https://huggingface.co/. [Accessed 20-01-2025].

Hussain, A., Malik, A., Halim, M. U., and Ali, A. M. (2014). The use of robotics in surgery: a review. *International journal of clinical practice*, 68(11):1376–1382.

Iatropoulou, S., Anastasiou, T., Karagiorgou, S., Petrou, P., Alexandrou, D., and Bouras, T. (2023). Privacy-preserving data federation for trainable, queryable and actionable data. In *2023 IEEE 39th International Conference on Data Engineering Workshops (ICDEW)*, pages 44–48. IEEE.

jmix.io (2024). What is Rapid Application Development (RAD)? - Jmix blog — jmix.io. https://www.jmix.io/rapid-application-development/. [Accessed 09-04-2024].

Jordan, M. I. and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260.

Kaissis, G. A., Makowski, M. R., Rückert, D., and Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6):305–311.

Kaur, S., Singla, J., Nkenyereye, L., Jha, S., Prashar, D., Joshi, G. P., El-Sappagh, S., Islam, M. S., and Islam, S. R. (2020). Medical diagnostic systems using artificial intelligence (ai) algorithms: Principles and perspectives. *IEEE Access*, 8:228049–228069.

Kenyalaw (2019). Data protection act. http://kenyalaw.org:8181/exist/kenyalex/actview.xql?actid=CAP. Laws of Kenya.

Kumar, G. S., Premalatha, K., Maheshwari, G. U., and Kanna, P. R. (2023). No more privacy concern: A privacy-chain based homomorphic encryption scheme and statistical method for privacy preservation of user's private and sensitive data. *Expert Systems with Applications*, 234:121071.

Lee, S.-I., Celik, S., Logsdon, B. A., Lundberg, S. M., Martins, T. J., Oehler, V. G., Estey, E. H., Miller, C. P., Chien, S., Dai, J., et al. (2018). A machine learning approach to integrate big data for precision medicine in acute myeloid leukemia. *Nature communications*, 9(1):42.

Li, Q., Yue, Y., and Wang, Z. (2020). Deep robust cramer shoup delay optimized fully homomorphic for iiot secured transmission in cloud computing. *Computer Communications*, 161:10–18.

Lv, Y. (2021). Data privacy protection based on homomorphic encryption. In *Journal of Physics: Conference Series*, volume 2037, page 012129. IOP Publishing.

Maimuţ, D. and Teşeleanu, G. (2020). A new generalisation of the goldwasser-micali cryptosystem based on the gap 2 k-residuosity assumption. In *International Conference on Information Technology and Communications Security*, pages 24–40. Springer.

Maple, C. (2017). Security and privacy in the internet of things. *Journal of cyber policy*, 2(2):155–184.

Martin, J. (1991). *Rapid application development*. Macmillan Publishing Co., Inc.

Matias, C., Ivaki, N., and Moraes, R. (2023). Exploring the impact of homomorphic encryption on the performance of machine learning algorithms. In *Proceedings of the 12th Latin-American Symposium on Dependable and Secure Computing*, pages 120–125.

Menezes, A. J., Van Oorschot, P. C., and Vanstone, S. A. (2018). *Handbook of applied cryptography*. CRC press.

Menon, S. J. and Wu, D. J. (2022). Spiral: Fast, high-rate single-server pir via fhe composition. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 930–947. IEEE.

Meyer, A. N., Giardina, T. D., Spitzmueller, C., Shahid, U., Scott, T. M., and Singh, H. (2020). Patient perspectives on the usefulness of an artificial intelligence–assisted symptom checker: cross-sectional survey study. *Journal of medical Internet research*, 22(1):e14679.

Oh, E. N., Baharon, M. R., Yassin, S., Idris, A., and MacDermott, A. (2022). Preserving data privacy in mobile cloud computing using enhanced homomorphic encryption scheme. In *Journal of Physics: Conference Series*, volume 2319, page 012024. IOP Publishing.

Omollo, R., Raburu, G., Omolo-Ongati, N., and Okelo, B. (2017). Enhancing data security in cloud computation using addition-composition fully homomorphic encryption scheme.

Park, J., Kim, D. S., and Lim, H. (2020). Privacy-preserving reinforcement learning using homomorphic encryption in cloud computing infrastructures. *IEEE Access*, 8:203564–203579.

Rao, B. S., Chattopadhyay, S., Singh, P., Hazela, B., Sabarinathan, G., and Yamini, K. (2023). Privacy-aware artificial intelligence with homomorphic encryption using machine learning. In *2023 International Conference on Sustainable Computing and Smart Systems (ICSCSS)*, pages 259–265. IEEE.

Rivest, D. R., Shamir, A., and Adleman, L. (1977). Rsa (cryptosystem). *Arithmetic Algorithms And Applications*, page 19.

Schadendorf, D., Van Akkooi, A. C., Berking, C., Griewank, K. G., Gutzmer, R., Hauschild, A., Stang, A., Roesch, A., and Ugurel, S. (2018). Melanoma. *The Lancet*, 392(10151):971–984.

Scheibner, J., Raisaro, J. L., Troncoso-Pastoriza, J. R., Ienca, M., Fellay, J., Vayena, E., and Hubaux, J.-P. (2021). Revolutionizing medical data sharing using advanced privacy-enhancing technologies: technical, legal, and ethical synthesis. *Journal of medical Internet research*, 23(2):e25120.

Sen, J. (2013). Homomorphic encryption-theory and application. *Theory and practice of cryptography and network security protocols and technologies*, 31.

Shaheen, M. Y. (2021). Applications of artificial intelligence (ai) in healthcare: A review. *ScienceOpen Preprints*.

Shen, S., Xu, M., Zhang, F., Shao, P., Liu, H., Xu, L., Zhang, C., Liu, P., Yao, P., and Xu, R. X. (2022). A low-cost high-performance data augmentation for deep learning-based skin lesion classification. *BME frontiers*.

Shollo, A., Hopf, K., Thiess, T., and Müller, O. (2022). Shifting ml value creation mechanisms: A process model of ml value creation. *The Journal of Strategic Information Systems*, 31(3):101734.

Sommerville, I. (2011). *Software Engineering, 9/E*. Pearson Education India.

streamlit.io (2024). Streamlit Docs — docs.streamlit.io. https://docs.streamlit.io/. [Accessed 09-04-2024].

Su, Y., Yang, B., Yang, C., and Tian, L. (2020). Fpga-based hardware accelerator for leveled ring-lwe fully homomorphic encryption. *IEEE Access*, 8:168008–168025.

tensorflow.org (2024). tf.keras.applications.ResNet50 | TensorFlow v2.16.1 — tensorflow.org. https://www.tensorflow.org/api_docs/python/tf/keras/applications/ResNet50. [Accessed 09-04-2024].

Turan, F., Roy, S. S., and Verbauwhede, I. (2020). Heaws: An accelerator for homomorphic encryption on the amazon aws fpga. *IEEE Transactions on Computers*, 69(8):1185–1196.

Voigt, P. and Von dem Bussche, A. (2017). The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, 10(3152676):10–5555.

Wang, Z. and Hong, T. (2020). Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy*, 269:115036.

Xiao, Y., Liu, A., Zhang, T., Qin, H., Guo, J., and Liu, X. (2023). Robustmq: benchmarking robustness of quantized models. *Visual Intelligence*, 1(1):30.

Zama.ai (2024). What is Concrete ML? | 1.5 | Concrete ML — docs.zama.ai. https://docs.zama.ai/concrete-ml/getting-started/getting-started. [Accessed 09-04-2024].

Zhang, C. and Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23:100224.

Zhang, L., Li, C., Hu, Q., Lang, J., Huang, S., Hu, L., Leng, J., Chen, Q., and Lv, C. (2023). Enhancing privacy in large language model with homomorphic encryption and sparse attention. *Applied Sciences*, 13(24):13146.

Appendix A
Similarity Index

Enhancing Healthcare Data Privacy with Homomorphic Encryption.pdf

by John Mugo

Submission date: 30-Mar-2025 02:15PM (UTC+0300)

Submission ID: 2629385975

File name: 44577_John_Mugo_Enhancing_Healthcare_Data_Privacy_with_Homomorphic_Encryption_229873_720353364.pdf (2.72M)

Word count: 14056

Character count: 85922

Enhancing Healthcare Data Privacy with Homomorphic Encryption.pdf

ORIGINALITY REPORT

15%	13%	13%	10%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	su-plus.strathmore.edu Internet Source	1%
2	B. Srinivasa Rao, Saumitra Chattopadhyay, Prashant Singh, Bramah Hazela, G. Sabarinathan, Kalva Yamini. "Privacy-Aware Artificial Intelligence with Homomorphic Encryption using Machine Learning", 2023 International Conference on Sustainable Computing and Smart Systems (ICSCSS), 2023 Publication	1%
3	peerj.com Internet Source	1%
4	hdl.handle.net Internet Source	1%
5	flore.unifi.it Internet Source	<1%
6	www.mdpi.com Internet Source	<1%
7	www.econstor.eu Internet Source	<1%
8	arxiv.org Internet Source	<1%

45	Internet Source	<1%
46	jurnal.itscience.org Internet Source	<1%
47	re.public.polimi.it Internet Source	<1%
48	www.ecohumanism.co.uk Internet Source	<1%

Exclude quotes	Off	Exclude matches	< 25 words
Exclude bibliography	Off		

Appendix B
Ethical Approval

Strathmore UNIVERSITY

11th September 2024

Mr Mugo John,
john.mugo@strathmore.edu

Dear Mr Mugo,

RE: Enhancing Healthcare Data Privacy with Homomorphic Encryption

This is to inform you that SU-ISERC has reviewed and **approved** your above **SU-masters** proposal. Your application reference number is **SU-ISERC2371/24.** The approval period is from **11th September 2024 to 10th September 2025.**

This approval is subject to compliance with the following requirements:

i. Only approved documents including (informed consents, study instruments, MTA) will be used.
ii. All changes including (amendments, deviations, and violations) are submitted for review and approval by SU-ISERC.
iii. Death and life-threatening problems and serious adverse events or unexpected adverse events whether related or unrelated to the study must be reported to SU-ISERC within 72 hours of notification.
iv. Any changes anticipated or otherwise that may increase the risks or affected safety or welfare of study participants and others or affect the integrity of the research must be reported to SU-ISERC within 72 hours.
v. Clearance for the export of biological specimens must be obtained from relevant institutions.
vi. Submission of a request for renewal of approval at least 60 days prior to the expiry of the approval period. Attach a comprehensive progress report to support the renewal.
vii. Submission of an executive summary report within 90 days of completion of the study to SU-ISERC.

Before commencing your study, you will be expected to obtain a research license from National Commission for Science, Technology, and Innovation (NACOSTI) https://research-portal.nacosti.go.ke/ and obtain other clearances needed.

Yours sincerely,

Mr Ambrose Rachier,
Chairperson; SU-ISERC

Ole Sangale Rd, Madaraka Estate. PO Box 59857-00200, Nairobi, Kenya. Tel +254 (0)703 034000
Email admissions@strathmore.edu www.strathmore.edu

Appendix C
Python Code

C.1 Main Page Code

app.py > is_skin_lesion_image

```python

import patch

import streamlit as st
import cv2
import os
import numpy as np
from skimage import io as skio
from pathlib import Path
import io
import gc
from encryption_utils import FHEPipeline
from inference import MelanomaFHESystem
from educational_content import display_educational_content

# Initialize components
fhe_system = MelanomaFHESystem()
fhe_client = FHEPipeline()

def is_skin_lesion_image(image):
    """Check if the image contains skin using HSV color thresholds"""
    # Convert to RGB if it's RGBA
    if len(image.shape) == 3 and image.shape[2] == 4:
        image = image[:, :, :3]

    # Convert to HSV for better skin detection
    hsv_img = cv2.cvtColor(image, cv2.COLOR_RGB2HSV)

    lower_threshold = np.array([0, 10, 60], dtype=np.uint8)
    upper_threshold = np.array([35, 255, 255], dtype=np.uint8)

    skin_mask = cv2.inRange(hsv_img, lower_threshold, upper_threshold)

    skin_percentage = (np.sum(skin_mask > 0) / (image.shape[0] * image.shape[1])) * 100

    return skin_percentage > 30

```

app.py > main

```python
def main():
    st.set_page_config(
        page_title="Secure Melanoma Diagnosis",
        page_icon="🔒",
        layout="wide"
    )

    st.title("🔐 FHE-Enhanced Melanoma Detection")
    st.markdown("""
    This application uses Fully Homomorphic Encryption (FHE) to securely analyze skin lesions
    while preserving patient privacy. Your image is encrypted before analysis.
    """)

    nav = st.sidebar.radio("Navigation", ["Diagnosis", "Educational Content"])

    if nav == "Diagnosis":
        uploaded_file = st.file_uploader("Upload dermoscopic image", type=["jpg", "jpeg", "png"])

        if uploaded_file:
            # Check if the uploaded image is a skin lesion
            try:
                pos = uploaded_file.tell()
                image = skio.imread(uploaded_file)
                uploaded_file.seek(0)

                if len(image.shape) < 2:
                    st.warning("Invalid image format. Please upload a proper image file.")
                    st.stop()

                if not is_skin_lesion_image(image):
                    st.error("The uploaded image does not appear to be a skin lesion. Please upload a dermoscopic image of a skin
                    st.stop()

            except Exception as e:
                st.error(f"Error validating image: {str(e)}")
                st.stop()
```

C.2 Encryption Code

encryption_utils.py > FHEPipeline > encrypt

```python
class FHEPipeline:

    def encrypt(self, image_array):
        """Encrypt an image array for FHE processing"""
        if self.client is None:
            raise ValueError("FHE client not initialized. Call initialize() first.")

        if isinstance(image_array, bytes):
            raise ValueError("Input must be a numpy array, not bytes")

        if not isinstance(image_array, np.ndarray):
            raise ValueError("Input must be a numpy array")

        if image_array.shape[1:3] != (128, 128):
            logger.warning(f"Expected input shape with dimensions (128, 128), got {image_array.shape}. Resizing...")
            import cv2
            if len(image_array.shape) == 4:
                resized = np.array([cv2.resize(img, (128, 128)) for img in image_array])
            else:
                resized = cv2.resize(image_array, (128, 128))
            image_array = resized

        if len(image_array.shape) == 3 and image_array.shape[2] == 4:
            logger.info("Converting RGBA to RGB by removing alpha channel")
            image_array = image_array[:, :, :3]
        elif len(image_array.shape) == 4 and image_array.shape[3] == 4:
            logger.info("Converting batch of RGBA to RGB by removing alpha channel")
            image_array = image_array[:, :, :, :3]

        image_array = image_array.astype(np.float32)

        if len(image_array.shape) == 3:
            image_nhwc = np.expand_dims(image_array, axis=0)
        else:
            image_nhwc = image_array
```

encryption_utils.py > FHEPipeline > decrypt

```python
class FHEPipeline:

    def decrypt(self, encrypted_result, threshold=None):
        """Decrypt FHE result and return diagnosis information

        Args:
            encrypted_result: The encrypted prediction from the FHE model
            threshold: Optional confidence threshold (uses calibrated threshold if available)
        """
        if self.client is None:
            raise ValueError("FHE client not initialized. Call initialize() first.")

        # Ensure encrypted_result is not None
        if encrypted_result is None:
            raise ValueError("Encrypted result is None")

        # Handle the encrypted result
        try:
            decrypted_result = self.client.deserialize_decrypt_dequantize(encrypted_result)

            gc.collect()

            confidence = float(decrypted_result[0][0])

            if threshold is None:
                threshold = 0.85

            # Generate diagnosis result
            diagnosis = {
                "malignant": confidence > threshold,
                "confidence": confidence,
                "risk_factors": {
                    "biopsy_recommended": confidence > threshold * 1.05,
                    "monitoring_interval": "3 months" if confidence > threshold * 0.7 else "1 year"
                }
            }
```

C.3 Inference Code

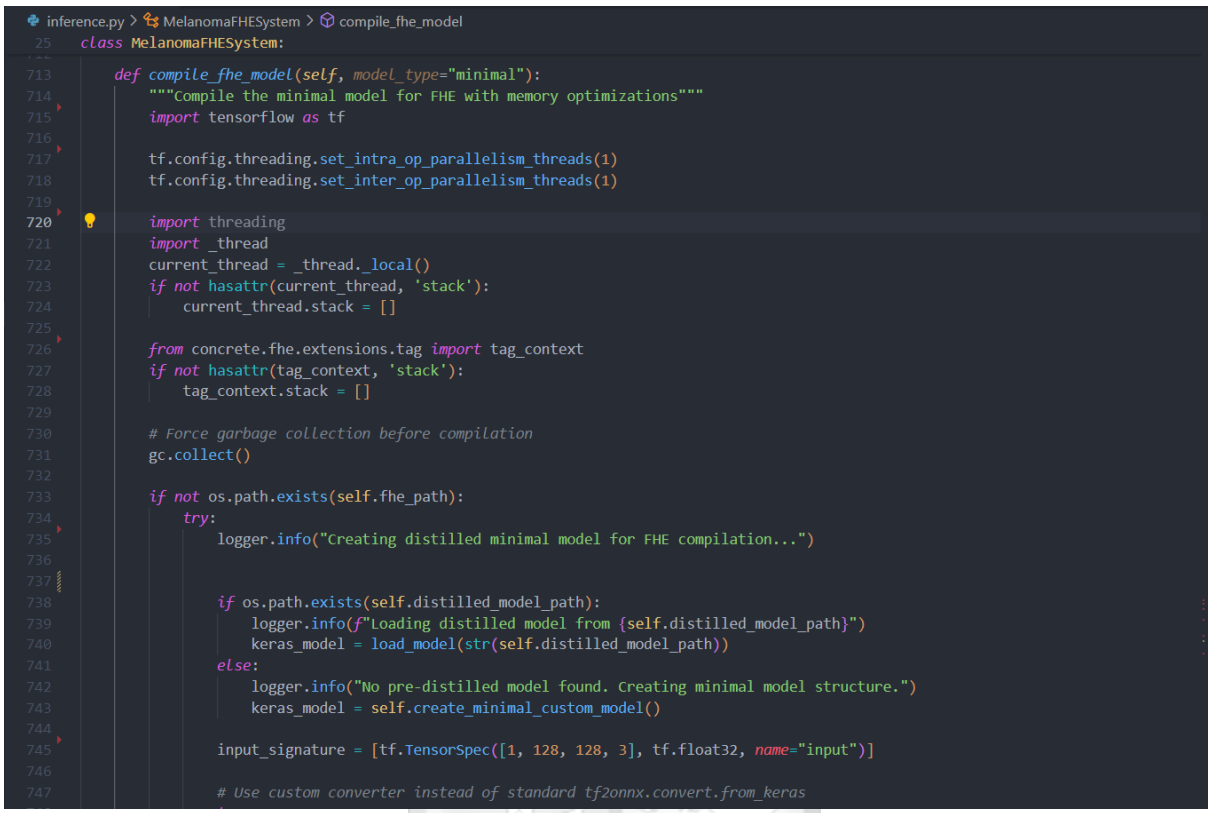

```
inference.py > MelanomaFHESystem > compile_fhe_model
class MelanomaFHESystem:

    def compile_fhe_model(self, model_type="minimal"):
        """Compile the minimal model for FHE with memory optimizations"""
        import tensorflow as tf

        tf.config.threading.set_intra_op_parallelism_threads(1)
        tf.config.threading.set_inter_op_parallelism_threads(1)

        import threading
        import _thread
        current_thread = _thread._local()
        if not hasattr(current_thread, 'stack'):
            current_thread.stack = []

        from concrete.fhe.extensions.tag import tag_context
        if not hasattr(tag_context, 'stack'):
            tag_context.stack = []

        # Force garbage collection before compilation
        gc.collect()

        if not os.path.exists(self.fhe_path):
            try:
                logger.info("Creating distilled minimal model for FHE compilation...")


                if os.path.exists(self.distilled_model_path):
                    logger.info(f"Loading distilled model from {self.distilled_model_path}")
                    keras_model = load_model(str(self.distilled_model_path))
                else:
                    logger.info("No pre-distilled model found. Creating minimal model structure.")
                    keras_model = self.create_minimal_custom_model()

                input_signature = [tf.TensorSpec([1, 128, 128, 3], tf.float32, name="input")]

                # Use custom converter instead of standard tf2onnx.convert.from_keras
```

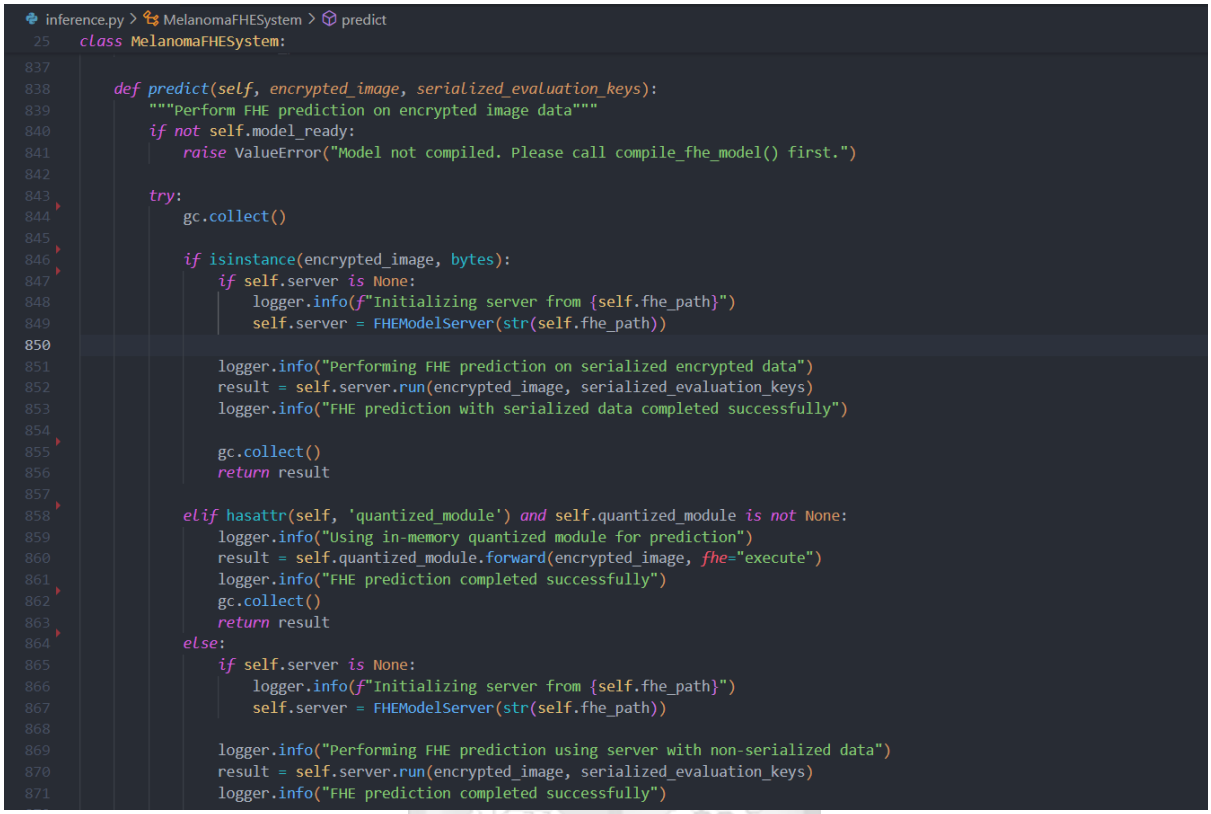

C.4 Educational Content Code

educational_content.py > display_educational_content

```python
import streamlit as st

def display_educational_content():
    st.header("Medical Resource Center")

    # FHE and Technology Section
    st.markdown("## How FHE Enhances Privacy")
    st.write("""
    Fully Homomorphic Encryption (FHE) allows computation on encrypted data without decryption:

    1. Your image is encrypted
    2. Analysis happens on the encrypted data
    3. Only the final result is decrypted
    4. Your original image data remains private throughout

    **Benefits of FHE:**
    - Protects sensitive medical data from unauthorized access
    - Allows for secure outsourcing of computations without compromising privacy
    - Enhances trust in AI-assisted medical diagnostics
    - Ensures HIPAA compliance with encryption
    """)

    # Clinical Workflow in Columns
    col1, col2 = st.columns([3, 2])

    with col1:
        st.markdown("## Clinical Information")
        st.write("""
        **Technology Details:**
        - **Model**: Optimized lightweight CNN for minimal resource usage
        - **Confidence Threshold**: Automatically calibrated for optimal accuracy
        - **Security**: FHE-enabled encryption, HIPAA-compliant

        **Clinical Workflow:**
        1. **Image Upload:** Users upload dermoscopic images
        2. **Homomorphic Inference:** The image is encrypted and processed by the FHE-enabled model
        3. **Secure Result Decryption:** Results are decrypted and provided
```

educational_content.py > display_educational_content

```python
def display_educational_content():

        **Clinical Workflow:**
        1. **Image Upload:** Users upload dermoscopic images
        2. **Homomorphic Inference:** The image is encrypted and processed by the FHE-enabled model
        3. **Secure Result Decryption:** Results are decrypted and provided
        4. **Risk Assessment Generation:** Provides clinical guidance based on the diagnosis
        """)

    with col2:
        st.markdown("## About Melanoma")
        st.write("""
        Melanoma is a serious form of skin cancer that begins in melanocytes (cells that make melanin, the pigment that g

        Early detection is critical for successful treatment. Regular skin checks and prompt medical attention for suspic

        ### Early Detection Matters
        - **99% Survival Rate:** When detected early, melanoma has a high survival rate
        - **Monthly Self-Exams:** Regular self-examinations can help identify changes in moles
        - **Annual Dermatologist Visits:** Regular check-ups are crucial for early detection
        """)

    # Video Section
    st.markdown("## Educational Video")
    video_data = "https://youtube.com/watch?v=-uf1mOu98V8"
    _, video_container, _ = st.columns([1, 3, 1])
    with video_container:
        st.video(video_data)

```

Appendix D
Project Timeline