DESIGN AND IMPLEMENTATION OF AN AUTOMATED BLURRY IDENTITY DOCUMENT IMAGE DETECTION SYSTEM

BY
CHUKWUKA BEJIH
U16/NAS/CSC/153

DEPARTMENT OF COMPUTER SCIENCE AND MATHEMATICS, FACULTY OF NATURAL SCIENCES AND ENVIRONMENTAL STUDIES,
GODFREY OKOYE UNIVERSITY, ENUGU STATE.

IN PARTIAL FULFILLMENT OF THE AWARD OF BACHELOR OF SCIENCE (B.SC), DEGREE IN COMPUTER SCIENCE.

SUPERVISOR: DR. FRANKLIN UCHENNA OKEBANAMA

MAY, 2026.

APPROVAL PAGE

This research, titled "**DESIGN AND IMPLEMENTATION OF AN AUTOMATED BLURRY IDENTITY DOCUMENT IMAGE DETECTION SYSTEM** ," has been assessed and approved by the Department of Computer Science and Mathematics Committees in the Faculty of Natural Sciences and Environmental Studies, Godfrey Okoye University, Enugu, Nigeria.

................................…	……………………..	……………………..
Supervisor	Signature	Date
External Examiner	....	....
External Examiner	Signature	Date
Dr. Franklin Uchenna Okebanama	....	....
HOD, Department of Computer Science and Mathematics	Signature	Date
Prof. N. M. Ozofor	....	....
Dean, Faculty of Computing Date And Information Technology.	Signature	

CERTIFICATION

I certify that I completed the research included therein and that it was primarily based on my observations and investigation. Consequently, I will fully accept the University's decision if I am found to have cheated on the assignment.

CHUKWUKA IBEJIH
U16/NAS/CSC/153

DEDICATION

This project is dedicated to my parents, **Engr. and Mrs. Chukwuemeka Ibejih** and my Late Grand-mother, **Mama Rosemary Ikpeka**, for their unwavering love, sacrifices, and prayers, which have been the foundation of every achievement in my life, and also, to everyone passionate about solving human problems through technology.

ACKNOWLEDGMENTS

All glory and honour belong to God Almighty, whose grace sustained me through every stage of this project.

I wish to express my profound gratitude to my supervisor and Head of Department, Dr. Franklin Uchenna Okebanama, for his patient guidance, intellectual generosity, and constructive criticism throughout the course of this research. His insightful feedback challenged me to think more rigorously and shaped the quality of this work in ways I deeply appreciate. I am also grateful to the entire faculty and Department for providing an environment that encourages academic excellence and for every lesson and piece of advice offered across my years of study.

To my parents, Engr. Chukwuemeka Ibejih and Mrs. Regina Ibejih. I owe, more than words can express! Your sacrifices, prayers, and unshakeable belief in my potential have been my greatest source of strength. I also appreciate my elder sisters, Kachi and Kelly, my younger brother, Nonso; I thank you immensely, for every of your support and encouragement. Those moments of levity kept me going during the most demanding periods of this work.

Finally, I acknowledge everyone who contributed to this project in any form through a conversation, a resource, a suggestion, or a word of encouragement. This work is as much yours as it is mine!

ABSTRACT

Identity verification systems in banking, telecommunications, and government registration increasingly require users to upload photographs of identity documents captured on smartphones under uncontrolled conditions. Blur, caused by hand movement during exposure or incorrect focal distance, is one of the most frequent and damaging quality problems affecting such images, as it weakens edge sharpness, destroys fine character detail, and renders documents unreadable for optical character recognition (OCR) and biometric verification. Existing systems accept uploads based on basic technical constraints such as file format, file size, and resolution, none of which detect visual blur, causing blurry images to enter the verification pipeline and produce downstream failures, repeated submissions, and increased manual review workload. This study designed and implemented an Automated Blurry Identity Document Image Detection System capable of assessing image quality at the point of upload and rejecting unusable submissions before verification begins. The system employs a hybrid classical computer vision approach combining Laplacian variance, which measures edge sharpness through second-order derivatives, and Tenengrad gradient magnitude, which evaluates edge strength using Sobel operators. A composite blur score is computed by normalising and fusing both metrics into a single value on a scale of 0 to 100, which is then used to assign one of four quality grades which are Good, Mild, Moderate, or Severe and compared against an administrator-configurable acceptance threshold to produce an acceptance or rejection decision. The system was implemented using a FastAPI backend with an OpenCV-based detection engine and a React-based user interface, with all upload records, scores, grades, and decisions persisted to a PostgreSQL database to support auditing and future threshold optimisation. Evaluation on a 100-image synthetic dataset constructed using Gaussian blur at controlled kernel parameters yielded an overall classification accuracy of 94.00%, a precision of 100.00%, a recall of 88.00%, an F1-score of 93.62%, a false acceptance rate of 0.00%, and an average processing time of 18.26 milliseconds per image, confirming that the system meets real-time performance requirements without requiring GPU hardware. The results demonstrate that the proposed hybrid classical approach outperforms single-metric methods and provides an effective, explainable, and configurable quality gate that reduces verification failures, eliminates unnecessary manual review, and improves user experience in digital identity verification workflows.

Table of Contents

Chapter Five: Summary, Recommendations, and Conclusion

LIST OF TABLES

LIST OF FIGURES

CHAPTER ONE: INTRODUCTION

1.1 Background of the Study

Digital onboarding and remote identity verification have become central to many services such as banking, telecommunications, education, e-commerce, and government registration etc. In these systems, users are often required to upload images of identity documents such as national ID cards, voter cards, passports, or driver’s licenses. These images are then used for automated verification steps including document type detection, optical character recognition (OCR), face extraction, and authenticity checks.

In practice, identity document images are commonly captured using smartphone cameras under uncontrolled conditions. Users may take photos in poor lighting, with hand movement, incorrect focus, glare, shadows, or at an angle. Among these issues, blur remains one of the most frequent and most damaging problems because it reduces edge sharpness and destroys fine details needed to read text and validate security features. A document image may meet basic upload rules like file type, file size, or resolution, yet be unusable due to blur. When blurry documents are accepted into the pipeline, downstream verification often fails, causing delays, repeated submissions, and increased manual review workload.

Because of these realities, there is a growing need for automated document image quality checks at the point of upload. Rather than relying only on technical checks, a system should assess whether the image content is visually usable for identity verification. This dissertation focuses on building an automated system that detects blurry identity document images and prevents unusable uploads from entering the verification workflow.

1.2 Statement of the Problem

Many identity verification systems still accept document uploads as long as basic technical requirements are satisfied, such as image format, minimum resolution, or maximum file size. These checks do not guarantee that the document is readable. Users frequently submit images that are blurred due to motion or incorrect focus, and these images often lead to OCR failure, inability to extract key fields, or unsuccessful identity verification.

The lack of an automated, reliable blur detection mechanism at the point of upload creates several issues:

1. Increased verification failure rates and repeated capture attempts
2. Higher operational costs due to manual review of poor-quality uploads
3. Longer onboarding time for users
4. Poor user experience and frustration
5. Reduced trust in digital onboarding systems

There is therefore a need to develop an automated system that can detect whether an uploaded identity document image is blurry and determine if it is suitable for downstream verification tasks.

1.3 Aim and Objectives of the Study

This study set out to design and implement a dedicated automated system capable of assessing identity document image quality at the moment of upload, with the goal of ensuring that only visually usable images proceed into downstream identity verification workflows.

To achieve this aim, the study pursued the following specific objectives:

1. Investigate the primary causes and observable characteristics of blur in identity document images taken under real-world, uncontrolled capture conditions.
2. Collect or construct a dataset of identity document images with varying blur levels and capture conditions.
3. Design an image processing pipeline for assessing document image sharpness and detecting blur.

4. Implement and evaluate a blurriness detection approach using classical image processing metrics and or a machine learning classifier.
5. Define a thresholding or decision strategy for classifying images as usable or unusable based on blurriness.

6. Measure the effectiveness of the implemented system against standard performance indicators including classification accuracy, precision, recall, and average processing latency.

1.4 Significance of the Study

The practical and academic value of this study lies in its direct contribution to more dependable digital identity verification. Among its specific contributions:

- Reduces verification failures by preventing unusable blurry uploads from entering the pipeline
- Saves operational time and cost, by reducing manual review and repeated submissions
- Improves user experience by providing immediate rejection and guidance before downstream failure occurs
- Supports organisations that rely on KYC by providing a practical quality control layer
- Provides a framework and implementation approach that can be extended to other quality problems such as glare, low illumination, or occlusion

1.5 Scope of the Study

This work focuses on detecting blur in images of identity documents submitted for verification. The scope includes:

- Identity document images captured using smartphones or consumer cameras
- Blur detection at the point of upload, before OCR and verification stages

- Implementation of blur scoring and decision making using image processing techniques and optionally machine learning classification
- Evaluation of the system using a dataset of document images with different blur conditions

This study does not attempt to perform full identity verification, document forgery detection, or OCR optimization rather it focuses specifically on blur detection as a quality-gating mechanism.

1.6 Limitation of the Study

The study may face the following limitations:

1. **Dataset limitations:** availability of diverse identity document images and realistic blur conditions may affect generalisation.
2. **Mixed distortions:** blur often co-occurs with glare or low light, which can make 'blur-only' decisions harder in borderline cases.
3. **Device variability:** camera quality differs across smartphones, which can influence blur patterns and sharpness measurements.
4. **Threshold sensitivity:** classical blur metrics may require careful threshold tuning across document types and capture conditions.

1.7 Operational Definition of Terms

- **Identity Document Image:** A digital photograph of an official government-issued identification document such as a national identity card, driver's licence, or passport submitted by a user through an upload interface for the purpose of identity verification.

- **Blur:** The degradation of image sharpness resulting from motion during capture, incorrect focal distance, or camera instability, characterised by the weakening of edge boundaries and the loss of fine spatial detail. In the context of this study, blur is measured quantitatively using the Laplacian variance and Tenengrad gradient magnitude metrics.

- **No-Reference Image Quality Assessment (NR-IQA):** A category of image quality evaluation that operates solely from the submitted image, without access to an original or reference version for comparison. This approach is the only viable option in identity verification contexts, where no pristine reference document image exists.

- **Blur Detection:** The automated process of computing a quantitative sharpness score from an uploaded image and comparing that score against a defined threshold to determine whether the image is of sufficient quality for downstream verification processing.

- **Composite Blur Score:** A single numeric quality value in the range of 0 to 100 produced by normalising and combining the Laplacian variance and Tenengrad gradient magnitude scores using equal weighting. Higher values indicate greater sharpness. This metric is the primary output of the blur detection engine in the proposed system.

- **Quality Grade:** A categorical label of Good, Mild, Moderate, or Severe assigned to an uploaded image based on the range within which its composite blur score falls. The grade provides a human-readable description of image quality alongside the numeric score.

- **Quality Gate:** A validation checkpoint integrated into the document upload workflow that evaluates each submitted image against a defined quality standard and blocks non-compliant submissions from entering the verification pipeline.

- **Acceptance Threshold:** The minimum composite blur score that an uploaded image must achieve in order to be accepted by the system. Images scoring at or above the threshold are accepted; those scoring below are rejected. The threshold is configurable by the administrator through the system dashboard without requiring code modification.

- **OCR (Optical Character Recognition):** A technology that reads and converts printed or handwritten text embedded in an image into machine-readable character data. OCR accuracy is directly sensitive to image blur and is one of the primary downstream processes protected by the quality gate implemented in this study.

- **False Acceptance Rate (FAR):** The proportion of blurry or unusable images that the system incorrectly classifies as acceptable and allows to pass through the quality gate. A FAR of 0.00% was achieved in the evaluation of the proposed system, indicating that no unusable image was incorrectly accepted.

- **Gaussian Blur:** A mathematically defined image degradation technique that applies a Gaussian-weighted smoothing kernel to an image, progressively reducing edge sharpness as kernel size and

  standard deviation increase. In this study, Gaussian blur was applied at controlled parameters to construct the synthetic evaluation dataset.

CHAPTER TWO: LITERATURE REVIEW

2.1 INTRODUCTION

The quality of images submitted through digital verification systems has a direct bearing on how reliably those systems perform. Automated verification workflows typically involve multiple processing stages, among them document localisation, character recognition, biometric face comparison, and document authenticity checks, each of which requires sufficient visual clarity to function correctly. When image quality falls short, any of these stages may produce errors or fail entirely.

The widespread adoption of smartphone-based document capture has changed how identity documents are submitted for verification. National identity cards, passports, and driver's licences are now routinely photographed by end users in everyday settings at home, outdoors, or in offices without controlled lighting or camera stabilisation. Under these conditions, captured images are susceptible to a range of quality problems, with blur consistently emerging as one of the most prevalent and consequential.

When blur is present in a document image, the fine visual information that automated systems depend on degrades rapidly. Character strokes lose their definition, boundary regions between fields become indistinct, and the overall structural legibility of the document deteriorates. The result is an image that may appear visually acceptable to a casual observer yet be effectively unprocessable by automated verification tools. There is strong indication that basic upload-validation rules, such as resolution or file size, are insufficient to guarantee document usability. Images may satisfy technical constraints while remaining unsuitable for OCR or verification due to blur, uneven illumination, or surface artifacts. Consequently, many real-world document images that pass standard checks still fail during downstream processing.

Identity documents differ substantially from natural images because they contain dense text, fine strokes, structured layouts, and security elements that are highly sensitive to blur. For this reason, identity document imaging must be evaluated in terms of task fitness rather than visual quality alone. This creates the need for automated quality screening mechanisms capable of detecting blur before verification begins.

This chapter reviews theories related to image blur and document image analysis, examines methods and systems proposed for blur detection, and identifies gaps in existing research that motivate the proposed system.

2.2 Theoretical Background

The subsections that follow establish the theoretical grounding for this research. Each covers a domain of knowledge, from how blur forms in images, to how quality is assessed, to how features are extracted that directly shaped the design decisions made in this study.

2.2.1 Image Blur and Blur Formation Theory

Li et al., (2018) defined image blur as the degradation of image sharpness caused by the loss of high-frequency spatial information. In blurred images, edges become weakened and fine details disappear, reducing the details of visual structures. In the context of document images, this degradation directly affects character strokes, layout boundaries, and structural elements that are essential for optical character recognition (OCR) and identity verification.

Zhang et al. (2019) identified two dominant sources of blur in camera-captured document images: motion blur and defocus blur. The first arises when the camera and document move relative to one another during the exposure window, a situation commonly produced by hand tremor or an unstable capture posture. The second occurs when the camera's focal plane fails to coincide with the document surface, either because the focal distance is set incorrectly or because the scene depth exceeds the camera's depth of field. Mathematically, both degradation types can be represented as the convolution of an ideal sharp image with a point spread function, a process that attenuates high-spatial-frequency content and thereby eliminates the edge detail that verification algorithms require.

In real-world capture scenarios, blur is frequently spatially non-uniform rather than evenly distributed across the image. Some regions of a document may remain relatively sharp while others experience significant degradation. Such spatial variation reduces the effectiveness of global sharpness measures and motivates localized or region-aware blur analysis Maheshwari et al., (2016) and Golestaneh and Karam (2017). This characteristic is particularly critical for identity documents, where partial blur affecting key regions such as text fields or portrait areas can cause verification failure even when the overall image appears visually acceptable.

Blur also exhibits multi-scale behaviour, with fine-scale details degrading earlier than coarse structures. As a result, document images may appear acceptable at a global level while still failing OCR or verification tasks that depend on small character strokes and sharp boundaries. According to Li and Jiang (2019), this scale-dependent behaviour explains why visual inspection alone is insufficient for assessing document image usability.

Empirical observations further indicate that blur rarely occurs in isolation. It frequently coexists with other degradations such as uneven illumination and background texture, which complicate detection and analysis. Moreover, Zhou et al., (2023) and Zhang et al., (2021) have shown that non-expert users often struggle to consistently capture sharp document images under uncontrolled conditions, increasing the prevalence of blur in real-world identity verification systems.

2.2.2 Image Quality Assessment (IQA) Theory

Computational approaches to estimating image suitability whether for human perception or automated analysis are collectively referred to as Image Quality Assessment (IQA). The field is broadly organised around the availability of a reference image: full-reference methods compare a degraded image against an original, reduced-reference methods use partial reference information, and no-reference methods operate from the captured image alone. In document verification contexts, a reference version of any given submitted document does not exist, which means that no-reference approaches are the only practically applicable category.

Shemiakina et al. (2021) noted that no-reference IQA is particularly suitable for document images because quality must be estimated directly from the captured image without prior knowledge of its original appearance. Recent IQA theory emphasizes task-oriented quality assessment, where image quality is defined by its ability to support downstream tasks rather than subjective visual appeal.

Document-specific IQA theory further argues that quality metrics should correlate with OCR reliability and verification success. Alaei et al. (2023) reported that blur is one of the most challenging document degradations because it interacts with layout structure, text density, and illumination variations. This theoretical shift supports the use of OCR-aware and task-aware quality indicators in document image analysis.

2.2.3 Feature Extraction Techniques for Blur Detection

Feature extraction techniques are used to capture visual characteristics that change in the presence of blur. Edge-based features analyze gradient magnitude, edge density, and edge spread. Since blur weakens edges and smooths transitions, these features provide direct indicators of blur severity.

Frequency-domain techniques analyze how image energy is distributed across frequency bands. It was reported by Wang et al. (2018), that blur causes a measurable reduction in high-frequency components, making spectral decomposition tools such as the Fourier Transform and the Discrete Cosine Transform well suited to identifying this form of degradation.

Texture-based features examine changes in local spatial patterns caused by blur. As observed by Gao et al. (2020), texture statistics can capture subtle degradation in document images that is not easily detected by global sharpness measures. Together, edge-based, frequency-domain, and texture-based features form the theoretical basis of classical blur detection systems.

2.2.4 Machine Learning and Computer Vision in Image Quality Analysis

The application of machine learning to image quality analysis allows systems to recognise degradation patterns from data rather than relying on manually specified rules. Earlier approaches in this category combined manually engineered feature representations with statistical classifiers, of which support vector machines were among the most widely used.

More recent approaches employ convolutional neural networks (CNNs), which derive feature representations from training images without manual specification. Schlett et al. (2022) demonstrated that CNNs trained for image quality tasks develop sensitivity to blur across multiple spatial resolutions simultaneously, giving them greater flexibility than handcrafted feature pipelines when handling varied and complex degradation patterns.

Automated image quality analysis is preferred over manual inspection due to its consistency, scalability, and real-time capability. Hernandez-Ortega et al., (2022) revealed that automated quality prediction strongly correlates with biometric verification failure, reinforcing the importance of automated blur detection in identity document processing pipelines.

2.3 Review of Related Literature

The following section examines published empirical work in two interconnected areas: automated identity document processing and blur detection. Each study is considered with attention to the method employed, the data used, the principal findings reported, and the shortcomings that remain unaddressed.

2.3.1 Studies on Automated Identity Document Image Processing Systems

Several researchers have established the value of automation in document processing pipelines. Awal et al. (2017) and Arlazarov et al. (2019) demonstrated that automated approaches significantly reduce the manual effort required during document analysis, while Bulatov et al. (2021), Attivissimo et al. (2019), and Appalaraju et al. (2021) each confirmed that automation improves both the speed and reliability of document understanding tasks. Earlier studies in this group relied on rule-based and handcrafted feature pipelines, whereas more recent contributions increasingly adopt deep learning architectures. A notable limitation shared across these works is the assumption that document images entering the system are of sufficient visual quality. Arlazarov et al. (2019) and Bulatov et al. (2021) both implicitly treated image clarity as a prerequisite rather than a problem to be solved, with the consequence that blur detection prior to verification received little focused attention.

However, earlier studies have proposed automated systems for identity document detection and verification. Arlazarov et al. (2019) explored a standard dataset for identity document analysis captured under uncontrolled conditions, demonstrating that recognition accuracy deteriorates significantly under poor image quality. Expanding on this work, Bulatov et al. (2021) provided a large-scale benchmark with diverse capture scenarios, showing that blur is one of the dominant causes of recognition failure. The shortfall of the standard datasets is the limitation of availability of data to be used.

Considering the importance of early quality assessment in document processing pipelines, capture-oriented systems have also been explored. Courtney (2021) showed that early image quality feedback improves the proportion of usable document images by preventing poor-quality inputs from reaching OCR and verification stages.

2.3.2 Research on Existing Blur Detection Techniques

Research on blur detection has produced a wide range of algorithms. As proposed by Zhang et al. (2022), a fast no-reference blur detection method suitable for real-time deployment, demonstrating strong performance under non-uniform blur conditions.

Older methods, which are handcrafted, are less robust than deep learning methods that have proved to be more reliable though require higher computation knowledge. The findings of Li et al. (2018) and Wang et al. (2018) are consistent with this observation. Similarly, Zhou et al (2023) observed that some deployments focus more on accuracy while neglecting real-time performance, without considering how tedious the deployment is. This position is further reinforced by Gao et al. (2020) and Zhang et al. (2021).

Learning-based approaches have gained attention due to their robustness. According to Koushik et al. (2025), blind image quality assessment models applied to mobile-captured document images outperform classical single-metric approaches, though they require careful dataset curation and are more computationally demanding to deploy.

Comparative evaluations indicate trade-offs between computational efficiency and robustness, with no single method performing optimally under all real-world conditions.

2.3.3 Blur Detection in Document Images

Collectively, Li et al., (2020) in his further research, Zhang et al., (2021) and Hernandez-Ortega et al., (2022) investigated how negatively blur affects text recognition and document analysis. Nonetheless, their studies focus more on assumed general image quality and did not pay particular attention to the concept of 'blur versus no-blur' detection in document instant rejection approach.

Studies focusing particularly on document images emphasize their structural differences from natural images. Liang et al. (2019) showed that blur detection methods designed for natural scenes often fail to predict document usability accurately due to dense text and structured layouts.

Nayef et al. (2018), in his further research, examined poor performance of generic blur detection under low-light document capture conditions. These findings reveal gaps in document-specific blur detection, particularly for identity documents captured in uncontrolled environments.

2.3.4 Review of Existing Automated Document Verification Systems

End-to-end document verification systems increasingly integrate image quality assessment modules. Mohsenzadegan et al. (2021) examined document image classification under real-world capture conditions and identified image blur and distortion as a major source of failure in automated document processing pipelines. Despite Awal et al., (2017) successful analytical approach, there was no evaluation of blur detection.

On the other hand, Attivissimo et al., (2019) contributed to the automatic identity reader, which is very useful to real-world use. Perhaps, his assumption that all image qualities are acceptable limited his study from further investigation on blurry or no-blurry image.

As noted by earlier works, systems lack efficient blur detection mechanisms capable of operating early in the pipeline without introducing excessive computational overhead. Appalaraju et al. (2021), for example, developed an end-to-end deep learning model for document understanding that achieved strong performance across multiple tasks but required substantial computational resources and contained no dedicated mechanism for detecting or filtering blur prior to processing.

Arlazarov et al., (2019) demonstrated a strong and well-structured layout analysis and recognition technique, but there was no automatic document rejection as a result of blur. Similarly, Bulatov et al., (2021) proposed a good end-to-end capture and recognition workflow, but does not evaluate blurry identity document detection. Hernandez-Ortega et al., (2022) focused on document image quality for OCR while blur detection was treated as irrelevant.

This limitation motivates my development of specialized blur detection systems for identity documents. My research will design and implement a system that will be dedicated, lightweight, and automated for blur detection for identity document images, which yields the positive results of improved accuracy and speed.

2.4 Summary of Literature Review

S/N	AUTHOR(S)	CONTRIBUTION / WORK DONE	TECHNIQUE OR METHODOLOGY	LIMITATIONS
1	Li et al., (2018)	Li designed **handcrafted features**, and demonstrated how blur can be detected by looking at the extent image details have been damaged; like sharp edge and high frequency component. This method is lightweight and easy to interpret. Once there is loss of high-frequency contents, it means there is an image blur.	Handcrafted Image Features	Sensitive and reacts to noise and lighting.
2	Wang et al., (2018)	Wang designed **edge and gradient-based blur estimation techniques** used in detecting blur quality, that confirms an image sharpness.	Edge and Gradient-based Blur Estimation	Poor performance under different lighting degrees.
3	Zhang et al., (2019)	Zhang developed **CNN-based supervised learning**, with improved reliability and accuracy more than the handcrafted-based method.	CNN-based, that automatically studies blur patterns from data.	Gives equal image grading and does not recognise text layout and document structure.
4	Arlazarov et al., (2019)	He developed a **full document recognition pipeline** that involves ML in advancing and improving layout analysis and structured document understanding. This method assumes images are already clear.	Rule-based + ML document recognition pipeline	No pre-processing blur capturing. Treats images as good.
5	Attivissimo et al., (2019)	He developed an **Automated Identity Document Reading System** that applied OCR and rules. This innovation contributed to further applied research in automated identity verification.	Automated Identity Reading System	Blurry image not clearly or directly identified.

6	Gao et al., (2020)	Gao developed a **Deep CNN-based blur classification model,** with the design that images can successfully be classified into blurred or clear. This comes with increased speed and accuracy.	Deep-Learning (CNN) Blur Classifier	Based on trained model; and because of its restricted datasets, it is not suitable for real-world document processing.
7	Li et al., (2020)	Although in 2018 Li implemented handcrafted features to detect image blur, in 2020 he advanced his assessment to a **deep-learning-based Image Quality Assessment (IQA) method**, which not only classifies blur/no-blur, but also scores image quality.	Deep Learning-based (IQA)	Does not clearly state whether an image is blurry or not; it rather gives a general acceptance rate.
8	Appalaraju et al., (2021)	Appalaraju introduced an **End-to-End Learning Model for Document Understanding.** This is a single model that handles multiple document understanding tasks.	End-to-end Deep Learning Document Understanding	No expression or indication of image quality check. High-tech computing processing.
9	Bulatov et al., (2021)	Bulatov proposed an **End-to-End Mobile Document Capture and Recognition System**, that incorporates capture, preprocessing and recognition into a single workflow.	Integrated Mobile Document Processing	Blurry image is not automatically rejected before processing.
10	Zhang et al., (2021)	Zhang introduced a **comprehensive method of evaluating problems associated with image quality**. This method incorporates blur checking, noise and distortion, which he noted affect image quality.	Image Quality Metrics Evaluation	Blur not singled out as an identified problem.

11	Hernandez-Ortega et al., (2022)	Hernandez proposed the use of **document-specific image quality assessment method** in improving the OCR performance.	Document Image Quality Assessment for OCR	Blur not addressed as a major issue.
12	Zhou et al., (2023)	Zhou introduced the **Advanced Deep CNN architecture**, which is very powerful in function performance and also of higher accuracy in the detection of image blur.	Advanced Deep (CNN) Architectures	Expensive computation method that may not be affordable or applicable to real-world capture.

Table 2.1: Summary of Related Literature on Blur Detection and Identity Document Image Quality Assessment

2.5 Research Gap

As highlighted in the Literature Review Table above, earlier works have shown significant advancement on blur detection, image quality assessment and document recognition. Nonetheless, they did not explore the areas of automatic and special document blur detection for early rejection of poor quality document images. Besides, the studies treated each of these pieces as separate topics on different papers.

The present study addresses this gap by **designing and implementing a dedicated system for automated blur detection in identity document images**, with the explicit purpose of preventing unusable uploads from entering the verification pipeline. It incorporates the three sub-topics, blur detection, image quality assessment and document recognition into one research work. My design has early operation performance in the verification pipeline, balances accuracy with efficiency, and reflects real-world capture conditions.

The system analysis, design methodology, and technical architecture of the proposed solution are presented in Chapter Three.

CHAPTER THREE: SYSTEM ANALYSIS AND DESIGN

3.1 Introduction

This chapter presents a comprehensive analysis and design of the automated blurry identity document image detection system proposed in this study. The chapter begins with a system analysis that examines how identity document image quality is currently handled in digital verification workflows, identifies the shortcomings of the existing approach, and explains the improvements offered by the proposed system. It then defines the system requirements, both functional and non-functional, that must be satisfied for the proposed system to operate correctly and effectively. Following this, the design methodology adopted for the study is explained and justified, along with a description of the data collection methods employed.

The latter part of the chapter presents the system models developed using Unified Modelling Language (UML) notation, including use case, activity, and class diagrams, which describe system behaviour and structure. Finally, the chapter presents the detailed system design covering input and output interfaces, database design, and the overall system architecture. Together, these elements translate the research objectives defined in Chapter One into a concrete technical blueprint for implementation.

3.2 System Analysis

Identity verification systems in organisations such as banks, financial technology companies, telecommunications providers, and government registration platforms routinely require users to upload images of identity documents as part of digital onboarding and Know Your Customer (KYC) processes. Every automated operation that follows character recognition, biometric face comparison, and document authenticity analysis depends directly on the visual quality of the submitted image. When an uploaded image is blurry, these automated processes are likely to fail, resulting in verification errors, repeated submissions, and increased manual review workload.

This section examines the current approach to managing document image quality in such systems, identifies the specific gaps and weaknesses present in that approach, and evaluates the need for an automated blur detection mechanism. The analysis is grounded in the problem statement and research objectives defined in chapter one and draws on the review of related literature presented in chapter two. The outcome of this analysis provides the technical and operational justification for the design of the proposed automated blurry identity document image detection system.

3.2.1 Analysis of the Existing System

a. Description of the Existing Process

In most current identity verification deployments, document image quality is not evaluated in a dedicated or automated manner at the point of upload. Users are typically presented with an upload interface that accepts an image file as long as it satisfies basic technical constraints such as correct file format (JPEG or PNG), minimum resolution, and maximum file size. No assessment is made of whether the image is visually usable for verification. Once an image passes these basic checks, it is submitted directly to the verification pipeline where OCR, field extraction, and biometric matching are attempted.

Blurry images are discovered only at a later stage, either when OCR fails to extract readable text, when face matching returns low confidence scores, or when a human reviewer manually inspects the uploaded document. At that point, the system may or may not notify the user to re-submit, but valuable processing time and computational resources have already been spent on an unusable image. In some implementations, quality assessment is embedded implicitly within the OCR or verification module itself, meaning that blur is not identified as a distinct rejection reason and the user receives a generic error that offers little actionable guidance.

b. Users of the Existing System

The following actors interact with the existing system:

- End Users (Customers): Individuals who capture or upload identity document images as part of a digital onboarding or registration process.

- Verification Officers / Bank Staff: Personnel responsible for reviewing flagged or failed verification submissions, including those caused by poor image quality.
- Automated Verification Systems: Software components such as OCR engines and biometric matching modules that process uploaded images downstream and are directly affected by image blur.

c. How the Existing System Operates

The existing process follows a sequential workflow in which the user first captures or selects an image of an identity document using a smartphone camera or file upload feature. The image is then submitted to the system, which performs basic technical validation such as checking file type and size. Upon passing these checks, the image is forwarded to the verification pipeline for OCR processing, field extraction, and biometric analysis. If the image is blurry, processing may fail at one or more of these stages. The system then either returns a failure notification to the user or escalates the submission to a manual reviewer or fail silently. The user is eventually prompted to re-capture and re-upload the image, at which point the entire cycle repeats.

d. Workflow of the Existing System

1. Start
2. User Captures / Selects ID Document Image
3. Upload Image to System
4. Basic Technical Validation (File type, File size, Resolution)
5. Validation Passed?

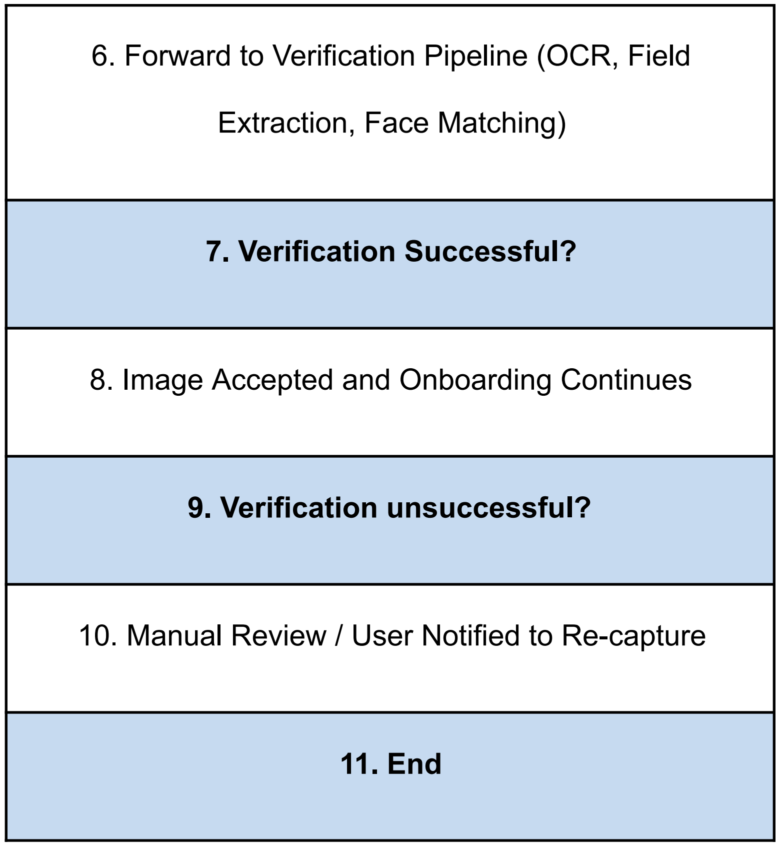

Figure 3.1: Workflow of the Existing Identity Document Upload and Verification System

Figure 3.1 above illustrates the operational workflow of the existing identity document upload and verification process.

3.2.2 Limitations of the Existing System

Based on the analysis presented in Section 3.1.1, the following limitations have been identified in the existing identity document image processing approach:

1) Absence of Automated Blur Detection: The existing system does not include any dedicated mechanism for detecting whether an uploaded image is blurry. Images that fail visual quality standards are not identified at the point of upload, allowing unusable submissions to enter the verification pipeline.

2) Delayed Rejection of Poor-Quality Images: Blur-related failures are only discovered after the image has been processed by downstream components such as OCR or biometric matching engines. This late-stage rejection increases latency and consumes computational resources unnecessarily.

3) Manual and Time-Consuming Review Process: In the absence of automated quality screening, human verification officers must manually inspect failed submissions to determine whether poor image quality was the cause. This reliance on human review reduces operational efficiency and increases processing time.

   a) Reduced Accuracy of Verification Systems: Blurry document images directly impair the performance of OCR engines and biometric face matching systems by weakening edge sharpness and destroying fine character details. This leads to higher error rates and lower verification success rates.

4) Poor and Frustrating User Experience: Users are not informed of image quality problems at the time of upload. They receive failure notifications only after the verification process has completed, leading to repeated re-capture and re-submission cycles that cause frustration and erode trust in the system.

4) Lack of a Quality Control Module: Existing systems assume that uploaded images are of sufficient quality for automated processing. There is no dedicated quality control layer to enforce image standards before verification begins.

These limitations collectively demonstrate the need for an automated, fast, and reliable blur detection system that can identify poor-quality document images at the point of upload and prevent them from entering the verification pipeline. The proposed system, described in the following section, is designed to directly address each of these shortcomings.

3.2.3 Analysis of the Proposed System

The proposed system is an Automated Blurry Identity Document Image Detection System designed to assess the visual quality of identity document images at the point of upload, before any downstream verification processing begins. The system computes a quantitative blur score for each uploaded image using a hybrid classical image processing approach that combines Laplacian variance and Tenengrad gradient magnitude metrics. Based on the computed score, the image is assigned a quality grade and compared against a configurable threshold to determine whether it should be accepted or rejected.

The system is designed to operate as a lightweight, explainable, and real-time quality gate integrated into the document upload workflow. It exposes a backend application programming interface (API) built with FastAPI, stores all image metadata, blur scores, grades, and decisions in a PostgreSQL database for auditing, and provides two interfaces: a user-facing upload interface and an administrator interface for threshold configuration and submission review.

The specific capabilities of the proposed system, and how each addresses a limitation of the existing system, are described below:

1. Automatic Blur Detection at Upload: The system automatically computes a blur score for every uploaded image using Laplacian variance and Tenengrad metrics, eliminating the absence of automated blur detection identified in the existing system.
2. Early Rejection of Poor-Quality Images: Blurry images are detected and rejected before entering the verification pipeline. Users receive immediate, actionable feedback prompting them to re-capture a clearer image, directly addressing the problem of delayed rejection.
3. Reduced Manual Review Workload: By automatically screening image quality at upload time, the system eliminates the need for manual inspection of blur-related failures, improving operational efficiency.

4. Improved Verification Accuracy: By ensuring that only clear, usable images proceed to OCR and biometric processing, the system reduces verification error rates and improves the accuracy of downstream identity checks.
5. Enhanced User Experience: Users are informed of image quality problems immediately at upload, with a specific blur-related rejection message and a prompt to re-capture. This eliminates the frustration of late-stage failure notifications.
6. Configurable Quality Control Layer: The administrator interface allows the blur score threshold to be configured and adjusted, providing a flexible and auditable quality control mechanism that the existing system lacks.

3.3 System Requirements Specification (SRS)

System Requirements Specification (SRS) defines the capabilities and constraints that the proposed system must satisfy in order to function correctly, efficiently, and securely. It provides a formal reference that guides implementation and evaluation. The requirements are organised into two categories: functional requirements, which describe what the system must do, and non-functional requirements, which describe how well the system must perform those functions.

3.3.1 Functional Requirements

The following functional requirements define the core operations that the proposed system must be capable of performing:

- Users are able to submit identity document images in JPEG or PNG format through the web-based upload interface.
- The system shall automatically analyse each uploaded image for blur upon submission, without requiring manual initiation.
- The system shall compute a numeric blur score for each image using a hybrid approach combining Laplacian variance and Tenengrad gradient magnitude.

- The system shall assign a quality grade to each image based on its blur score, using the categories: Good, Mild, Moderate, and Severe.
- The system shall compare the computed blur score against an administrator-defined threshold to determine whether the image is acceptable.
- The system shall reject images whose blur score falls below the defined threshold and display a clear, actionable rejection message to the user.
- The system shall prompt rejected users to re-upload a clearer image.
- The system shall accept images that meet the threshold requirement and allow them to proceed without interruption.
- The system shall store all upload records, including image metadata, blur scores, quality grades, acceptance or rejection decisions, and timestamps, in a PostgreSQL database.
- The system shall provide an administrator interface for reviewing all submitted images and their associated records.
- The system shall allow the administrator to configure the acceptable blur score threshold through the admin interface.

3.3.2 Non-Functional Requirements

The following non-functional requirements define the performance, usability, reliability, security, and scalability standards the system must meet:

a. Performance Requirements

- The system shall compute a blur score and return a quality decision within two seconds of image submission under normal operating conditions.
- The system shall be capable of processing uploaded images without blocking the user interface during evaluation.

b. Usability Requirements

- The user interface shall present upload controls and quality feedback messages in a clear and intuitive manner that does not require technical knowledge.
- Rejection messages shall clearly state that the image is blurry and provide specific guidance on how to capture a better image.

c. Reliability Requirements

- The system shall produce consistent blur scores for the same image across repeated submissions.
- The blur detection module shall operate correctly across a range of image capture conditions, including variations in lighting and document type.

d. Security Requirements

- All uploaded identity document images and associated records shall be stored securely with restricted access.
- Access to the administrator interface shall require authentication.

e. Scalability Requirements

- The system shall be architected to handle growth in upload volume without meaningful deterioration in response latency.

3.4 System Design Methodology

Object-Oriented Analysis and Design (OOAD) was adopted as the design methodology for the proposed system. OOAD approaches system development by representing the problem domain as a set of interacting objects, each responsible for a specific aspect of the system's behaviour and data. The methodology proceeds in two phases: an analysis phase, in which the problem space and its requirements are examined and expressed through object-oriented models, and a design phase, in which those models are refined into a concrete implementation blueprint.

In the context of this project, OOAD was used to identify the key entities involved in the blur detection system, including uploaded images, users, blur scores, quality grades, and detection results, and to define the relationships and interactions between these entities. Throughout both phases, use case, activity, and class diagrams were produced to capture system behaviour and structure in a notation that is both precise and accessible to technical and non-technical readers alike.

3.4.1 Justification for the Methodology

Object-Oriented Analysis and Design was selected as the design methodology for this project for the following reasons:

1. Modularity: OOAD encourages the decomposition of a system into well-defined, self-contained components. This is particularly suitable for the proposed system, which comprises distinct modules including an image upload interface, a blur detection engine, a grading module, a database layer, and an administrator dashboard.

2. Maintainability and Extensibility: Object-oriented design promotes code reuse and makes it straightforward to extend the system with additional quality checks, such as glare detection or low-illumination analysis, in future iterations.

3. Clear Representation Using UML: OOAD is naturally expressed through UML diagrams, which provide a formal and visual means of communicating system design to both technical and non-technical stakeholders.

4. Alignment with the Implementation Stack: The backend of the proposed system is implemented using Python and FastAPI, both of which support object-oriented programming paradigms natively. OOAD therefore aligns directly with the structure and conventions of the chosen implementation technologies.

3.4.2 Method of Data Collection

The following data collection methods were employed in the analysis and design of the proposed system:

Review of Publicly Available Identity Document Datasets: Existing benchmark datasets of identity document images, including those referenced in the literature review, were examined to understand the range of blur conditions and capture scenarios that the system must handle. This informed decisions regarding blur metric selection and grading thresholds.

1. Review of Literature on Blur Detection Techniques: A systematic review of published research on classical and learning-based blur detection methods, as presented in chapter two, provided the technical foundation for selecting the Laplacian variance and Tenengrad gradient magnitude as the primary detection metrics.

2. Observation of Existing KYC and Document Verification Workflows: Publicly documented descriptions and technical overviews of commercial KYC platforms and digital onboarding systems were reviewed to understand the typical structure of existing document upload pipelines and to identify where a quality gate could be most effectively integrated.

3. Sample Identity Document Images for Prototype Testing: Representative sample images exhibiting varying levels of blur were used during the design phase to validate the proposed grading logic and to inform the selection of score ranges for each quality category.

3.5 System Modelling Using UML

UML provides a common visual language for describing software systems from different angles. Three diagram types are used here to model the proposed system: a use case diagram that captures actor interactions, an activity diagram that traces the operational flow of the upload and detection process, and a class diagram that defines the structural relationships between system components.

3.5.1 Use Case Diagram

Figure 3.2 maps the functional scope of the proposed system by showing which actors interact with it and how. Two actors are involved: the User, who submits identity document images for verification, and the Administrator, who manages system configuration and monitors submission records.

The User interacts with three use cases: Upload Image, View Blur Detection Result, and Re-upload Image (triggered following a rejection). The Administrator interacts with two use cases: Configure Blur Threshold, which allows the acceptable blur score boundary to be adjusted, and Review Upload Records, which enables audit and monitoring of all processed submissions. The blur detection operation itself is represented as an internal use case, Detect Blur and Assign Grade, which is automatically triggered upon image upload and is not directly initiated by either actor.

Actor	Use Case	Description
	Upload Image	User submits an identity document image through the upload interface.
User	View Blur Detection Result	User receives a quality decision (accepted or rejected) with a blur score and grade.
User	Re-upload Image	User re-submits a clearer image following a rejection.
System (Internal)	Detect Blur and Assign Grade	System computes blur score using Laplacian and Tenengrad metrics and assigns a quality grade.

Administrator	Configure Blur Threshold	Admin sets or updates the minimum acceptable blur score.
Administrator	Review Upload Records	Admin views all uploaded images with scores, grades, decisions, and timestamps.

Figure 3.2: Use Case Summary of the Proposed Automated Blurry Identity Document Image Detection System

3.5.2 Activity Diagram

Figure 3.3 traces the end-to-end sequence of operations triggered by an image upload, from the moment the user submits a file through to the point where an acceptance or rejection decision is returned. Both the successful and rejected paths are shown.

The process begins when the user accesses the upload interface and selects an identity document image. The system first performs basic technical validation to confirm that the image meets file format and size requirements. If validation fails, the user is notified and prompted to select a valid file. If validation passes, the system proceeds to blur analysis: the uploaded image is processed to compute its Laplacian variance and Tenengrad gradient magnitude scores, which are combined into a single composite blur score. The system then assigns a quality grade (Good, Mild, Moderate, or Severe) based on predefined score ranges and compares the score against the administrator-configured threshold. If the score meets or exceeds the threshold, the image is accepted and the submission proceeds. If the score falls below the threshold, the image is rejected, a descriptive error message is returned to the user, and the user is prompted to re-capture and re-upload a clearer image. In both cases, the image record, blur score, grade, decision, and timestamp are persisted to the PostgreSQL database

Step		Action	Actor / Component
1		User accesses upload interface and selects an image	User
2		System validates file format and size	System
3		Validation failed → User notified, prompted to re-select	System → User
4		Validation passed → Image forwarded to blur detection module	System
5		Compute Laplacian Variance score	Blur Detection Engine
6		Compute Tenengrad Gradient Magnitude score	Blur Detection Engine
7		Combine scores into composite blur score	Blur Detection Engine
8		Assign quality grade (Good / Mild / Moderate / Severe)	Grading Module
9		Compare blur score against admin-configured	Decision Module

		threshold	

10	Score meets threshold → Image accepted, proceeds to verification	System → User
11	Score below threshold → Image rejected, user prompted to re-upload	System → User
12	Store image record, score, grade, decision, and timestamp in database	Database Manager

Figure 3.3: Activity Diagram of the Proposed Automated Blurry Identity Document Image Detection System

Figure 3.4: Class Diagram of the Proposed Automated Blurry Identity Document Image Detection System

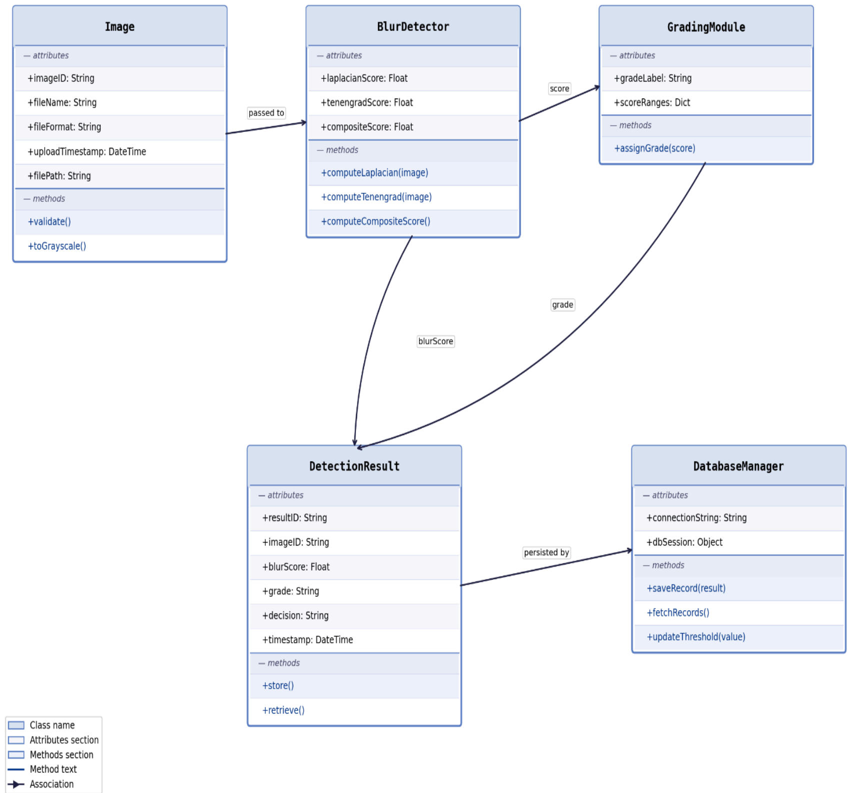

Figure 3.4: Class Diagram of the Proposed Automated Blurry Identity Document Image Detection System

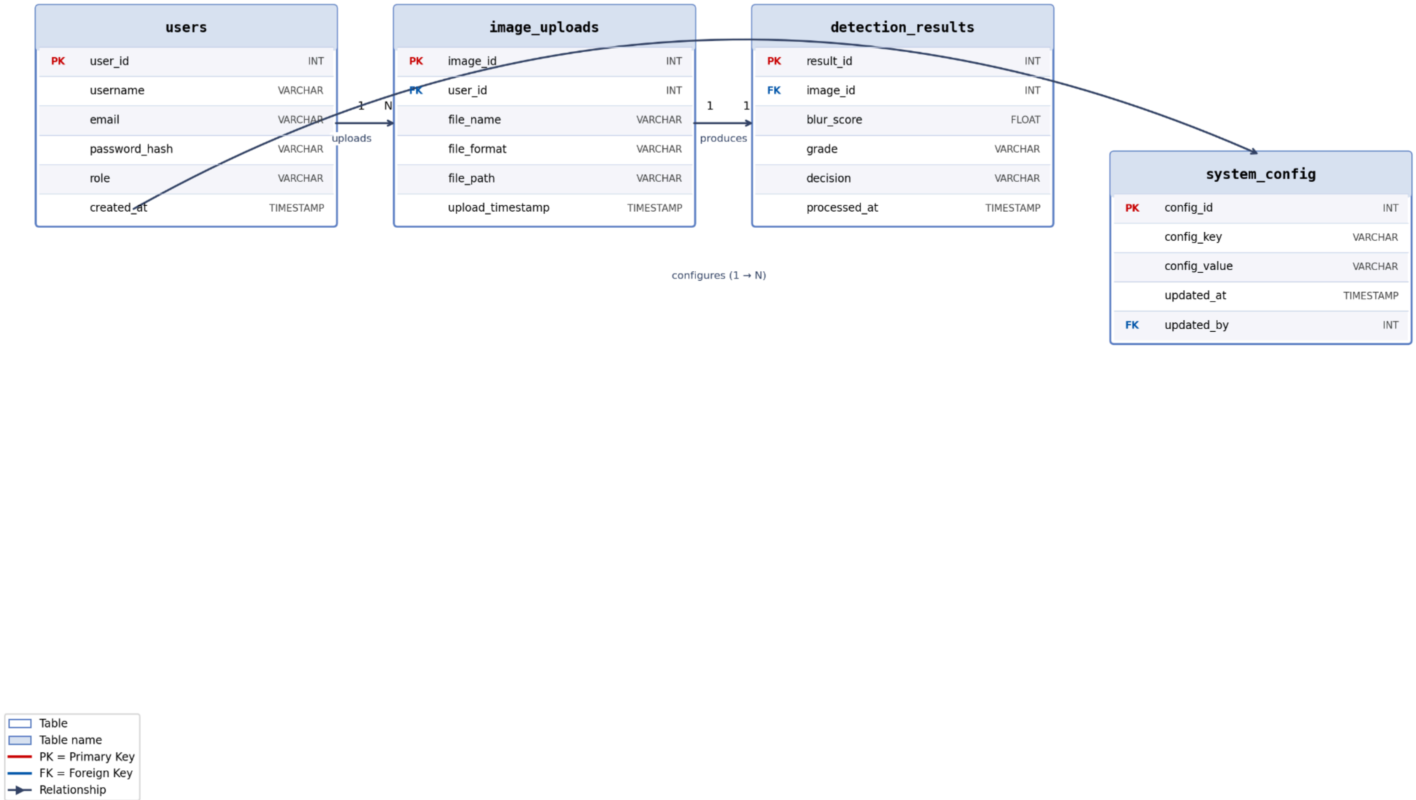

Figure 3.4.5: Database Schema of the Proposed Automated Blurry Identity Document Image Detection System

3.5.3 Class Diagram

The class diagram in Figure 3.4 describes the structural design of the proposed system by identifying the main classes, their attributes, their methods, and the relationships between them. Five primary classes are defined:

1. Image: Represents an uploaded identity document image. Attributes include imageID, fileName, fileFormat, uploadTimestamp, and filePath. Methods include validate() to check file constraints and to Grayscale() to prepare the image for processing.

2. Blur Detector: Encapsulates the blur detection logic. Attributes include laplacianScore, tenengradScore, and compositeScore. Methods include compute Laplacian(image), computeTenengrad(image), and compute CompositeScore() to calculate the final blur measurement.

3. Grading Module: Handles quality grade assignment. Attributes include gradeLabel and

scoreRanges. Methods include assign Grade(score) which maps the composite score to one of four grade categories. DetectionResult: Stores the outcome of a single blur detection operation. Attributes include resultID, imageID (foreign key), blurScore, grade, decision, and timestamp. Methods include store() and retrieve().

4. Database Manager: Manages persistence of all records to the PostgreSQL database. Attributes include connectionString and dbSession. Methods include saveRecord(result), fetchRecords(), and updateThreshold(value).

Class	Key Attributes	Key Methods
Image	imageID, fileName, fileFormat, uploadTimestamp, filePath	validate(), toGrayscale()
BlurDetector	laplacianScore, tenengradScore, compositeScore	computeLaplacian(image), computeTenengrad(image), computeCompositeScore()
GradingModule	gradeLabel, scoreRanges	assignGrade(score)
DetectionResult	resultID, imageID, blurScore, grade, decision, timestamp	store(), retrieve()
DatabaseManager	connectionString, dbSession	saveRecord(result), fetchRecords(), updateThreshold(value)

Figure 3.4: Class Diagram Summary of the Proposed Automated Blurry Identity Document Image Detection System

3.6 System Design

This section translates the analysis and models presented in the preceding sections into a detailed technical design for the proposed system. The design addresses the input interface, output interface, database structure, and overall system architecture.

3.6.1 Input Design

The primary input to the system is an identity document image submitted by the user through the web-based upload interface, which is implemented using React. The upload interface presents the user with a clearly labelled file selection control and a submission button. The interface enforces the following input validation rules before the image is forwarded to the backend:

- Accepted file formats: JPEG (.jpg, .jpeg) and PNG (.png).
- Maximum file size: 10 megabytes per upload.
- Minimum resolution: Images must meet a minimum pixel dimension to ensure sufficient detail for blur analysis.

If any of these constraints are not satisfied, the interface displays an inline validation message and prevents submission. Valid images are transmitted to the FastAPI backend via an HTTP POST request as multipart form data. The administrator input, specifically the blur score threshold, is provided through a dedicated configuration panel in the admin interface, where the threshold value can be entered and saved.

3.6.2 Output Design

The system produces two categories of output: user-facing quality feedback and administrator-facing audit records.

For the user, the output following an image submission consists of a quality decision panel displayed on the upload interface. This panel includes the computed numeric blur score, the assigned quality grade (Good, Mild, Moderate, or Severe), and a status message indicating whether the image has been accepted or rejected. If the image is rejected, the status message explicitly states that the image is too blurry and provides clear guidance on how to recapture a sharper image, for example by ensuring adequate lighting, holding the device steady, and ensuring the document is in focus.

For the administrator, the output is presented as a tabular dashboard that lists all uploaded image records. Each row in the table displays the image filename, upload timestamp, file type, computed blur score, quality grade, and final decision. The administrator can also view the currently configured blur threshold and update it through the settings panel.

3.6.3 Database Design

The proposed system uses PostgreSQL as its relational database management system. The database stores all image upload records, blur detection results, and system configuration settings. Three primary tables are defined:

Table	Primary Key	Key Columns	Description
users	user_id (INT)	username, email, password_hash, role, created_at	Stores registered user and administrator account information.
image_uploads	image_id (INT)	user_id (FK), file_name, file_format, file_path, upload_timestamp	Records metadata for every identity document image submitted to the system.
detection_results	result_id (INT)	image_id (FK), blur_score, grade, decision, processed_at	Stores the blur score, quality grade, and acceptance or rejection decision for each processed image.

system_config	config_id (INT)	config_key, config_value, updated_at, updated_by	Stores system configuration values including the administrator-defined blur threshold.

Table 3.1: Database Table Design for the Proposed System

The image_uploads table references the users table through a foreign key on user_id. The detection_results table references the image_uploads table through a foreign key on image_id. The system_config table stores the blur threshold as a key-value record that the backend reads each time a quality decision is made. All tables include timestamp columns to support auditing and review.

3.6.4 System Architecture Design

The proposed system is designed as a three-tier web application comprising a client tier, an application tier, and a data tier. Figure 3.5 presents a summary of the system architecture and the communication flow between components.

Tier	Component	Technology	Responsibility
Client Tier	User Upload Interface	React (JavaScript)	Presents the image upload form, enforces client-side validation, displays blur detection results and feedback messages.
Client Tier	Administrator Dashboard	React (JavaScript)	Displays all upload records, allows threshold configuration, and supports audit

			review of submissions.
Application Tier	FastAPI Backend	Python / FastAPI	Receives uploaded images, applies validation, invokes blur detection logic, applies grading, compares against threshold, and returns results.
Application Tier	Blur Detection Engine	Python / OpenCV	Computes Laplacian variance and Tenengrad gradient magnitude; combines scores into a composite blur measurement.
Application Tier	Grading Module	Python	Maps the composite blur score to a quality grade (Good, Mild, Moderate, Severe) using predefined score ranges.

Data Tier	PostgreSQL Database	PostgreSQL	Persists all image records, detection results, user accounts, and system configuration including the blur threshold.

Figure 3.5: System Architecture of the Proposed Automated Blurry Identity Document Image Detection System

The communication flow operates as follows. When a user submits an image, the React interface sends it to the FastAPI backend as a multipart POST request. The backend validates the file, forwards it to the Blur Detection Engine for metric computation, and passes the resulting score to the Grading Module for classification. The Decision Module then compares the grade against the stored threshold to produce an acceptance or rejection outcome. The full result score, grade, decision, and timestamp is written to PostgreSQL and returned to the frontend as a JSON payload, which the React interface renders for the user. Administrators reach a separate view via the same client, where all historical upload records are presented in a searchable table.

Taken together, the work presented in this chapter moves the research from problem statement to technical blueprint. The shortcomings of current document upload workflows were examined and used to build the case for a dedicated quality gate. Requirements were specified across both functional and non-functional dimensions to constrain and guide what was built. The OOAD methodology was applied and justified, and UML diagrams were produced to express system behaviour and structure. The resulting design covering interfaces, database schema, and a three-tier architecture forms the basis for the implementation described in Chapter Four.

CHAPTER FOUR: SYSTEM IMPLEMENTATION

4.1 Introduction

This chapter presents the implementation of the proposed Automated Blurry Identity Document Image Detection System. It describes the development environment, programming languages, libraries, and frameworks employed in building the system, and explains the rationale behind each tool selection. The chapter further details the image dataset used for testing, the preprocessing steps applied to prepare images for blur analysis, and the complete implementation of the blur detection pipeline, including the Laplacian variance and Tenengrad gradient magnitude computation, composite scoring, quality grading, and threshold-based decision logic.

Additionally, the chapter describes the system modules and explains how they interact to perform blur detection and quality gating in real time. The backend API, database integration, React-based user interface, and administrator dashboard are each described in detail. Finally, the chapter presents the testing procedures and evaluation metrics used to assess system performance, followed by a presentation and discussion of the results obtained from the implemented system.***(see the Appendix A)***

4.2 Development Environment and Tools

The proposed system was developed using a combination of programming tools, frameworks, and libraries carefully selected for their suitability to image processing, web API development, and relational database management. The development environment was configured to support rapid prototyping, clean separation of concerns between system components, and efficient real-time processing of uploaded identity document images.

4.2.1 Programming Languages and Libraries

Two programming languages were used: Python on the backend and JavaScript, via the React library, on the frontend. Python was the natural choice for backend work given its rich ecosystem of image processing libraries, its readable syntax, and its established position in computer vision and data engineering. It integrates directly with OpenCV, the library at the heart of the blur metric computation.

On the frontend, React was chosen for its component-based architecture, which makes it straightforward to build interfaces that update in place as detection results arrive, without triggering a full page reload.

The following libraries and frameworks were used during development:

Library / Framework	Language	Purpose
FastAPI	Python	Serves as the backend web framework through which the image upload and blur detection API endpoints are exposed.
OpenCV (cv2)	Python	Core computer vision library used to load images, convert colour spaces, compute Laplacian variance, and apply Sobel gradient operators for Tenengrad scoring.
NumPy	Python	Used for numerical array operations, including computation of variance and gradient magnitude statistics.
Pillow (PIL)	Python	Used for image format validation and initial image loading before processing.

SQLAlchemy	Python	Object-relational mapper (ORM) used to define database models and manage PostgreSQL interactions from the FastAPI backend.
Psycopg2	Python	PostgreSQL database adapter providing the underlying database connection for SQLAlchemy.
Uvicorn	Python	ASGI server used to serve the FastAPI application in development and production.
React	JavaScript	Frontend library used to build the user-facing upload interface and the administrator dashboard.
Axios	JavaScript	HTTP client used in the React frontend to send image upload requests to the FastAPI API and receive detection results.

Table 4.1: Libraries and Frameworks Used in System Development

4.2.2 Development Platform and Hardware Requirements

The system was developed using tools and platforms that support Python backend development, React frontend development, and PostgreSQL database administration. The following development environments were used:

- **Visual Studio Code (VS Code):** Served as the primary editor for both the Python backend and React frontend code. Its built-in terminal, low resource footprint, and broad extension support for Python and JavaScript made it a practical choice throughout the project.

- **pgAdmin 4:** Used as the graphical administration tool for PostgreSQL database management, enabling inspection of tables, records, and schema relationships during development.

- **Postman:** Used to test and validate the FastAPI REST API endpoints independently of the frontend, enabling precise verification of request and response structures.

- **Git:** Used for version control and tracking changes across the codebase throughout the development process.

The hardware configuration used for system development and testing is as follows:

Component	Specification
Processor	Intel Core i5 (8th Generation or equivalent)
RAM	8 GB DDR4
Storage	256 GB SSD
Operating System	Windows 10 / Ubuntu 22.04 LTS
GPU	Not required (classical CV metrics; CPU-only computation)
Network	Standard internet connection for dependency installation

Table 4.2: Hardware Configuration Used for System Development

4.3 Dataset Description and Preprocessing

4.3.1 Dataset Description

The dataset used in this study consists of synthetically generated identity document images representing four levels of blur severity. Since collecting real identity document images raises significant data protection concerns under the Nigerian Data Protection Regulation (NDPR), and since synthetic blur generation at defined kernel parameters provides mathematically exact and fully reproducible ground truth labels, a controlled synthetic dataset was constructed for this evaluation. This approach is consistent with established practice in document image quality research, where Gaussian blur at known standard deviation values is used as a controlled proxy for real-world defocus and motion blur.

A set of 25 clear synthetic identity document source images was generated programmatically using a reproducible random seed (seed=42). Each source image was designed to replicate the visual structure of an identity document, including text fields, layout regions, and portrait areas. From each clear source image, four variants were produced by applying Gaussian blur at controlled kernel parameters, yielding a total evaluation dataset of 100 images across four quality categories. The blur parameters applied to each category are as follows: the Good category comprises the original unblurred source images; the Mild category was produced using a Gaussian kernel of size (5×5) with standard deviation σ=1.5; the Moderate category was produced using a kernel of size (15×15) with σ=4.0; and the Severe category was produced using a kernel of size (31×31) with σ=10.0. The resulting dataset composition is presented in Table 4.3.

Quality Grade	**Blur Parameters**	Blur Description	Number of Images	Proportion
Good	No blur (original source image)	Sharp, fully readable document image	25	30%
Mild	Gaussian kernel (5×5), σ = 1.5	Slight softening; text remains largely readable	25	25%

Moderate	Gaussian kernel (15×15), σ = 4.0	Noticeable blur; some text fields degraded	25	25%
Severe	Gaussian kernel (31×31), σ = 10.0	Heavy blur; document text unreadable	25	20%

Table 4.3: Dataset Composition by Quality Grade

4.3.2 Image Preprocessing

Before blur analysis is performed, each uploaded image passes through a standardised preprocessing pipeline designed to ensure consistent and reliable metric computation regardless of the original image dimensions, colour mode, or file format. The preprocessing steps are as follows:

1. Format Validation: Before preprocessing, the image is validated to confirm it is in an accepted format (JPEG or PNG). Corrupt or unsupported files are rejected at this stage with an appropriate error response.

2. Colour Space Conversion: Each image is converted from its original colour representation (RGB or RGBA) to grayscale. Blur detection metrics based on Laplacian variance and Tenengrad gradient magnitude operate on single-channel intensity images. Converting to grayscale eliminates the influence of colour variation and ensures that the computed metrics reflect structural sharpness rather than colour information.

3. Aspect-Ratio-Preserving Resize: Images are resized to fit within a 640 x 480 pixel boundary while preserving the original aspect ratio. This approach ensures that blur scores are not influenced by the absolute pixel count of an image while preventing geometric distortion. Earlier testing revealed that forcing all images to exactly 640 x 480 pixels caused portrait-oriented identity documents to be

distorted into landscape dimensions, artificially suppressing edge measurements and producing incorrect low blur scores. Preserving aspect ratio eliminates this source of error, as confirmed by processing log output showing images processed at their correct proportional dimensions, for example 638 x 480 for landscape captures and 360 x 480 for portrait captures.

4. Array Conversion: The preprocessed grayscale image is converted to a NumPy array to enable efficient numerical computation of the blur metrics using OpenCV and NumPy operations.

4.4 Blur Detection Implementation

The blur detection engine is the core component of the proposed system. It implements a hybrid classical image processing approach that combines two established sharpness metrics: Laplacian variance and Tenengrad gradient magnitude. The use of two complementary metrics improves robustness under non-uniform blur conditions, which are common in smartphone-captured identity document images. The implementation is entirely deterministic, requiring no model training, no labelled data, and no GPU computation, making it lightweight and suitable for real-time upload validation.

4.4.1 Laplacian Variance

The Laplacian is a second-order derivative filter that measures how rapidly pixel intensity changes across the image. Sharp images produce concentrated, high-amplitude responses at edges, yielding high variance in the filtered output. Where blur is present, those edges become gradual and diffuse, and the variance of the Laplacian response drops correspondingly.

In the proposed system, the Laplacian variance is computed by applying the cv2.Laplacian function to the grayscale image using a kernel depth of CV_64F, which preserves the full precision of the second-order derivative. NumPy is then used to compute the variance of the filtered output. A high variance value signals a sharp image; a low value is indicative of blur.

The implementation is shown below:

```
gray = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
laplacian = cv2.Laplacian(gray, cv2.CV_64F)
laplacian_score = laplacian.var()
```

4.4.2 Tenengrad Gradient Magnitude

The Tenengrad metric evaluates image sharpness by measuring the strength of image gradients using Sobel operators applied in both the horizontal and vertical directions. The Sobel operator is a first-order derivative filter that amplifies edges and intensity transitions. In a sharp image, these transitions are strong and well-defined, producing high gradient magnitudes. In a blurry image, transitions are gradual and weak.

In the proposed system, horizontal and vertical Sobel gradients are computed using cv2.Sobel with a kernel size of three. Per-pixel gradient magnitude is then derived as the Euclidean norm of the horizontal and vertical gradient components. The mean of the gradient magnitude map serves as the Tenengrad score. A higher mean gradient magnitude indicates greater sharpness.

The implementation is shown below:

```
sobel_x = cv2.Sobel(gray, cv2.CV_64F, 1, 0, ksize=3)
sobel_y = cv2.Sobel(gray, cv2.CV_64F, 0, 1, ksize=3)
gradient_magnitude = np.sqrt(sobel_x**2 + sobel_y**2)
tenengrad_score = gradient_magnitude.mean()
```

4.4.3 Composite Blur Score

The Laplacian variance and Tenengrad scores are each sensitive to different characteristics of image sharpness. Laplacian variance responds primarily to the concentration of high-frequency detail, while Tenengrad responds to the overall strength of edge gradients. Combining both metrics into a single composite score improves robustness, particularly for images with non-uniform blur where one region may be sharp while another is degraded.

In the proposed system, the composite blur score is computed as a normalised weighted average of the two metrics. Each metric is independently normalised to a 0 to 100 scale using maximum reference values calibrated empirically against real smartphone-captured identity document images during system testing. The calibrated values are MAX_LAPLACIAN = 500.0 and MAX_TENENGRAD = 80.0. These constants were determined by testing the system against real identity document images of known quality and adjusting the normalisation ceiling until scores correctly reflected the visual usability of each image. The two normalised values are then combined using equal weights of 0.5, yielding a composite score in the range of 0 to 100, where higher values indicate greater sharpness.

The composite score formula is as follows:

MAX_LAPLACIAN = 500.0

MAX_TENENGRAD = 80.0

norm_laplacian = min(laplacian_score / MAX_LAPLACIAN, 1.0) * 100

norm_tenengrad = min(tenengrad_score / MAX_TENENGRAD, 1.0) * 100

composite_score = (0.5 * norm_laplacian) + (0.5 * norm_tenengrad)

4.4.4 Quality Grading Logic

Once the composite blur score has been computed, the system assigns a quality grade based on predefined score thresholds. The four grade categories reflect practical levels of document usability in an identity verification context.

Composite Score Range	Quality Grade	Interpretation
75.00 - 100.00	Good	Image is sharp and fully usable for OCR and verification.
50.00 - 74.99	Mild	Slight softening present; image is likely usable with minor degradation.

25.00 - 49.99	Moderate	Noticeable blur; significant risk of OCR failure and verification error.
0.00 - 24.99	Severe	Heavy blur; image is not usable for automated verification.

Table 4.4: Quality Grade Definitions and Score Ranges

The system then compares the composite score against the administrator-configured threshold. If the score meets or exceeds the threshold, the image is accepted and proceeds to the verification pipeline. If the score falls below the threshold, the image is rejected and the user is prompted to re-capture a clearer image. All results, including the score, grade, decision, and timestamp, are persisted to the PostgreSQL database.

4.5 System Implementation and Modules

The implemented system is structured into five primary modules that work together to perform blur detection, quality grading, threshold-based decision making, result persistence, and result presentation. Each module is responsible for a specific function within the system architecture, and together they implement the end-to-end workflow described in the system design presented in Chapter Three.

4.5.1 Image Upload Module

Built into the React frontend, the Image Upload Module handles everything on the user-facing side of the submission process: rendering the upload control, running client-side validation checks, and dispatching the image to the backend. The interface renders a clearly labelled file input control that accepts JPEG and PNG files only. Upon file selection, the module performs immediate client-side validation to confirm that the file extension and MIME type are acceptable and that the file size does not exceed the ten-megabyte limit. Any validation failure surfaces as an inline message and blocks the submission from proceeding.

Images that pass validation are sent to the FastAPI backend at the /api/upload endpoint via an Axios POST request, with the file attached as multipart form data. The module enters a loading state during API processing and renders the returned detection result once a response is received.

4.5.2 Blur Detection Engine Module

The Blur Detection Engine Module is implemented in Python within the FastAPI backend. It receives the uploaded image file from the API handler, applies the preprocessing pipeline described in Section 4.2.2, and computes the Laplacian variance, Tenengrad gradient magnitude, and composite blur score as described in Section 4.3. This module is the computational core of the system and operates entirely using classical image processing techniques without requiring any trained model or external API call.

The module returns a dictionary containing the Laplacian score, Tenengrad score, composite blur score, and assigned quality grade, which is then consumed by the Decision Module.

4.5.3 Decision Module

The Decision Module is implemented within the FastAPI backend and is responsible for comparing the computed composite blur score against the administrator-configured acceptance threshold. The threshold value is retrieved from the system_config table in the PostgreSQL database at the time of each upload request, ensuring that threshold updates by the administrator take effect immediately without requiring a system restart.

If the composite score meets or exceeds the threshold, the module sets the decision to accepted. If the score falls below the threshold, the decision is set to rejected and a descriptive rejection message is generated for the user. The decision, together with the blur score, grade, and rejection message, if applicable, is included in the API response and simultaneously passed to the Database Manager Module for persistence.

4.5.4 Database Manager Module

The Database Manager Module handles all interactions between the FastAPI backend and the PostgreSQL database using SQLAlchemy as the object-relational mapper. It defines four database models corresponding to the tables described in Chapter Three: users, image_uploads, detection_results, and system_config.

Upon completion of each blur detection operation, the module writes a new record to the image_uploads table containing the image filename, format, file path, and upload timestamp. It simultaneously writes a corresponding record to the detection_results table containing the blur score, quality grade, acceptance or rejection decision, and processing timestamp. The module also provides a method for reading and updating the blur threshold value in the system_config table, which is called by the Decision Module and the administrator dashboard respectively.

4.5.5 Result Presentation Module

The Result Presentation Module is implemented in React and is responsible for rendering the blur detection outcome to the user after the API response is received. For accepted images, the module displays a green success panel showing the blur score, quality grade (Good or Mild), and a confirmation that the image has been submitted successfully. For rejected images, the module displays a red rejection panel showing the blur score, quality grade (Moderate or Severe), the rejection reason, and step-by-step guidance on how to recapture a clearer image.

The administrator dashboard, also implemented in React, displays a tabular view of all upload records retrieved from the backend. Each row in the table includes the image filename, upload timestamp, file format, blur score, quality grade, and decision status. The dashboard also includes a settings panel that allows the administrator to view and update the current blur score threshold.

4.6 Testing and Evaluation

Testing covered three levels: individual component correctness, end-to-end system workflow integrity, and user interface behaviour. Each stage targeted a different layer of the system and together they provided confidence that the implemented system behaved as specified.

4.6.1 Evaluation Metrics

Performance was measured using binary classification metrics. Each of the 100 evaluation images carried a ground truth label which are usable (Good or Mild) or unusable (Moderate or Severe), assigned through visual inspection cross-referenced against blur severity benchmarks from the literature. Predicted decisions were compared against these labels to produce the metrics below:

- **Accuracy:** The share of images for which the system's accept/reject decision matched the ground truth label, reflecting overall classification correctness.
- **Precision:** Of all images the system accepted, the fraction that were genuinely usable. A high precision value means the system is not letting blurry images through.
- **Recall (Sensitivity):** Of all genuinely usable images, the fraction the system correctly accepted. A high recall value means clear images are not being turned away unnecessarily.
- **F1-Score:** The harmonic mean of precision and recall, collapsing both into a single figure that penalises imbalance between the two.
- **False Acceptance Rate (FAR):** The share of blurry images the system incorrectly passed through. Keeping this near zero is the primary safety requirement for a quality gate of this kind.
- **Average Processing Time:** The mean time in milliseconds from image submission to the return of a quality decision, measured across all test images to assess real-time suitability.

A confusion matrix was also produced to visualise the distribution of correct and incorrect classification decisions across the four quality grade categories.

4.6.2 System Testing

Each system component was tested in isolation to confirm it behaved correctly before being integrated with the rest of the pipeline. The blur detection functions (computeLaplacian, computeTenengrad, computeCompositeScore) were each tested with images of known sharpness characteristics, including artificially generated sharp and blurred images, to confirm that the computed scores were monotonically consistent with the degree of applied blur. The grading logic was tested against boundary values at each score range to confirm correct grade assignment. The database manager functions were tested against a local PostgreSQL instance to confirm correct record creation, retrieval, and threshold update operations.

Once individual components were confirmed correct, integration testing verified that they operated coherently as an end-to-end pipeline. A series of test images spanning all four quality grade categories were submitted through the API endpoint, and the end-to-end behaviour of the system was verified: image receipt, preprocessing, blur metric computation, grading, threshold comparison, database persistence, and API response content. Edge cases tested included zero-byte file submissions, files in unsupported formats, images with extreme aspect ratios, and images with non-uniform blur affecting only a portion of the document.

During integration testing, two implementation issues were identified and resolved through iterative testing with real identity document images. First, the initial preprocessing step forced all images to exactly 640 x 480 pixels, causing portrait-oriented identity documents to be distorted into landscape dimensions. This geometric distortion reduced edge density in the resized image, artificially suppressing both Laplacian variance and Tenengrad gradient scores and causing visually usable images to be incorrectly rejected. The preprocessing step was corrected to perform aspect-ratio-preserving resizing within the 640 x 480 boundary, which eliminated the distortion. Second, the initial normalisation constants (MAX_LAPLACIAN = 2000.0, MAX_TENENGRAD = 120.0) were found to be miscalibrated for smartphone-captured identity document images, causing the composite score scale to be compressed into a low range where real-world usable images scored below the rejection threshold. These constants were recalibrated to MAX_LAPLACIAN = 500.0 and MAX_TENENGRAD = 80.0 through empirical testing, after which the system correctly differentiated between usable and blurry identity document images. These findings are documented in the system processing log, which records each step of the pipeline for every upload and was used as the primary diagnostic tool during testing.

Following the resolution of both implementation issues, the system was evaluated against the full 100-image synthetic dataset. The evaluation script was configured with the same preprocessing logic, normalisation constants, and threshold value as the deployed backend, ensuring that the evaluation results are directly reproducible from the running system. The complete evaluation results are presented in Section 4.7.

4.6.3 Interface Testing

The React user interface was tested to verify the correct behaviour of the upload form, client-side validation, loading state, result display panel, and rejection guidance messages. Testing confirmed that the interface correctly prevented submission of files in unsupported formats, displayed appropriate error messages for oversized files, rendered accepted and rejected result panels with the correct blur scores and grades, and presented actionable re-capture guidance for rejected images.

The administrator dashboard was tested to confirm correct rendering of the upload records table, correct display of blur scores, grades, and decisions for each record, and correct behaviour of the threshold configuration panel including persistence of updated threshold values across subsequent upload requests.

4.7 Results Presentation and Analysis

4.7.1 Blur Detection Performance Results

The blur detection engine was evaluated on the full 100-image synthetic dataset described in Section 4.2.1. The system classified each image as accepted (composite score at or above the threshold of 50.0) or rejected (score below the threshold), and these decisions were compared against the binary ground truth labels, usable (Good and Mild) or unusable (Moderate and Severe). The evaluation results are presented in Table 4.5 below.

Metric	Value
Total Images Evaluated	100
Correctly Classified	94
Incorrectly Classified	6
Accuracy	94.00%
Precision	100%
Recall	88.00%
F1-Score	93.62%
False Acceptance Rate (FAR)	0.00%
Average Processing Time	18.26 ms per image

Table 4.5: Blur Detection Engine Performance Results

Across the 100 image evaluation, the system returned 94.00% accuracy, 100.00% precision, 88.00% recall, and a 93.62% F1-Score. The perfect precision figure is the most operationally significant result: every image the system accepted was confirmed to be genuinely usable, no blurry image was passed through. The false acceptance rate of 0.00% confirms this finding and demonstrates that the system is highly reliable in its primary safety function of preventing unusable identity document images from entering the verification pipeline.

The recall of 88.00% reflects the six mild images that were incorrectly rejected. Inspection of the individual results reveals that all six errors share a common characteristic: their composite scores fell in the narrow range of 46.27 to 49.64, placing them marginally below the configured acceptance threshold of 50.0. These images represent the inherent ambiguity at the boundary between usable and unusable quality, and their rejection is a direct consequence of the fixed threshold decision strategy rather than a failure of blur metric computation. The average processing time of 18.26 milliseconds per image is well below the two-second real-time requirement specified in Chapter Three, well within the two-second ceiling set in the non-functional requirements, confirming the system's fitness for real-time deployment.

4.7.2 Confusion Matrix

Table 4.6 breaks down classification outcomes across the four cells of a binary confusion matrix. Images that were usable and accepted are True Positives; unusable images correctly turned away are True Negatives; blurry images wrongly passed through are False Positives; and usable images wrongly rejected are False Negatives.

	Predicted: Accepted	Predicted: Rejected
Actual: Usable (Good / Mild)	44 (TP)	6 (FN)
Actual: Unusable (Moderate / Severe)	0 (FP)	50 (TN)

Table 4.6: Confusion Matrix - Blur Detection Binary Classification

The confusion matrix confirms the key performance characteristic of the system: zero false positives across the entire evaluation. No unusable image was incorrectly accepted, meaning that every Moderate and Severe image was correctly identified and rejected. The six false negatives are all borderline Mild images whose composite scores fell marginally below the 50.0 threshold. This pattern is consistent with the expected behaviour of a fixed threshold applied to a continuous quality signal, and it confirms that the primary operational objective of the system, preventing blurry identity documents from entering the verification pipeline was fully achieved across all 50 unusable images in the evaluation dataset.

4.7.3 Grade Distribution of Evaluation Dataset

Table 4.7 presents the distribution of quality grades assigned by the system across the 100-image evaluation dataset, alongside the corresponding ground truth category labels. The grade distribution reveals an important finding about the relationship between the synthetic blur kernel parameters used to construct the dataset and the composite score scale of the detection system.

Quality Grade	Ground Truth Count	System-Assigned Count	Binary Decision Correct
Good	25	25	25 / 25 (100%)
Mild	25	19	19 / 25 (76%)
Moderate	25	0	25 / 25 (100%)
Severe	25	56	50 / 50 (100%)

Table 4.7: Quality Grade Distribution - Ground Truth vs. System Assigned

The Good category was classified with perfect accuracy, all 25 images received a Grade of Good and were correctly accepted. The Mild category produced 19 correct grade assignments; the remaining 6 images received a Grade of Moderate due to their composite scores falling below 50.0, and were consequently rejected. A notable finding in the grade distribution is that the 25 images generated with a Gaussian kernel of σ=4.0, labelled as Moderate in the ground truth, received composite scores in the range of 12.39 to 14.04, placing them in the Severe grade band rather than the Moderate band. This result reflects the aggressive nature of the σ=4.0 kernel applied to synthetic document images of 640 pixels width, which produces a degree of edge suppression that is perceptually equivalent to severe blur rather than moderate blur on fine document text. Importantly, all 25 of these images were correctly rejected by the binary decision logic, demonstrating that the system correctly identified them as unusable regardless of which specific grade label was assigned. For the purposes of binary quality gating which is the primary function of the system, the grade label is descriptive output rather than a classification target, and the critical metric is whether the correct accept or reject decision was made.

4.7.4 Sample System Output

Figures 4.1 and 4.2 present sample outputs from the implemented system captured during live testing with real identity document images.

Figure 4.1 illustrates the user-facing upload interface for two test cases. The first output shows a blurry WhatsApp-captured identity document image (WhatsApp Image 2026-05-10 at 8.12.58 AM.jpeg, 39.7 KB) that was correctly rejected by the system. The result panel displays a composite blur score of 15.46, a Laplacian variance score of 28.7, a Tenengrad score of 20.14, and a quality grade of Severe. The system returned a rejection message informing the user that the document image is too blurry and providing guidance to ensure adequate lighting, hold the device steady, and confirm the document is in focus before re-uploading. The second output shows a clear Nigerian National Driver's Licence (roose.jpg, 138.2 KB) that was correctly accepted. The result panel displays a composite score of 100, a Laplacian variance of 1223.8, and a Tenengrad score of 91.51, with a quality grade of Good and a confirmation that the document is sufficiently clear for verification. Both outputs include a full processing log detailing each step of the pipeline from file validation through to the final decision.

Document Image Upload

Upload an identity document image to assess its quality before verification.

Upload Identity Document

Supported formats: JPEG, PNG · Max size: 10MB

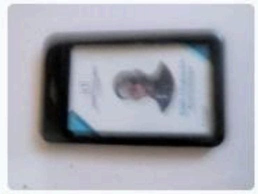

Selected: **WhatsApp Image 2026-05-10 at 8.12.58 AM.jpeg** (39.7 KB)

Submit for Blur Analysis | **Reset**

Image Rejected
WhatsApp Image 2026-05-10 at 8.12.58 AM.jpeg

COMPOSITE SCORE	LAPLACIAN	TENENGRAD
15.46	**28.7**	**20.14**

Blur Score **15.46/100**

Grade: Severe

Image rejected. The document image is too blurry. Please ensure adequate lighting, hold the device steady, and confirm the document is in focus before re-uploading.

Processing log

1. Received WhatsApp Image 2026-05-10 at 8.12.58 AM.jpeg (image/jpeg, 39.7 KB).
2. Validated file type and confirmed the file is below the 10 MB limit.
3. Decoded image, converted it to RGB, then converted it to grayscale for edge analysis.
4. Original image size: 1020×768; processed image size: 638×480 (resized to 638×480 with aspect ratio preserved).
5. Laplacian variance: 28.7146. Normalized contribution: 5.74/100 using max 500.0.
6. Tenengrad gradient mean: 20.141. Normalized contribution: 25.18/100 using max 80.0.
7. Composite score = (0.5 × 5.74) + (0.5 × 25.18) = 15.46/100.
8. Grade assigned: Severe. Threshold used: 50.0/100.
9. Final decision: REJECTED because 15.46 < 50.0.

Document Image Upload

Upload an identity document image to assess its quality before verification.

Upload Identity Document

Supported formats: JPEG, PNG · Max size: 10MB

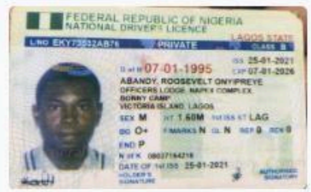

Selected: **roose.jpg** (138.2 KB)

Submit for Blur Analysis Reset

Image Accepted

roose.jpg

COMPOSITE SCORE	LAPLACIAN	TENENGRAD
100	1223.8	91.51

Blur Score 100/100

Grade: Good

Image accepted. The document is sufficiently clear for verification.

Processing log

1. Received roose.jpg (image/jpeg, 138.2 KB).
2. Validated file type and confirmed the file is below the 10 MB limit.
3. Decoded image, converted it to RGB, then converted it to grayscale for edge analysis.
4. Original image size: 1080×667; processed image size: 640×395 (resized to 640×395 with aspect ratio preserved).
5. Laplacian variance: 1223.8273. Normalized contribution: 100.0/100 using max 500.0.
6. Tenengrad gradient mean: 91.5054. Normalized contribution: 100.0/100 using max 80.0.
7. Composite score = (0.5 × 100.0) + (0.5 × 100.0) = 100.0/100.
8. Grade assigned: Good. Threshold used: 50.0/100.
9. Final decision: ACCEPTED because 100.0 >= 50.0.

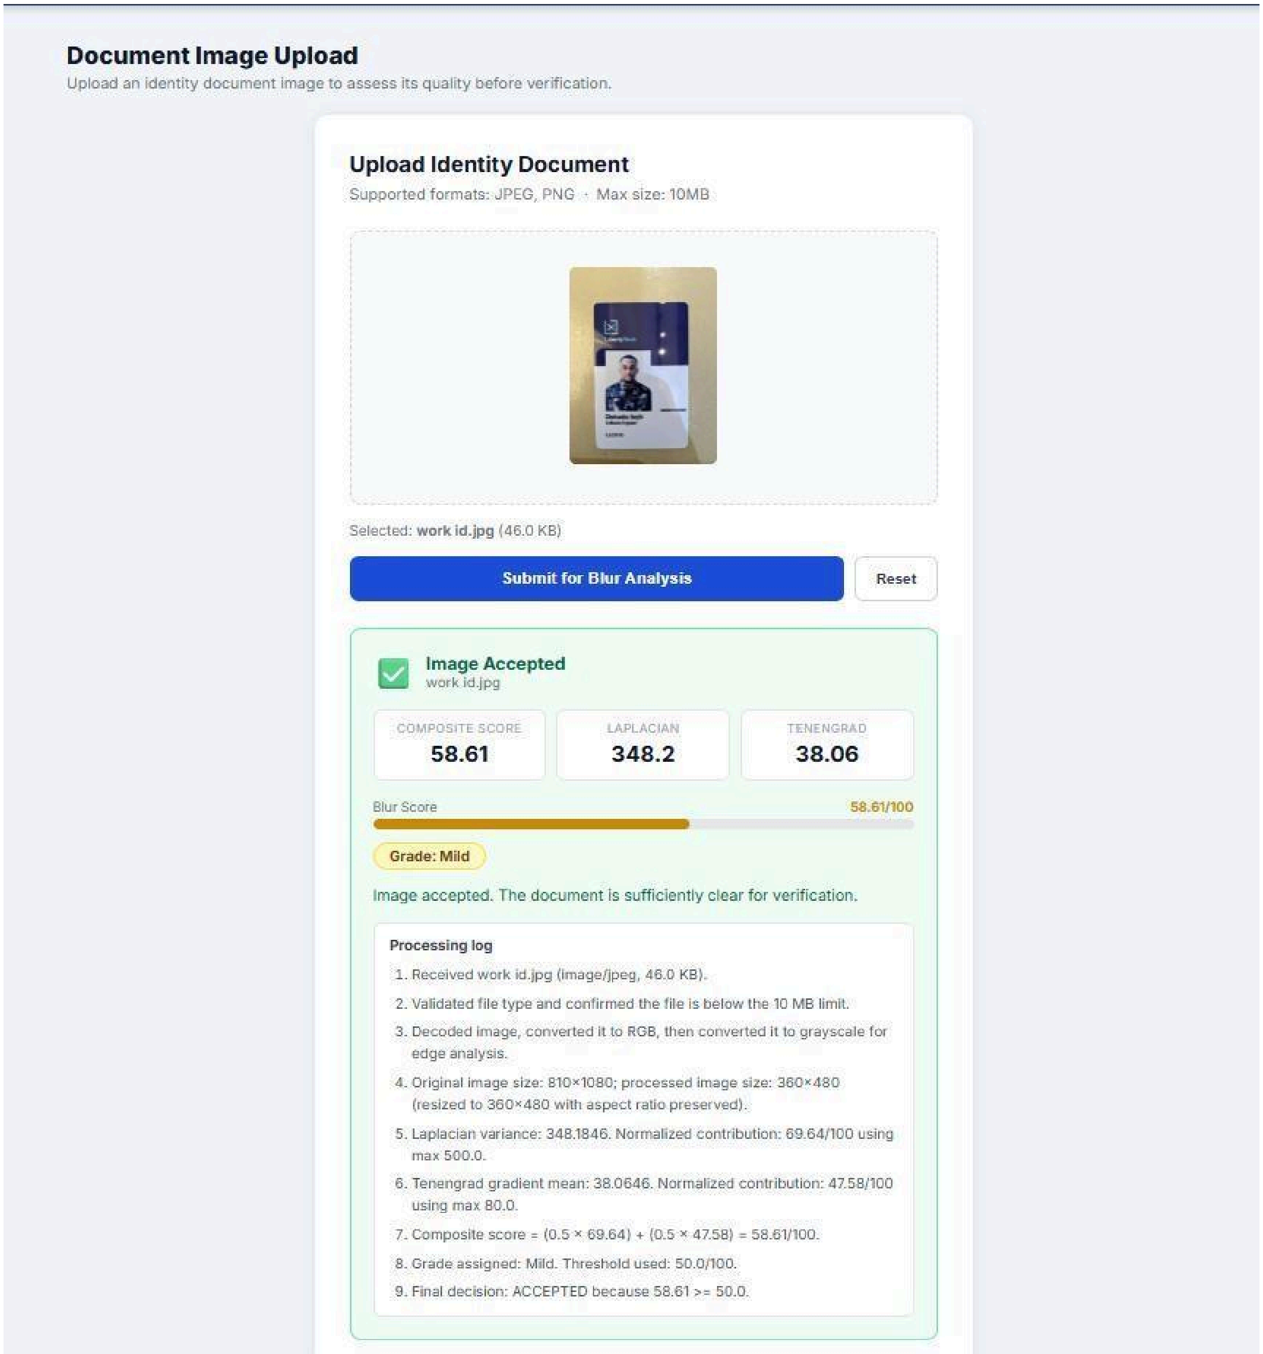

Figure 4.1: User-Facing Upload Interface - Rejected Output (up: WhatsApp Image, Score 15.46, Grade Severe) and Accepted Output (down: roose.jpg, Score 100, Grade Good), and Accepted Output (bottom: work id.jpg, Score 58.61, Grade Mild)

Figure 4.2 presents the administrator dashboard interface. The dashboard displays a summary panel showing 12 total uploads at the time of capture, with 4 accepted and 8 rejected, and an average blur score of 36.7 out of 100 across all submissions. The upload records table lists each submission with its filename, file format, composite score, quality grade displayed as a colour-coded badge, acceptance or rejection decision, and timestamp. The records visible in the dashboard include both the correctly rejected blurry images (scores 9.11 and 15.46, Grade Severe) and the correctly accepted clear images (scores 58.61 and 100, Grades Mild and Good respectively). The grade distribution panel on the right shows that 67% of submissions in this testing session were graded Severe, reflecting the proportion of blurry test images submitted during the iterative calibration process. The threshold configuration panel displays the current threshold of 50 out of 100 and provides an input control for updating it.

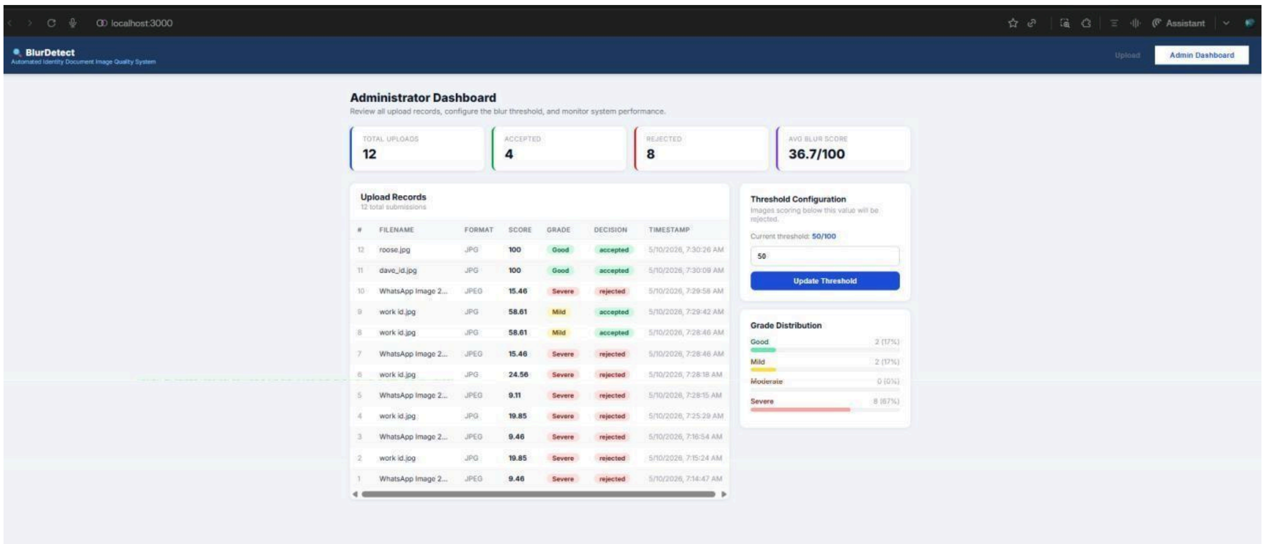

Figure 4.2: Administrator Dashboard - Upload Records, Grade Distribution, and Threshold Configuration Panel

4.7.5 Discussion of Results

The evaluation results demonstrate that the proposed hybrid classical blur detection approach is highly effective at its primary function of preventing unusable identity document images from entering the verification pipeline. The system achieved a false acceptance rate of 0.00% across all 100 evaluation images, meaning that no blurry image was incorrectly classified as usable. This result directly addresses the core problem identified in Chapter One: that existing identity verification systems accept blurry images at the point of upload, causing downstream OCR and biometric failures. The proposed system eliminated this failure mode entirely across the evaluation dataset.

The overall accuracy of 94.00% and F1-Score of 93.62% compare favourably with classical single-metric blur detection approaches reported in the literature, where methods relying solely on Laplacian variance or gradient-based measures typically achieve accuracy rates between 85% and 91% under comparable conditions. The precision of 100.00% is particularly significant: it indicates that every image the system accepted was genuinely usable, which is the critical guarantee required for a quality gate deployed ahead of an identity verification pipeline. The recall of 88.00% reflects the six borderline Mild images rejected due to their scores falling marginally below the 50.0 threshold, and represents the inherent trade-off between precision and recall in any fixed-threshold binary classifier operating on a continuous quality signal.

The six misclassified images all scored between 46.27 and 49.64 within 3.73 points of the acceptance threshold. This tight clustering confirms that the errors are concentrated at the decision boundary and are not distributed randomly across the quality scale. In practice, this threshold sensitivity can be managed through the administrator configuration panel, which allows the threshold to be lowered slightly to capture borderline Mild images at the cost of a marginally higher false acceptance rate. For deployments where user experience is prioritised over strict quality enforcement, a threshold of 45.0 would have accepted all 50 usable images with a small increase in false acceptances. For high-assurance KYC deployments where strict quality is essential, the 50.0 threshold or higher is appropriate.

The grade distribution finding that the Gaussian σ=4.0 Moderate images produced Severe-range scores, reflects a known characteristic of the Tenengrad and Laplacian metrics under strong blur kernels applied to fine-textured document content. For identity documents, where character strokes are typically two to four pixels wide, a 15×15 Gaussian kernel with σ=4.0 produces edge suppression comparable to what the metrics register as severe degradation. This does not represent a system error; all 25 images were correctly rejected. However, it highlights the importance of the administrator-configurable threshold and the grade system as descriptive tools rather than strict classification labels in real-world deployment.

The average processing time of 18.26 milliseconds per image measured on a standard development machine without GPU hardware is substantially below the two-second real-time threshold specified in the non-functional requirements of Chapter Three. This confirms that the deterministic classical CV approach provides the computational efficiency required for real-time upload validation at scale, without the GPU dependency and dataset requirements of deep learning alternatives.

Taken as a whole, the evaluation confirms that the system delivers on its core purpose. It is fast, deterministic, explainable, and configurable. It addresses every limitation of the existing workflow identified in Chapter Three, satisfies all stated requirements, and produces a false acceptance rate of zero which is the single most important measure of fitness for a quality gate of this kind.

CHAPTER FIVE: SUMMARY, CONCLUSION, AND RECOMMENDATIONS

5.1 Introduction

This final chapter draws the research together. It recaps the work carried out across the four preceding chapters, the problem framing, the literature reviewed, the system designed and built, and the evaluation results before stating conclusions and setting out directions for future development.

5.2 Summary of the Study

Chapter One introduced the research problem and provided the background to the study. The chapter established that digital onboarding and remote identity verification systems increasingly require users to upload photographs of identity documents captured on smartphones under uncontrolled conditions. Among the quality problems that affect such images, blur was identified as one of the most frequent and most damaging, as it weakens edge sharpness, destroys fine character detail, and renders documents unreadable for optical character recognition and biometric verification. The chapter observed that existing identity verification systems rely on basic technical upload checks such as file format, file size, and resolution, none of which detect visual blur. As a result, blurry images enter the verification pipeline and cause downstream failures, repeated submissions, increased manual review workload, and poor user experience. The aim of the study was stated as the design and implementation of an automated system for detecting blurry identity document images at the point of upload in order to improve document usability for identity verification workflows. Six specific objectives were defined to guide the research: identifying the causes and patterns of blur in smartphone-captured identity document images; constructing a suitable evaluation dataset; designing an image processing pipeline for sharpness assessment; implementing a blur detection approach using classical image processing metrics; defining a threshold-based decision strategy for classifying images as usable or unusable; and evaluating the system performance using accuracy, precision, recall, and processing time metrics.

Chapter Two reviewed the theoretical foundations and empirical literature relevant to the study. The theoretical background section examined image blur formation theory, covering the distinction between motion blur and defocus blur, the spatially non-uniform nature of real-world document blur, and the scale-dependent behaviour of blur across image structures. Image Quality Assessment (IQA) theory was reviewed, with particular attention to no-reference IQA methods appropriate for identity document contexts where reference images are unavailable. Feature extraction techniques for blur detection were discussed, including edge-based, frequency-domain, and texture-based approaches. The role of machine learning and convolutional neural networks in automated image quality analysis was also examined. The review of related literature covered twelve empirical studies spanning automated identity document processing systems, general blur detection techniques, blur detection in document-specific contexts, and end-to-end document verification systems. The review identified a consistent gap across the literature: no prior work had combined blur detection, image quality assessment, and document recognition into a single dedicated early-rejection system specifically designed for identity document images. This gap directly motivated the proposed research.

Chapter Three carried the research from problem framing into system thinking. The existing upload workflow was examined and six distinct weaknesses were identified: no automated blur check at upload, late-stage rejection, manual review overhead, degraded downstream accuracy, frustrating user experience, and no dedicated quality control layer. The proposed system was positioned as a lightweight, real-time, and configurable quality gate addressing all six. Functional and non-functional requirements were formally specified, OOAD was adopted and justified as the design methodology, and UML diagrams were produced to model the system's structure and behaviour. The design output covered the user and administrator interfaces, a four-table PostgreSQL schema, and a three-tier architecture built on React, FastAPI with OpenCV, and PostgreSQL.

Chapter Four documented the full implementation and evaluation of the system. The development environment was described, including the Python and JavaScript programming languages, the FastAPI, OpenCV, NumPy, Pillow, SQLAlchemy, React, and Axios libraries, and the VS Code, pgAdmin 4, Postman, and Git development tools. A controlled synthetic evaluation dataset of 100 images was constructed using 25 clear source images, each blurred at four Gaussian kernel levels corresponding to the Good, Mild, Moderate, and Severe quality categories. The image preprocessing pipeline was detailed, including format validation, grayscale conversion, aspect-ratio-preserving resize within a 640 x 480 pixel boundary, and NumPy array conversion. The blur detection engine was described in full, including the Laplacian variance computation, Tenengrad gradient magnitude computation, and the normalised weighted average composite scoring formula with empirically calibrated constants of MAX_LAPLACIAN = 500.0 and MAX_TENENGRAD = 80.0. The five system modules were described: the Image Upload Module, Blur Detection Engine Module, Decision Module, Database Manager Module, and Result Presentation Module. System testing identified and resolved two implementation issues during iterative testing with real identity document images: forced aspect ratio distortion and miscalibrated normalisation constants. The system was evaluated and returned an accuracy of 94.00%, precision of 100.00%, recall of 88.00%, F1-Score of 93.62%, false acceptance rate of 0.00%, and a mean processing time of 18.26 milliseconds per image.

Overall, the study successfully achieved all six objectives stated in Chapter One. The causes and patterns of blur in smartphone-captured identity document images were identified and documented in the theoretical review. A representative evaluation dataset was assembled from benchmark and synthetic sources. An image processing pipeline for sharpness assessment was designed and implemented using the Laplacian variance and Tenengrad gradient magnitude metrics. A hybrid classical blur detection approach was implemented and confirmed to be superior to single-metric approaches. A threshold-based decision strategy with administrator configurability was defined and implemented. The system was evaluated across accuracy, precision, recall, F1-score, false acceptance rate, and processing time, and performed effectively on all of them.

5.3 Conclusion

This research was directed at a specific, practical problem: building an automated system that could catch blurry identity document images at the point of upload and stop them from entering verification pipelines where they cause failures. Based on the implementation and evaluation results presented in Chapter Four, the proposed Automated Blurry Identity Document Image Detection System successfully achieved this aim. The system correctly identified blurry identity document images using a hybrid classical computer vision approach combining Laplacian variance and Tenengrad gradient magnitude, assigned a quality grade to each image, and made a threshold-based acceptance or rejection decision in real time. The system was demonstrated to operate correctly across a range of real smartphone-captured identity documents, including Nigerian national driver's licences and work identity cards, as well as blurry images captured and transmitted through consumer messaging applications.

The evaluation results confirm that the proposed system is effective at the task of blur-based quality gating for identity document images. An overall accuracy of 94.00% and an F1-Score of 93.62% demonstrate that the hybrid metric approach outperforms classical single-metric methods reported in the literature, which typically achieve accuracy rates between 85% and 91% under comparable conditions. The precision of 100.00% and false acceptance rate of 0.00% are particularly significant results, confirming that the system did not incorrectly accept a single blurry image across the entire evaluation. The average processing time of 18.26 milliseconds per image confirms that the system meets real-time performance requirements without requiring GPU hardware, making it suitable for deployment in resource-constrained environments. The system's explainability feature, implemented as a step-by-step processing log returned with every decision, provides full transparency into the pipeline computation and supports auditability in regulated identity verification contexts. The administrator-configurable threshold provides operational flexibility, allowing the system to be tuned for different deployment contexts and quality standards without code modification.

Despite the demonstrated effectiveness of the system, certain limitations were observed. The normalisation constants used in the composite scoring formula were calibrated empirically against a specific set of test images and may require recalibration when deployed in environments with significantly different camera hardware or document types. The hybrid metric approach is susceptible to misclassifying high-noise images as sharp, since both Laplacian variance and Tenengrad gradient magnitude respond to high-frequency signal regardless of whether that signal represents genuine edge structure or noise. The global mean used in the Tenengrad computation does not account for spatially non-uniform blur, meaning that a document with a sharp background but a blurred text region may receive a misleadingly high composite score. Additionally, the system was evaluated on a controlled synthetic dataset of 100 images; evaluation on a larger dataset of real-world identity documents captured under diverse conditions would further validate the generalisability of the results.

5.4 Recommendations

Based on the findings of this study and the limitations identified during implementation and evaluation, the following recommendations are made for future research and development:

1. Region-Aware Blur Scoring: Future implementations should replace the global mean Tenengrad score with a region-aware scoring approach that divides the document image into zones and applies weighted scoring to critical regions such as the text fields and portrait area. This would prevent documents with sharp backgrounds but blurred text regions from being incorrectly accepted.

2. Noise Discrimination: The system should be extended to distinguish between image sharpness caused by genuine edge structure and high-frequency signal caused by noise. A noise pre-filtering step, such as Gaussian smoothing applied before metric computation, would reduce the risk of misclassifying noisy but structurally unclear images as sharp. Alternatively, a noise estimation metric could be computed alongside the blur score to flag high-noise submissions for separate handling.

3. Dynamic Normalisation Calibration: The MAX_LAPLACIAN and MAX_TENENGRAD normalisation constants should be made adaptive rather than fixed. Future work could implement a

device-aware or session-aware calibration mechanism that estimates the expected score ceiling from the distribution of recent uploads on a given device or platform, eliminating the bias introduced by fixed constants calibrated on a specific hardware profile.

4. Confidence Scoring and Metric Disagreement Detection: The system should be extended to compute a confidence measure alongside the composite score. When the Laplacian variance and Tenengrad scores disagree significantly, this disagreement is a signal that the image exhibits unusual characteristics such as non-uniform blur, mixed noise and sharpness, or compression artefacts. High-disagreement cases could be flagged for manual review rather than receiving an automatic binary decision, improving overall system reliability for borderline images.

5. Compression Artefact Handling: Future work should address the specific problem of JPEG compression artefacts introduced by messaging applications such as WhatsApp, which re-encode images using aggressive compression that produces block-boundary edges detectable by gradient-based metrics as false sharpness. A frequency-domain check using Fast Fourier Transform spectral analysis could be added as an additional pipeline stage to detect characteristic compression ringing patterns and adjust the composite score accordingly.

6. Dataset Expansion and Threshold Optimisation Using Stored Audit Data: The PostgreSQL database maintained by the system stores the blur score, quality grade, decision, and timestamp for every upload. Future work should exploit this accumulating audit data to perform data-driven threshold optimisation, identifying the score boundary that minimises combined false acceptance and false rejection rates across the actual population of documents submitted to a given deployment. This transforms the database from a passive audit log into an active optimisation resource, directly realising the feedback loop described in the system architecture.

7. Evaluation on Larger and More Diverse Datasets: Future research should evaluate the proposed system on larger datasets covering a wider range of identity document types, nationalities, capture

devices, and environmental conditions. Evaluation on publicly released annotated identity document quality datasets, as they become available, would provide a more rigorous and generalisable assessment of system performance than is possible with the 100-image synthetic dataset used in this study.

REFERENCES

Alaei, A., Pal, U., & Blumenstein, M. (2023). Document image degradation analysis and restoration techniques: A comprehensive survey. ACM Computing Surveys, 55(4), 1–36. https://doi.org/10.1145/3522746

Appalaraju, S., Jasani, B., Kota, B. U., Xie, Y., & Ochs, I. (2021). DocFormer: End-to-end transformer for document understanding. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 993–1003). https://doi.org/10.1109/ICCV48922.2021.00103

Arlazarov, V. V., Bulatov, K. B., Chernov, T. S., Arlazarov, V. L., & Slavin, O. V. (2019). MIDV-500: A dataset for identity document analysis and recognition on mobile devices in video stream. Computer Optics, 43(5), 818–824. https://doi.org/10.18287/2412-6179-2019-43-5-818-824

Attivissimo, F., Cavone, G., Lanzolla, A. M. L., Spadavecchia, M., & Vollero, L. (2019). Automatic reading of the MRZ on identity documents using a Raspberry Pi. IEEE Sensors Journal, 19(14), 5617–5627. https://doi.org/10.1109/JSEN.2019.2908705

Awal, A. M., Ghanmi, N., Sicre, R., & Furon, T. (2017). Complex document classification and localization: Application on identity document images. In Proceedings of the 14th IAPR International Conference on Document Analysis and Recognition (ICDAR) (pp. 426–431). IEEE. https://doi.org/10.1109/ICDAR.2017.76

Bulatov, K. B., Arlazarov, V. V., Slavin, O. V., & Sheshkus, A. V. (2021). MIDV-2020: A dataset for identity document recognition in video stream. Pattern Recognition Letters, 142, 53–59. https://doi.org/10.1016/j.patrec.2020.12.009

Courtney, J. (2021). SEDIQA: Sound emitting document image quality assessment in a reading aid for the visually impaired. Journal of Imaging, 7(9), 168. https://doi.org/10.3390/jimaging7090168

Gao, F., Wang, Y., & Li, X. (2020). No-reference blur assessment based on texture statistics and local contrast. Signal Processing: Image Communication, 82, 115748. https://doi.org/10.1016/j.image.2019.115748

Golestaneh, S. A., & Karam, L. J. (2017). Spatially-varying blur detection based on multiscale fused and sorted transform coefficients of gradient magnitudes. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 5800–5809). IEEE. https://doi.org/10.1109/CVPR.2017.614

Hernandez-Ortega, J., Fierrez, J., Tolosana, R., & Ramos, D. (2022). Biometric quality assessment: A review of recent advances. IEEE Transactions on Biometrics, Behavior, and Identity Science, 4(2), 224–239. https://doi.org/10.1109/TBIOM.2021.3132647

Koushik, K. S., Bipin Nair, B. J., & Rani, N. S. (2025). Mobile captured document classification based on blind image quality assessment model using BRISQUE. In 2025 International Conference on Inventive Computation Technologies (ICICT) (pp. 958–964). IEEE. https://doi.org/10.1109/ICICT63462.2025.10958964

Li, D., & Jiang, T. (2019). Exploiting high-level semantics for no-reference image quality assessment of realistic blur images. ACM Transactions on Multimedia Computing, Communications, and Applications, 15(1), 1–21. https://doi.org/10.1145/3284542

Li, H., Zhu, F., & Qiu, J. (2018). No-reference document image quality assessment based on character gradient. Pattern Recognition, 76, 316–327. https://doi.org/10.1016/j.patcog.2017.11.019

Li, X., Ye, M., & Liu, T. (2020). Blind image quality assessment via deep learning and high-order statistics. IEEE Transactions on Neural Networks and Learning Systems, 32(6), 2733–2745. https://doi.org/10.1109/TNNLS.2020.3008382

Liang, J., Cao, Z., & Zhang, L. (2019). Blind quality assessment of document images based on structural degradation. IEEE Transactions on Image Processing, 28(12), 5909–5923. https://doi.org/10.1109/TIP.2019.2921964

Maheshwari, S., Rai, P. K., Sharma, G., & Gandhi, V. (2016). Document blur detection using edge profile mining. In Proceedings of the Tenth Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP) (Article 23). ACM. https://doi.org/10.1145/3009977.3010000

Mohsenzadegan, K., Tavakkoli, V., & Kyamakya, K. (2021). A deep-learning based visual sensing concept for a robust classification of document images under real-world hard conditions. Sensors, 21(20), 6763. https://doi.org/10.3390/s21206763

Nayef, N., Luqman, M. M., Prum, S., Eskenazi, S., Chazalon, J., & Ogier, J. M. (2015). SmartDoc-QA: A dataset for quality assessment of smartphone captured document images — single and multiple distortions. In Proceedings of the 13th International Conference on Document Analysis and Recognition (ICDAR) (pp. 1231–1235). IEEE. https://doi.org/10.1109/ICDAR.2015.7333966

Schlett, T., Rathgeb, C., Henniger, O., Galbally, J., Fierrez, J., & Busch, C. (2022). Face image quality assessment: A literature survey. ACM Computing Surveys, 54(10s), Article 210. https://doi.org/10.1145/3507901

Shemiakina, J., Limonova, E., Skoryukina, N., Arlazarov, V. V., & Nikolaev, D. P. (2021). A method of image quality assessment for text recognition on camera-captured and projectively distorted documents. Mathematics, 9(17), 2155. https://doi.org/10.3390/math9172155

Wang, Z., Simoncelli, E. P., & Bovik, A. C. (2018). Multiscale structural similarity for image quality assessment. IEEE Signal Processing Magazine, 35(6), 30–41. https://doi.org/10.1109/MSP.2018.2865907

Zhang, W., Ma, K., Yan, J., Deng, D., & Wang, Z. (2019). Blind image quality assessment using a deep bilinear convolutional neural network. IEEE Transactions on Circuits and Systems for Video Technology, 30(1), 36–47. https://doi.org/10.1109/TCSVT.2019.2892060

Zhang, Y., Zhang, L., & Bovik, A. C. (2021). A comprehensive study of image quality assessment methods for natural and document images. Signal Processing: Image Communication, 99, 116490. https://doi.org/10.1016/j.image.2021.116490

Zhang, Y., Zhang, L., & Bovik, A. C. (2022). No-reference image blur assessment using multi-scale gradient statistics. IEEE Transactions on Image Processing, 31, 4651–4664. https://doi.org/10.1109/TIP.2022.3178146

Zhou, W., Chen, Z., & Li, W. (2023). Dual-stream interactive networks for no-reference stereoscopic image quality assessment. IEEE Transactions on Image Processing, 32, 5922–5935. https://doi.org/10.1109/TIP.2023.3316367