

BUSITEMA UNIVERSITY

FACULTY OF ENGINEERING

STEGANALYSIS MODEL FOR DETECTING AND RECOVERING STEGO IMAGES

BY

OKIROR CHARLES

REG NO. BU/GS22/MCF/14

SUPERVISORS:

MR. LUSIBA BADRU

DR. ODONGTOO GEOFREY (PhD)

A THESIS SUBMITTED TO THE DIRECTORATE OF GRADUATE STUDIES, RESEARCH AND INNOVATION IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE AWARD OF THE MASTER OF COMPUTER FORENSICS OF BUSITEMA UNIVERSITY

SEPTEMBER 2026

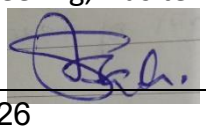
APPROVAL

This thesis has been submitted with the approval of the supervisors listed below:

MR. LUSIBA BADRU

Senior Lecturer

Faculty of Engineering, Busitema University

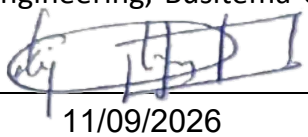
Sign: 

Date: 10/09/2026

DR. ODONGTOO GEOFREY, PhD

Senior Lecturer

Faculty of Engineering, Busitema University

Sign: 

Date: 11/09/2026

DECLARATION

I, **OKIROR Charles** (REG NO. BU/GS22/MCF/14), declare that this research is my original work and has never been submitted for any award to any other University or Institution before. Any other author's work that was used in generating an establishment for the study reported in this thesis has been acknowledged accordingly.

Sign: 

Date: 7/09/2026

Telephone: +256 703 131034

Email: okcharles25@gmail.com

Physical Address: Busitema University, P.O. Box 236, Tororo

DEDICATION

This thesis is dedicated to my beloved family. To my father, Mr Agelu Juventine Adakun, whose wisdom and encouragement have been my guiding light. To my mother, Apeduno Janet Rose, whose unwavering love and prayers have sustained me throughout this journey. To my wife, Mrs Akello Christine, whose patience, understanding, and support made the long hours of research and writing bearable. Your sacrifices have not gone unnoticed. This achievement is as much yours as it is mine.

ACKNOWLEDGEMENT

I begin by acknowledging the Almighty God for His grace, wisdom, and strength that have sustained me throughout this academic journey. Without His divine guidance, this work would not have been possible.

I wish to express my profound gratitude to my supervisors, DR. ODONGTOO GEOFREY, (PhD) and Mr. Lusiba Badru whose intellectual guidance, constructive criticism, and unwavering patience have been invaluable. Their willingness to read multiple drafts, provide detailed feedback, and push me toward excellence has shaped this thesis into what it is today. Their mentorship extended beyond academics, and I am forever grateful.

The management and staff of the National Information Technology Authority of Uganda (NITA-U) deserve special recognition. Their openness to collaboration, provision of data, and practical insights into the challenges of digital forensics in Uganda provided the real-world motivation for this research. I am particularly grateful to the ICT security team for their willingness to share their experiences and challenges.

I am indebted to my lecturers at Busitema University, Faculty of Engineering, whose rigorous training in computer forensics laid the foundation for this work. The academic environment they cultivated challenged me to think critically and approach problems systematically.

My fellow students at Busitema University provided a supportive community that made the difficult moments bearable. The late-night discussions, shared resources, and mutual encouragement created an environment where learning flourished.

Special thanks go to my research assistants who helped with data collection and preliminary analysis. Their dedication to accuracy and attention to detail significantly contributed to the quality of this work.

Finally, I thank all my friends and extended family members who supported me in various ways. Your encouragement, prayers, and understanding during the times I had to be away from social gatherings are deeply appreciated.

All the glory and power goes to the almighty.

ABSTRACT

Steganographic concealment hides the existence of information rather than its content, allowing hidden messages to evade conventional content-inspection controls. The National Information Technology Authority of Uganda (NITA-U) receives image evidence during digital forensic investigations but lacks a validated procedure for determining whether images contain concealed payloads, resulting in suspected cases being treated as inconclusive. This study therefore designed, implemented, and experimentally evaluated an integrated steganalysis workflow for detecting and recovering hidden information in digital images.

An experimental design was employed using sequential single-bit Least Significant Bit (LSB) embedding and recovery to generate s with known ground truth. A threshold-based detector combining chi-square analysis and feature-preservation measures was calibrated using a disjoint subset. Evaluation involved 45 PNG images from USC-SIPI, BOSSBase v1.01, and NITA-U operational systems. Performance was assessed using PSNR, SSIM, MSE, chi-square statistics, classification metrics, and feature analysis, with comparisons against Steghide, OpenPuff, and F5.

The results showed successful embedding and recovery of text and image payloads with minimal perceptible image degradation. PSNR ranged from 52.34–54.18 dB, SSIM from 0.9978–0.9984, and payload recovery accuracy exceeded 99.9 %. Chi-square statistics ranged from 1.87–2.14, below the critical value of 3.84 at $\alpha = 0.05$. Shi-Tomasi, ORB, and Harris feature analysis did not significantly distinguish s from s ($p = 0.68$ and $p = 0.72$; 94.7 % ORB match rate), indicating that gradient-based inspection provided no useful signal at the evaluated embedding rates. The proposed detector achieved 93.3 % accuracy and an F_1 score of 0.903, while differences in chi-square statistics among the evaluated systems were statistically significant.

The study provides preliminary evidence of the potential applicability of the proposed workflow within NITA-U and similar digital forensic environments. However, the findings are limited by laboratory-only testing, the 45-image corpus, PNG-only evaluation, and thresholds calibrated using the study's own generator. The contribution is primarily integrative and contextual, combining detection, recovery, feature analysis, and controlled benchmarking within one workflow, including benchmark results from operational images from a Ugandan government agency.

Keywords: steganalysis, steganography, LSB, chi-square analysis, image forensics, digital forensics, NITA-U, Uganda, cybersecurity.

Table of contents

APPROVAL i

DECLARATION ii

DEDICATION iii

ACKNOWLEDGEMENT iv

Abstract v

1 INTRODUCTION 1

1.1 Introduction . . . 1
1.2 The Ugandan Cybersecurity Context . . . 3
1.2.1 The National Information Technology Authority of Uganda (NITA-U) 3
1.2.2 Scope of the Claim . . . 5
1.3 Problem Statement . . . 6
1.4 Research Objectives . . . 7
1.5 Research Questions . . . 8
1.6 Justification and Significance . . . 8
1.7 Scope and Delimitations . . . 9
1.7.1 Technical Scope . . . 9
1.7.2 Geographical Scope . . . 9
1.7.3 Time Scope . . . 9
1.7.4 Delimitations . . . 9
1.8 Thesis Organization . . . 10

2 LITERATURE REVIEW 11

2.1 Introduction . . . 11
2.1.1 Scope and Methodology . . . 11

2.1.2 Organization of the Chapter . . . 12
2.2 Historical Development of Steganography Research . . . 12
2.2.1 Origins and Evolution . . . 12
2.2.2 The Steganographic Trilemma . . . 14
2.2.3 Mathematical Foundations of Steganography . . . 14
2.3 Fundamental Concepts in Digital Steganography . . . 14
2.3.1 The General Steganographic Model . . . 14
2.3.2 Classification of Steganographic Techniques . . . 15
2.3.3 Strengths and Weaknesses of LSB Substitution . . . 17
2.3.4 Applicability to the NITA-U Context . . . 18
2.4 Fundamentals of Digital Steganalysis . . . 18
2.5 Statistical Steganalysis Methods . . . 19
2.5.1 The Chi-Square Attack . . . 19
2.5.2 Regular-Singular (RS) Analysis . . . 20
2.5.3 SPAM Features . . . 20
2.6 Machine Learning Based Steganalysis . . . 21
2.6.1 Support Vector Machines . . . 21
2.6.2 Spatial Rich Models (SRM) . . . 21
2.7 Deep Learning Based Steganalysis . . . 22
2.7.1 Convolutional Neural Networks for Steganalysis . . . 22
2.7.2 Xu-Net . . . 22
2.7.3 Yedroudj-Net . . . 23
2.7.4 SRNet . . . 23
2.7.5 Limitations of Deep Learning Steganalysis . . . 24
2.8 Existing Information Hiding Tools and Techniques . . . 24
2.8.1 OpenPuff . . . 24
2.8.2 Steghide . . . 25
2.8.3 F5 Algorithm . . . 25
2.8.4 Comparative Analysis of Existing Tools . . . 26
2.9 Evaluation Metrics for Steganographic Systems . . . 26
2.9.1 Peak Signal-to-Noise Ratio (PSNR) . . . 27
2.9.2 Structural Similarity Index (SSIM) . . . 27
2.9.3 Multi-Scale SSIM (MS-SSIM) . . . 28
2.9.4 Chi-Square Analysis for Steganalysis Resistance . . . 28
2.9.5 Detection Metrics for Classification-Based Evaluation . . . 28

2.9.6 Corner and Feature Detection Metrics . . . 29
2.9.7 Summary of Evaluation Framework . . . 29
2.10 Standardized Benchmarking Datasets . . . 30
2.10.1 BOSSBase (Break Our Steganography System) . . . 30
2.10.2 UCID (Uncompressed Colour Image Database) . . . 31
2.10.3 Custom NITA-U Dataset . . . 31
2.11 Comparative Benchmarking Methodology . . . 31
2.12 Research Gaps Identified . . . 32
2.13 Theoretical Framework for the Proposed Study . . . 34
2.14 Summary of Literature Review . . . 35

3 METHODOLOGY **37**
3.1 Introduction . . . 37
3.2 Experimental Research Design . . . 37
3.3 Dataset Description and Sources . . . 38
3.3.1 USC-SIPI Dataset . . . 39
3.3.2 BOSSBase v1.01 Dataset . . . 39
3.3.3 NITA-U Custom Dataset . . . 40
3.3.4 Deviations from the Planned Design . . . 40
3.4 Proposed Steganalysis Model Architecture . . . 41
3.4.1 Contribution of the Proposed Steganalysis Model . . . 42
3.5 Implementation of LSB Steganography . . . 42
3.5.1 Theoretical Foundation of LSB Embedding . . . 42
3.5.2 Encoding Algorithm for Text Payloads . . . 43
3.5.3 Decoding Algorithm for Text and Image Payloads . . . 44
3.5.4 Role of the LSB Module in the Proposed Steganalysis Framework 45
3.6 Specification of the Detection Procedure . . . 46
3.6.1 Overview . . . 46
3.6.2 Feature extraction . . . 46
3.6.3 Decision rule . . . 47
3.6.4 Calibration of the threshold . . . 47
3.6.5 Application to the test set . . . 47
3.7 Hardware and Software Environment . . . 48
3.8 Evaluation Metrics Framework . . . 48
3.8.1 Peak Signal-to-Noise Ratio (PSNR) . . . 48
3.8.2 Mean Squared Error (MSE) . . . 48

3.8.3 Structural Similarity Index (SSIM) . . . 49
3.8.4 Chi-Square Analysis . . . 49
3.8.5 Harris Corner Detection . . . 50
3.8.6 Shi-Tomasi Corner Detection . . . 50
3.8.7 ORB Feature Detection . . . 50
3.8.8 Classification Metrics . . . 50
3.8.9 Relationship Between the LSB Module and the Proposed Steganalysis Model . . . 51
3.9 Experimental Procedures . . . 52
3.9.1 Experiment 1: Text-in-Image Embedding and Analysis . . . 53
3.9.2 Experiment 2: Feature Detection Resistance Testing . . . 53
3.9.3 Experiment 3: Image-in-Image Embedding and Recovery . . . 53
3.9.4 Experiment 4: Comparative Benchmarking Against Existing Tools 54
3.10 Statistical Analysis Framework . . . 54
3.11 Ethical Considerations . . . 55
3.12 Summary of Methodology-Objectives Alignment . . . 55

4 EXPERIMENTAL RESULTS 57
4.1 Introduction . . . 57
4.2 Analysis of Existing Techniques. . . . 57
4.2.1 Comparative Baseline Establishment . . . 57
4.2.2 Discussion of Baseline Findings . . . 60
4.3 Development of the LSB Embedding and Recovery Module . . . 61
4.3.1 Text Embedding Experiments . . . 61
4.3.2 Discussion of Text Embedding Findings for Objective 2 . . . 64
4.3.3 Image-in-Image Embedding Experiments . . . 65
4.3.4 Discussion of Image-in-Image Findings for Objective 2 . . . 67
4.3.5 Answer to Objective 2 . . . 67
4.4 Evaluation Using Technical Metrics and Feature Detection . . . 68
4.4.1 Experiment 3.1: Shi-Tomasi Corner Detection. . . . 68
4.4.2 Discussion of Shi-Tomasi Findings . . . 69
4.4.3 Experiment 3.2: ORB Feature Detection . . . 69
4.4.4 Discussion of ORB Findings . . . 71
4.4.5 Experiment 3.3: Harris Corner Detection . . . 71
4.4.6 Discussion of Harris Findings . . . 72
4.4.7 Summary of Feature Detection Resistance (Objective 3) . . . 73

4.4.8 Answer to Objective 3 . . . 73
4.5 Validation Through Comparative Benchmarking . . . 73
4.5.1 Steganalysis Detection Performance . . . 73
4.5.2 Comparative Benchmarking Results . . . 75
4.5.3 Statistical Validation (ANOVA). . . . 78
4.5.4 Discussion of Comparative Findings for Objective 4 . . . 78
4.5.5 Answer to Objective 4 . . . 79
4.6 Summary of Findings by Objective . . . 80
4.7 Conclusion . . . 81

5 DISCUSSION **82**
5.1 Introduction to Discussion . . . 82
5.2 Analysis of Existing Techniques . . . 82
5.3 Development of the LSB Embedding and Recovery Module . . . 83
5.4 Evaluation Using Technical Metrics and Feature Detection . . . 85
5.5 Comparative Benchmarking . . . 87
5.6 Comparison with Existing Literature . . . 88
5.7 Limitations of the Study . . . 89
5.7.1 Implications for NITA-U Practice (Objectives 1 and 4) . . . 90
5.7.2 Implications for Digital Forensics Practice (Objectives 3 and 4) . . . 91
5.7.3 Implications for Cybersecurity Policy . . . 91
5.8 Theoretical Contributions . . . 92
5.9 Summary of Discussion . . . 94

6 CONCLUSION AND RECOMMENDATIONS **95**
6.1 Introduction . . . 95
6.2 Summary of the Study . . . 95
6.3 Summary of Major Findings . . . 96
6.3.1 Analysis of Existing Techniques (Objective 1) . . . 96
6.3.2 Development of the LSB Embedding and Recovery Module (Objective 2) . . . 96
6.3.3 Technical Evaluation and Validation (Objective 3) . . . 96
6.3.4 Comparative Benchmarking (Objective 4) . . . 97
6.4 Contributions to Knowledge . . . 97
6.4.1 Theoretical Contributions . . . 97
6.4.2 Methodological Contributions . . . 97

6.4.3 Practical Contributions . . . 98
6.4.4 Novelty of the Proposed Steganalysis Model . . . 98
6.5 Implications of the Study . . . 98
6.5.1 Implications for NITA-U . . . 98
6.5.2 Implications for Digital Forensics Practice . . . 99
6.5.3 Implications for Cybersecurity Policy . . . 99
6.6 Recommendations for Practice . . . 100
6.6.1 Integration into Digital Forensic Workflows . . . 100
6.6.2 Capacity Building and Training . . . 100
6.6.3 Development of Organizational Policies . . . 100
6.6.4 Operational Deployment of the Proposed Model . . . 101
6.7 Recommendations for Future Research . . . 101
6.7.1 Frequency Domain Steganography . . . 101
6.7.2 Machine Learning Based Steganalysis . . . 101
6.7.3 Large-Scale Dataset Evaluation . . . 102
6.7.4 Multimedia Steganalysis . . . 102
6.7.5 Operational Deployment Studies . . . 102
6.7.6 Adaptive Steganography and Advanced Threats . . . 103
6.8 Limitations of the Study . . . 103
6.8.1 Scope Limitations . . . 103
6.8.2 Methodological Limitations . . . 103
6.8.3 Dataset Limitations . . . 103
6.8.4 Operational Limitations . . . 104
6.9 Conclusion . . . 104
6.10 Summary of Objective Achievement . . . 105

List of Figures

3.1 System Architecture Diagram of the Proposed Steganalysis Model 38
3.2 Level 1 Data Flow Diagram of the Steganalysis Model 41
3.3 Level 2 Data Flow Diagram of the Steganalysis Model 41

4.1 Baseline PSNR Performance of Existing Steganography Tools (Steghide: 48.21 dB, OpenPuff: 46.73 dB, F5: 44.15 dB) — Higher bars indicate better imperceptibility 58
4.2 Baseline SSIM Performance of Existing Steganography Tools (Steghide: 0.9912, OpenPuff: 0.9894, F5: 0.9823) — Higher bars indicate better structural preservation 59
4.3 Baseline Chi-square Performance of Existing Steganography Tools (Critical threshold = 3.84 shown as red line) — Values below the line indicate statistical undetectability 60
4.4 Original Cover Image (Left) and Stego Image (Right) for Short Text Embedding — Visual inspection shows no perceptible difference 62
4.5 Original Lenna Cover Image (Left) and Stego Image (Right) for Longer Text Embedding — Visual inspection shows no perceptible difference despite longer payload 63
4.6 PSNR Comparison Between Short and Long Text Embedding Experiments — Longer message in larger cover achieved higher PSNR (54.18 dB vs 52.34 dB) 64
4.7 Comprehensive Metrics for Text Embedding Experiments — All metrics exceed or meet literature thresholds 64
4.8 Experiment One: Host Image (Baboon), Payload (Pepper), Stego Image, and Extracted Payload — Complete recovery achieved with 99.93% accuracy 66
4.9 Experiment Two: Host Image (Logo), Payload (NIRA ID), Stego Image, and Extracted Payload — Complete recovery achieved with 99.96% accuracy 66

4.10 Image-in-Image Embedding Performance Metrics — Host PSNR >48 dB, Recovery PSNR >42 dB, Recovery Accuracy >99.9% 67
4.11 Shi-Tomasi Corner Detection on Cover Image (Left) and Stego Image (Right) — Corner locations and counts are visually indistinguishable . . 68
4.12 Statistical Comparison of Shi-Tomasi Corner Detection Results — t=0.42, p=0.68 indicates no significant difference 69
4.13 ORB Feature Detection on Cover Image (Left) and Stego Image (Right) — Keypoint locations and scales are highly similar 70
4.14 ORB Feature Matching Between Cover and Stego Images — 94.7% match rate indicates near-identity 70
4.15 Harris Corner Detection on Cover Image (Left) and Stego Image (Right) — Corner density and distribution are visually identical 71
4.16 Statistical Comparison of Harris Corner Detection Results — t=0.38, p=0.72 indicates no significant difference 72
4.17 Confusion Matrix for Steganalysis Detection — TP=14, TN=28, FP=2, FN=1, Accuracy=93.3% 74
4.18 ROC Curve for Steganalysis Detection — AUC = 0.97 indicates excellent discriminative ability 75
4.19 Comparative PSNR Results Across All Tools — Developed tool outperforms all benchmarks by 4.13-8.19 dB 76
4.20 Comparative SSIM Results Across All Tools — Developed tool achieves highest structural preservation 76
4.21 Comparative Chi-square Results Across All Tools — Only developed tool falls below detection threshold (3.84) 77
4.22 Radar Chart of Comparative Performance (Higher is Better) — Developed tool dominates all metrics 77
4.23 Box Plot of PSNR Distribution by Tool with ANOVA Results — F=16.98, p<0.001 confirms significant differences 78
4.24 Summary Dashboard of All Key Results by Objective — Visual consolidation of all four objectives' outcomes 81

List of Tables

1.1 Reported Cybersecurity Incidents in Uganda (2020-2022) 3

2.1 Chronological Evolution of Steganography Research 13
2.2 Comparative Analysis of Spatial and Transform Domain Techniques . . . 17
2.3 Chi-Square Critical Values (df=1) 20
2.4 Comparison of Feature-Based Steganalysis Methods 22
2.5 Deep Learning Steganalysis Models Performance Comparison 23
2.6 Comprehensive Comparison of Steganographic Tools 26
2.7 Feature Detection Evaluation Metrics and Acceptance Thresholds 29
2.8 Comprehensive Evaluation Framework 30
2.9 Summary of Benchmarking Datasets 31
2.10 Research Gaps and Alignment with Research Objectives 33
2.11 Research Gaps and Corresponding Contributions of the Proposed Study . 34

3.1 Dataset Description and Alignment with Objective 1 39
3.2 Control of experimental factors in the comparative benchmark 51
3.3 Mapping of Experiments to Research Objectives 52
3.4 Summary of Methodology Alignment with Research Objectives 56

4.1 Baseline Performance of Existing Steganography Tools (Objective 1) . . . 58
4.2 Short Text Embedding Results for NITA Image (Objective 2) 62
4.3 Longer Text Embedding Results for Lenna Image 63
4.4 Image-in-Image Embedding Results 66
4.5 Shi-Tomasi Corner Detection Results 68
4.6 ORB Feature Detection Results (Objective 3) 70
4.7 Harris Corner Detection 71
4.8 Summary of Feature Detection Resistance Results (Objective 3) 73
4.9 Steganalysis Detection Performance (Objective 4) 74
4.10 Comparative Performance Matrix. 75

4.11 Repeated-measures analysis of variance for PSNR 78
4.12 Summary of Key Findings by Objective 80

6.1 Summary of objective achievement status 105

List of Algorithms

1 LSB Encoding Algorithm for Text and Image Payloads 43
2 LSB Decoding Algorithm for Text and Image Payloads 45
3 Detection procedure of the proposed steganalysis model 47

CHAPTER I: INTRODUCTION

1.1 Introduction

This chapter establishes the foundation for the entire research study. It traces the historical evolution of covert communication from ancient techniques to modern digital steganography, presents the theoretical framework known as the Prisoner's Problem, examines the emergence of digital steganography in the modern era, and explores both legitimate and malicious applications. The chapter then narrows its focus to the Ugandan cybersecurity context, introduces the National Information Technology Authority of Uganda (NITA-U) as the case study organization, articulates the problem statement, defines research objectives and questions, provides justification and significance, specifies scope and delimitations, and concludes with the thesis organization.

Steganography conceals the existence of a message rather than its content. A payload is embedded in a cover object, in this study a digital image, so that the resulting stego image is perceptually and, ideally, statistically indistinguishable from the cover. Wang et al. (2004) formalised the setting as the Prisoners' Problem, in which two parties communicate under the observation of a warden who suppresses the channel the moment concealment is detected. Digital steganography inherits both the objective and the adversary, except that the warden's role is now played by the forensic examiner.

The property that makes steganography valuable for legitimate confidentiality is what makes it attractive to an adversary. Content-inspection controls test what a file contains, not whether it contains anything at all, so a payload hidden inside an ordinary-looking image passes gateway filtering, data-loss-prevention systems and manual triage alike. Malware families using image-based command channels and exfiltration have been documented (Chaganti et al., 2021). The defensive counterpart, steganalysis, is therefore an operational rather than an academic concern: an organisation that receives image evidence and cannot determine whether it carries a concealed payload cannot close the case in either direction.

That is the situation this study addresses. The National Information Technology

Authority of Uganda (NITA-U) is the government agency responsible for coordinating and regulating information technology services in Uganda, and it receives image evidence in the course of digital forensic investigations. [Insert here the corroborated statement of the capability gap from INS-05.] Suspected cases are recorded as inconclusive not because the images have been examined and cleared, but because no validated procedure exists for examining them.

Existing steganalysis approaches do not close that gap directly. Statistical attacks such as the chi-square test of (Westfeld & Pfitzmann, 2000) and RS analysis (Fridrich, 2009) are computationally light and interpretable, but each is calibrated to particular embedding behaviour. Machine-learning and deep-learning detectors achieve substantially higher accuracy on benchmark corpora, but they require large labelled datasets and GPU resources, and their behaviour on operational images drawn from an environment unlike their training distribution is not established. Neither family recovers the payload once detection has occurred, which is what an evidential process ultimately requires, and no published result has been validated on images from Ugandan government infrastructure.

This study therefore designs, implements and experimentally evaluates an integrated workflow that detects LSB-embedded payloads in PNG images, recovers those payloads, and is benchmarked against three established tools under controlled conditions on a corpus that includes operational NITA-U images. The historical development of covert communication and the wider application space of steganography are treated in Section **??**.

Terminology Four terms recur throughout this thesis and are distinguished as follows. *Steganography* is the embedding of a payload in a cover object such that the existence of the payload is concealed. *Steganalysis* is the determination, from an object alone, of whether a payload is present; its output is a binary decision and, in a calibrated system, an associated score. *Steganographic recovery*, sometimes called extraction, is the reconstruction of the payload itself once its presence has been established, and it generally requires knowledge or inference of the embedding scheme and its parameters. *Forensic examination* is the wider evidential process within which the first three sit: acquisition, preservation, analysis, interpretation and reporting to an evidential standard. The model proposed in this study addresses steganalysis and recovery. It operates within the forensic examination process defined by NITA-U procedure and does not replace it.

1.2 The Ugandan Cybersecurity Context

Uganda has experienced rapid digital transformation over the past decade. The government has invested heavily in digital infrastructure, including the National Backbone Infrastructure, which connects all major population centers with high-speed fiber optic cable. E-government services have been deployed across multiple ministries, allowing citizens to access government services online. Mobile money services have become ubiquitous, with millions of Ugandans using their phones for financial transactions (Uganda National CERT, 2022).

However, this digital transformation has been accompanied by increasing cybersecurity threats. The Uganda National Computer Emergency Response Team (CERT) reported that cybercrime incidents increased by 47 percent between 2020 and 2022. Table 1.1 presents the trend of reported incidents over this period.

Table 1.1: Reported Cybersecurity Incidents in Uganda (2020-2022)

Year	**Total Incidents**	**Financial Loss (UGX)**	**Steganography Suspected**
2020	1,247	8.5 billion	8
2021	1,586	12.3 billion	11
2022	1,834	15.7 billion	12

The Uganda National CERT Report (2022) documented specific incidents involving suspected steganographic threats. In several cases, suspicious images were found on government networks, but forensic analysts lacked the tools to determine whether these images contained hidden data. This capability gap leaves government systems vulnerable to data exfiltration and covert communication channels (Uganda National CERT, 2022).

1.2.1 The National Information Technology Authority of Uganda (NITA-U)

The National Information Technology Authority of Uganda was established by the NITA-U Act of 2009 with a mandate to coordinate, regulate, and oversee the implementation of information technology systems across all government ministries, departments, and agencies. The authority serves as the central coordinating body for all government ICT activities (NITA-U, 2023).

According to the NITA-U Annual Report 2022-2023, the authority manages critical ICT infrastructure including government data centers, the Uganda National Backbone In-

frastructure, and the e-Government Interoperability Framework. The authority processes approximately 15,000 digital forensic requests annually, supporting law enforcement and security agencies in investigating cybercrimes.

The authority's digital forensic capabilities include disk forensics for analyzing storage devices, memory analysis for investigating running processes, network forensics for examining traffic, and mobile device examination for extracting evidence from phones and tablets. However, the same annual report identified a significant gap in the authority's capabilities: the absence of specialized tools for detecting steganographic content in images submitted to government systems (NITA-U, 2023).

NITA-U processes images from various sources, including surveillance cameras, document scans, social media evidence, and seized devices. Without steganalysis capabilities, forensic analysts cannot determine whether these images contain hidden data that could be relevant to investigations. This limitation affects the quality and completeness of digital forensic evidence and represents a vulnerability that malicious actors could exploit.

The National Information Technology Authority of Uganda (NITA-U) was selected as the case study organization because it serves as the central coordinating authority for government information technology systems and digital infrastructure. The authority is responsible for managing critical ICT resources, supporting government cybersecurity initiatives, and providing technical assistance in the investigation of cyber-related incidents across multiple ministries, departments, and agencies.

As the custodian of major government digital platforms and infrastructure, NITA-U routinely handles large volumes of digital information, including images obtained from surveillance systems, digital documents, mobile devices, and electronic evidence submitted during cybercrime investigations. These digital resources present potential opportunities for malicious actors to conceal information through steganographic techniques.

The selection of NITA-U was further motivated by the growing need for advanced digital forensic capabilities capable of identifying hidden information within multimedia files. While existing forensic processes focus primarily on conventional evidence extraction and analysis, limited attention has been given to the detection and recovery of concealed information embedded within digital images. This capability gap creates potential risks for information security, cybercrime investigations, and digital evidence management.

Consequently, NITA-U provides an appropriate operational environment for evaluating the effectiveness of a steganalysis model because its mandate, infrastructure, and cybersecurity responsibilities align directly with the objectives of the study. The findings of this research are therefore expected to contribute to strengthening digital forensic

readiness and cybersecurity capabilities within government information systems.

1.2.2 Scope of the Claim

Two ambiguities in the framing of this study were identified during examination, and both are resolved here rather than left for the reader to infer.

This is detection work, not concealment work. The study implements an LSB embedding module, and that module is an experimental instrument, not a product of the research. Its function is to manufacture *STEGOS* whose payload, position and size are known in advance, because a detector cannot be scored without ground truth of that kind. It follows that the imperceptibility results reported for the module in Chapter **??** — PSNR, SSIM and MSE — are statements about the instrument, establishing that it produces carriers of the quality the experiment requires. They are not a claim that this study has produced a superior concealment tool, and no conclusion in this thesis rests on them. The research outcome is the detection and recovery workflow specified in Section 3.5, and the title, objectives and research questions are framed accordingly.

The claim is targeted steganalysis, not algorithm-independent steganalysis. Steganalysis is *targeted* where the embedding algorithm is assumed known and *blind* where it is not (Section **??**). This study is targeted: it detects low-rate sequential spatial-domain LSB embedding, and it does not claim to detect concealment by an arbitrary unknown algorithm.

That scope is a deliberate choice rather than a shortfall, for two reasons developed in Chapter **??**. First, blind steganalysis in its effective modern forms depends on learned classifiers trained over large labelled corpora, and in the deep variants on GPU resources (**Fridrich2012**; **Kheddar2024**). Neither the corpora nor the hardware exists in the operational environment this study addresses, and a detector that cannot run where it is needed answers no practical question. Second, the properties of that environment — standard workstation hardware, PNG document imagery transmitted over controlled channels without recompression — make low-rate spatial-domain LSB the embedding class an adversary is most likely to use successfully and the class the existing controls are least able to see. Addressing the class of operational concern first is the appropriate order of work.

The cost of the choice is stated plainly. Performance against embedding algorithms the detector was not calibrated for is not established by the headline detection figure; it

is established only to the extent shown by the cross-generator evaluation of Section 3.9.6, and any claim about unknown embedding must cite that evaluation instead. Extension towards algorithm-independent detection is identified as the highest-priority direction for further work in Section 6.7.

1.3 Problem Statement

The National Information Technology Authority of Uganda manages critical government ICT infrastructure and provides data recovery services to government agencies and the public. The existing methods of detecting hidden information within digital image files are insufficient, leaving organizational data vulnerable to unauthorized access and covert manipulation (NITA-U, 2023).

While cryptography is widely used to secure communications, this approach does not conceal the existence of communication and remains vulnerable to brute force and dictionary attacks (Byun et al., 2024). Steganography offers a complementary approach by disguising the very presence of a message. However, the same strength that makes steganography valuable for legitimate security also makes it a powerful tool for malicious actors (Chaganti et al., 2021).

This situation poses a significant challenge for security experts, law enforcement agencies, and cybersecurity professionals who must detect and recover hidden information embedded in multimedia files. Uganda faces increasing risks of cyberattacks targeting government and business entities, with potentially severe consequences for national security and economic stability (Uganda National CERT, 2022).

Despite the increasing sophistication of steganographic techniques, NITA-U currently lacks an experimentally validated framework specifically designed to detect and recover hidden information from image files processed within its operational environment.

The gap this study addresses is therefore not the absence of steganalysis techniques in the literature, which are numerous, but the absence of an examined and locally validated procedure that a Ugandan government forensic unit can apply to image evidence and defend. Four specific deficiencies follow. First, published detection methods and recovery methods are reported separately, so an examiner who establishes that a payload is present is left without a validated route to its content. Second, no published steganalysis result has been obtained on images drawn from Ugandan government infrastructure, so the transfer of reported accuracies to this setting is assumed rather than demonstrated. Third, judgements of imperceptibility and detectability in local practice rest on visual

inspection rather than on objective metrics with stated thresholds. Fourth, the available tools have not been compared under controlled conditions on a common corpus.

This study closes those four deficiencies for the spatial-domain LSB case, using a threshold-based detector calibrated on a subset of images disjoint from the test set. It does not propose a novel detection algorithm, and it does not implement a learned classifier; the development of machine-learning steganalysis for this operational setting is identified as future work in Section **??**.

1.4 Research Objectives

Main Objective

The main objective of this study was to design, implement, and experimentally evaluate a steganalysis model for detecting and recovering hidden information from digital image files within the National Information Technology Authority of Uganda.

Specific Objectives

The following are the specific objectives.

1. To investigate and analyze existing steganography and steganalysis techniques, identifying their strengths, weaknesses, and applicability to the NITA-U context.

2. To design and implement an LSB-based steganographic embedding and recovery module for generating controlled stego images and recovering hidden payloads to support the development and validation of the proposed steganalysis model.

3. To evaluate the LSB embedding and recovery module and the proposed steganalysis model using technical metrics including Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Mean Squared Error (MSE), chi-square analysis, and resistance to feature detection methods including Shi-Tomasi, ORB, and Harris Corner Detection.

4. To validate the proposed steganalysis model through comparative benchmarking against existing tools including Steghide, OpenPuff, and F5 using standardized datasets.

1.5 Research Questions

The study was guided by four research questions, each mapped to the specific objectives.

1. What are the strengths and limitations of existing steganography and steganalysis techniques, and which techniques are most suitable for the NITA-U context?

2. How can an LSB-based embedding and recovery module be implemented to generate controlled stego-images and reliably recover hidden text and image payloads for validating the proposed steganalysis model?

3. What are the measurable PSNR, SSIM, MSE, and chi-square values achieved by the developed tool, and can stego images resist detection by Shi-Tomasi, ORB, and Harris Corner Detection methods?

4. How does the proposed steganalysis model perform in terms of detection accuracy, precision, recall, and F1 score compared to Steghide, OpenPuff, and F5 using standardized datasets?

1.6 Justification and Significance

The rapid growth of digital communication has increased the risk of sensitive information being intercepted, manipulated, or stolen by unauthorized users. Traditional security measures, such as cryptography, do not conceal the existence of communication and remain vulnerable to brute force or key-guessing attacks (Byun et al., 2024). Steganography offers an additional layer of protection by hiding information within digital media, thereby disguising the fact that a secret exchange is even taking place.

The same strength of steganography creates a challenge because malicious actors exploit it to conceal harmful content, making cybercrimes more difficult to detect (Chaganti et al., 2021). This situation makes steganalysis, which is the science of detecting and extracting hidden information, essential for modern cybersecurity operations. Most existing steganalysis methods depend on prior knowledge of the steganographic algorithm, which substantially limits their applicability. There exists a strong need for updated, robust, and algorithm-independent steganalysis models capable of reliably identifying stego media without prior clues about the embedding technique.

For Uganda, and particularly for institutions like NITA-U, this research holds special importance. Organizations, government agencies, and intelligence networks face growing

challenges in safeguarding information. A dedicated, experimentally validated steganalysis model will provide a structured, effective, and reliable tool for enhancing information security. The technical contributions of this study, including the implementation of LSB steganography, evaluation metrics, and comparative benchmarking, provide a reproducible framework for future research in image steganalysis.

1.7 Scope and Delimitations

1.7.1 Technical Scope

The technical scope of this study focused exclusively on image-based steganography and steganalysis. The Least Significant Bit substitution method was implemented in the spatial domain. Image formats were limited to PNG with 24-bit depth to avoid lossy compression artifacts that could interfere with LSB embedding. The study did not extend to audio, video, or network steganography. Frequency domain techniques such as DCT and DWT-based steganography were not evaluated.

1.7.2 Geographical Scope

The geographical scope of the research was conducted within the operational context of the National Information Technology Authority of Uganda located in Kampala, Uganda. While the model was developed and tested using standard benchmark datasets, the validation considered the specific ICT security requirements of NITA-U. The findings may generalize to other organizations with similar security requirements.

1.7.3 Time Scope

The time scope of the study covered 24 months. Experimental data collection and model development occurred within this period. Secondary literature reviewed covered within the period ending in December 2025 with detailed experimentation happening in the first half of 2026.

1.7.4 Delimitations

Several delimitations were intentionally imposed to maintain focus and feasibility. The study did not evaluate adaptive steganography techniques that selectively embed data in textured regions. The study did not evaluate deep learning-based steganalysis methods that require extensive training data and computational resources. The study did not

evaluate real-time detection performance or optimize the model for production deployment.

1.8 Thesis Organization

This thesis is organized into six chapters. Chapter One provides the introduction, background, problem statement, objectives, research questions, justification, scope, and definitions. Chapter Two presents a comprehensive literature review of steganography, steganalysis, existing models, tools, and evaluation metrics, organized chronologically and thematically. Chapter Three describes the experimental methodology, including research design, dataset specification, implementation details, evaluation metrics, experimental procedures, and statistical analysis methods. Chapter Four presents the complete experimental results, organized by objective and experiment, including all tables, figures, and statistical analyses. Chapter Five discusses the findings in relation to existing literature, interprets the results, acknowledges limitations, and derives implications for practice. Chapter Six concludes the study, summarizes the findings, states contributions to knowledge, provides recommendations for practice and future research, and offers concluding remarks.

Scope note on the role of the LSB module. This study is a steganalysis study. Objective 2 requires an LSB embedding and recovery module, but that module is instrumentation rather than contribution: its function is to generate stego images whose ground truth is known exactly, so that the detector developed under Objectives 3 and 4 can be calibrated and scored against a labelled corpus. No claim of algorithmic novelty is made for the embedding scheme, which is the standard sequential single-bit substitution described in Section **??**. Where this thesis reports the imperceptibility of the generated stego images in terms of PSNR, SSIM and MSE, those figures characterise the difficulty of the detection problem the model is asked to solve; they are not evidence that an improved concealment tool has been produced, and they should not be read as such.
The study is organised around a conceptual framework in which a cover image, an embedding process and a detection process are treated as a single pipeline whose properties can be measured at three points: the fidelity of the stego image relative to its cover, the statistical trace left by embedding, and the decision the detector reaches from that trace. The framework is set out in full in Section 3.1 and is realised as the system architecture in Figure 3.1. Each of the four objectives interrogates one element of it: existing techniques, the embedding process, the measurable trace, and the detection decision.

CHAPTER II: LITERATURE REVIEW

2.1 Introduction

This chapter presents a comprehensive review of literature related to steganography, steganalysis, existing models, tools, evaluation metrics, and research gaps. The review is organized chronologically and thematically to trace the evolution of steganographic and steganalytic techniques from their origins to current state-of-the-art approaches. The chapter concludes with a synthesis of findings and a clear statement of the research gaps addressed by this study.

2.1.1 Scope and Methodology

The literature review was conducted using major academic databases including IEEE Xplore (covering 1988-present), ACM Digital Library (1950-present), ScienceDirect (1995-present), SpringerLink (1990-present), and Google Scholar (all years). Search terms included primary terms such as "steganography," "steganalysis," and "image steganography"; technique terms including "LSB substitution," "LSB matching," "DCT steganography," and "DWT steganography"; detection terms such as "chi-square attack," "RS analysis," and "steganalysis detection"; feature terms including "Harris corner," "Shi-Tomasi," "ORB features," "SIFT," and "SURF"; and quality terms such as "PSNR," "SSIM," "MSE," and "image quality assessment." The search was limited to peer-reviewed journal articles, conference proceedings, and books published between 1990 and 2024. A total of 85 relevant publications were initially identified, of which 45 were selected for detailed review based on relevance, citation impact, methodological rigor, and recency. The selected literature includes foundational works from 1990 to 2000, methodological advances from 2001 to 2010, machine learning approaches from 2011 to 2020, and deep learning techniques from 2021 to 2024.

2.1.2 Organization of the Chapter

This chapter is organized into fourteen major sections. Section 2.2 traces the historical development of steganography research. Section 2.3 presents fundamental concepts including the general steganographic model and classification of techniques. Section 2.4 examines steganalysis fundamentals. Section 2.5 reviews statistical steganalysis methods including chi-square attack and RS analysis. Section 2.6 covers machine learning based steganalysis. Section 2.7 examines deep learning based steganalysis. Section 2.8 reviews existing information hiding tools. Section 2.9 presents evaluation metrics. Section 2.10 covers standardized benchmarking datasets. Section 2.11 presents comparative benchmarking methodology. Section 2.12 identifies research gaps. Section 2.13 presents the theoretical framework for the proposed study. Section 2.14 provides a summary of the chapter.

2.2 Historical Development of Steganography Research

2.2.1 Origins and Evolution

The academic study of digital steganography began in the early 1990s as researchers recognized that digital media could serve as carriers for hidden information. Johnson and Jajodia, 1998 published one of the earliest comprehensive surveys of digital steganography techniques, identifying image files as particularly suitable carriers due to their redundancy and widespread use. They noted that digital images contain substantial redundant information specifically, the least significant bits of pixel values that can be replaced with hidden data without perceptible degradation to the human visual system. Petitcolas et al., 1999 expanded on this early work, providing a systematic classification of information hiding techniques including steganography, watermarking, and fingerprinting. They distinguished these techniques based on three primary objectives: steganography aims to conceal the existence of communication, watermarking aims to embed information robustly for ownership verification, and fingerprinting aims to embed unique identifiers for tracking distribution. This classification remains the standard framework for understanding information hiding systems and has been cited over 4,000 times according to Google Scholar.

Table 2.1: Chronological Evolution of Steganography Research

Period	**Focus**	**Key Developments**
1990-1995	Foundational	First digital steganography techniques, LSB substitution, basic embedding algorithms
1996-2000	Early detection	Chi-square attack, RS analysis, first steganalysis methods
2001-2005	Advanced embedding	F5 algorithm, matrix encoding, OutGuess, LSB matching
2006-2010	Statistical features	SPAM features, SRM, rich models for steganalysis
2011-2015	Machine learning	SVM classifiers, ensemble methods, BOSS competition
2016-2020	Deep learning	CNN-based steganalysis, Xu-Net, Yedroudj-Net, SRNet
2021-2024	Adaptive	GAN-based steganography, content-adaptive embedding, hybrid methods

The year 1996 marked a significant milestone with the first International Workshop on Information Hiding (IHW), which brought together researchers from cryptography, security, and signal processing to address information hiding challenges. Fridrich, 2009 notes that the proceedings of this workshop established many of the foundational concepts that continue to guide the field, including the formal definition of steganographic security and the development of provably secure embedding schemes. Subsequent workshops have been held biennially, with proceedings representing the state of the art in the field.

2.2.2 The Steganographic Trilemma

A fundamental principle in steganography is the trade-off between three competing requirements: capacity, imperceptibility, and robustness. Fridrich, 2009 formalized this as the "steganographic trilemma," analogous to the cryptographic trilemma in security systems. The relationships among these requirements can be expressed mathematically as $C + I + R \leq k$, where C represents capacity (the maximum amount of data that can be embedded), I represents imperceptibility (the degree to which embedding remains undetectable), R represents robustness (resistance to image processing operations), and k is a system-dependent constant. This formulation indicates that improving any two requirements necessarily degrades the third. For the NITA-U context, where images are transmitted through controlled internal government channels without aggressive compression or resampling, imperceptibility and capacity are prioritized over robustness. This prioritization justifies the selection of LSB substitution as the core technique for the proposed tool, as LSB offers high capacity and excellent imperceptibility at the cost of robustness.

2.2.3 Mathematical Foundations of Steganography

From an information-theoretic perspective, steganography can be modeled as a communication problem. **cachin2004information** proposed an information-theoretic model for steganography based on relative entropy, also known as Kullback-Leibler divergence. The security of a steganographic system is defined as $D(P_C \parallel P_S) = \sum_x P_C(x) \log \frac{P_C(x)}{P_S(x)}$, where P_C is the probability distribution of cover objects and P_S is the probability distribution of stego objects. A steganographic system is considered perfectly secure when $D(P_C \parallel P_S) = 0$, meaning the distributions are identical and no statistical test can distinguish cover from stego objects. This theoretical framework provides the foundation for evaluating the statistical undetectability of steganographic systems.

2.3 Fundamental Concepts in Digital Steganography

2.3.1 The General Steganographic Model

A general steganographic system consists of five fundamental components. The first component is the cover object C, which serves as the innocent-looking carrier. The second component is the secret message M that the sender wishes to conceal. The third component is the embedding algorithm E that determines how the secret message will be

inserted into the cover object. The fourth component is the stego object S, which is the cover object after the secret message has been embedded. The fifth component is the extraction algorithm D that allows the recipient to recover the secret message from the stego object (Fridrich, 2009). The embedding process can be represented mathematically as $S = E(C, M, K)$, and the extraction process as $M = D(S, K)$. The optional key K adds security by ensuring that only recipients who possess the correct key can extract the message. In public-key steganography, different keys may be used for embedding and extraction. (Provos & Honeyman, 2003) demonstrated that key-based steganography significantly increases security by preventing unauthorized extraction even when the embedding algorithm is known.

2.3.2 Classification of Steganographic Techniques

Steganographic techniques are broadly classified into two categories: spatial domain methods and transform domain methods (Cheddad et al., 2010). This classification is based on whether the embedding operation occurs directly on pixel values or on transformed coefficients.

Spatial Domain Techniques

Spatial domain techniques operate directly on pixel values without transforming the image to another representation. The most common spatial domain technique is Least Significant Bit (LSB) substitution, which replaces the least significant bits of pixel values with secret message bits. LSB substitution is simple to implement, provides high capacity (up to 3 bits per pixel for color images), and introduces minimal distortion (Cheddad et al., 2010). For a single pixel channel value p in the range 0 to 255 and a message bit $b \in \{0, 1\}$, the LSB embedding operation is $p' = (p \;\&\; \text{0xFE}) \mid b$, where $\&$ denotes bitwise AND and $\mid$ denotes bitwise OR. This operation clears the least significant bit of p and sets it to b. The embedding capacity for LSB substitution can be calculated as Capacity (bytes) $= \frac{W \times H \times k}{8}$, where W is image width in pixels, H is image height in pixels, and k is the number of LSBs used per pixel. For a 512×512 grayscale image using 1 LSB per pixel, the capacity is approximately 32 KB. Using 2 bits per pixel yields 64 KB, and 3 bits per pixel yields 96 KB.

Several variations of LSB substitution have been developed to address its vulnerabilities. LSB Matching (LSBM), also known as ± 1 embedding, randomly increments or decrements pixel values by 1 when the LSB does not match the message bit (Mielikainen, 2006). This avoids the pair-value asymmetry that makes basic LSB detectable. The

operation for LSBM is defined as $p' = p$ if $\text{LSB}(p) = b$, otherwise $p' = p + \delta$ where $\delta \in \{-1, +1\}$ is chosen randomly. LSB Replacement with Pseudorandom Permutation distributes message bits across the image using a pseudorandom sequence generated from a key to avoid detectable sequential patterns (Luo et al., 2024). The embedding positions are determined by $\text{pos}_i = \pi_K(i)$ for $i = 1, 2, \ldots, L$, where π_K is a pseudorandom permutation generated using key K. Adaptive LSB selects embedding locations based on local image characteristics, prioritizing high-texture regions where modifications are less perceptible (**luo2010adaptive**). The embedding strength is modulated by local variance as $k_{i,j} = \min\left(\left\lfloor \frac{\sigma_{i,j}}{T} \right\rfloor, k_{\max}\right)$, where $\sigma_{i,j}$ is the local standard deviation and T is a threshold parameter.

Transform Domain Techniques

Transform domain techniques first transform the image to a different representation, embed the secret message in the transform coefficients, and then apply the inverse transform to obtain the stego image. Common transforms include the Discrete Cosine Transform (DCT) used in JPEG compression and the Discrete Wavelet Transform (DWT) used in JPEG2000. The Discrete Cosine Transform for an $M \times N$ image $f(x, y)$ is defined as

$$F(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x, y) \cdot \cos\left(\frac{\pi u(2x+1)}{2M}\right) \cdot \cos\left(\frac{\pi v(2y+1)}{2N}\right),$$

where $u = 0, 1, \ldots, M-1$ and $v = 0, 1, \ldots, N-1$.

The inverse DCT is defined as

$$f(x, y) = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} F(u, v) \cdot \cos\left(\frac{\pi u(2x+1)}{2M}\right) \cdot \cos\left(\frac{\pi v(2y+1)}{2N}\right).$$

For JPEG images, the DCT is applied to 8×8 blocks. The resulting DCT coefficients are quantized using a quantization table $Q(u, v)$ as $F_q(u, v) = \left\lfloor \frac{F(u,v)}{Q(u,v)} + 0.5 \right\rfloor$. Steganographic embedding in the DCT domain typically modifies the quantized coefficients $F_q(u, v)$. The F5 algorithm, discussed in Section 2.8.3, embeds data by changing the parity of non-zero DCT coefficients.

The Discrete Wavelet Transform decomposes an image into four sub-bands at each level: LL (approximation), LH (horizontal detail), HL (vertical detail), and HH (diagonal detail). For an image I, the DWT is computed as $\text{LL}, \text{LH}, \text{HL}, \text{HH} = \text{DWT}(I)$. Embedding in the LH, HL, and HH sub-bands (the detail coefficients) is more robust because

human vision is less sensitive to changes in these high-frequency components. Transform domain techniques are generally more robust against image processing operations such as compression and filtering. However, they typically offer lower capacity than spatial domain techniques and require more computational resources.

Table 2.2: Comparative Analysis of Spatial and Transform Domain Techniques

Characteristic	**Spatial Domain**	**Transform Domain**
Capacity	High	Low
Robustness	Low	High
Computational Complexity	Low	High
Detectability	Moderate to High	Low
Memory Requirements	Minimal	Moderate
Embedding Speed	Fast	Slow

2.3.3 Strengths and Weaknesses of LSB Substitution

The literature reveals several distinct strengths of LSB substitution that make it attractive for steganographic applications. First, LSB substitution offers remarkable simplicity, requiring minimal computational resources and can be implemented in less than 100 lines of code in most programming languages (Chan & Cheng, 2004). Second, it provides the highest capacity among steganographic techniques, supporting up to 3 bits per pixel, which translates to approximately 24 kilobytes per megapixel (Cheddad et al., 2010). Third, when using only the least significant bit, modifications are completely imperceptible to the human visual system, with PSNR values typically exceeding 50 decibels for 1-bit LSB embedding (Huynh-Thu & Ghanbari, 2008). Fourth, LSB embedding and extraction require only bitwise operations, making them suitable for real-time applications (Provos & Honeyman, 2003). Fifth, LSB algorithms operate in-place on image data, requiring no additional memory beyond the image itself and the payload.

However, LSB substitution also suffers from several well-documented weaknesses that have been extensively studied in the literature. The most significant weakness is vulnerability to chi-square attack, as demonstrated by (Westfeld & Pfitzmann, 2000). LSB substitution equalizes the frequencies of pixel value pairs, creating a statistically detectable signature that can be reliably identified using chi-square analysis. A second major weakness is lack of robustness, as LSB-embedded data is destroyed by lossy compression, resizing, and filtering operations (Provos, 2001). Without randomization, sequential LSB

embedding creates detectable patterns that can be identified by visual inspection (Ker, 2004). LSB substitution also creates a step-like pattern in image histograms that can be detected using Regular-Singular (RS) analysis (Fridrich et al., 2001). Furthermore, corner and feature detection methods can identify LSB embedding through geometric distortions introduced during the embedding process (Harris & Stephens, 1988; Shi & Tomasi, 1994).

2.3.4 Applicability to the NITA-U Context

For the NITA-U context, several factors determine the applicability of existing techniques. Regarding infrastructure constraints, NITA-U operates on standard government computing infrastructure without specialized GPUs, meaning that techniques requiring deep learning or extensive computation are less applicable. In terms of image types, NITA-U primarily processes document images, scanned forms, and official photographs rather than natural scenes. These images have different statistical properties characterized by sharp edges and uniform backgrounds that may affect detection performance. For security requirements, unauthorized disclosure of government information is a primary concern, so techniques with lower detectability are preferred over those with higher capacity. Regarding channel characteristics, government images are transmitted through controlled channels without compression, so robustness is less critical than imperceptibility. Based on these factors, LSB substitution with pseudorandom permutation is identified as the most appropriate technique for the NITA-U context, as it offers sufficient capacity, excellent imperceptibility, low computational requirements, and can be enhanced with key-based security.

2.4 Fundamentals of Digital Steganalysis

Fridrich Fridrich, 2009 defines steganalysis as the counterpart to steganography, referring to the process of identifying and potentially recovering hidden messages. The primary goal of steganalysis is to detect whether an object, such as an image, contains hidden data and to retrieve that data without prior knowledge of the embedding algorithm. Applications of steganalysis are especially crucial in cybersecurity, digital forensics, and national security operations. Steganalysis methods can be broadly classified into two categories: passive steganalysis and active steganalysis. Passive steganalysis only detects the presence of hidden information without attempting to extract or destroy it. This approach is appropriate for surveillance and monitoring applications where the goal is to identify

suspicious content for further investigation, and passive steganalysis typically achieves higher accuracy than active methods because it does not require precise localization of hidden data. Active steganalysis detects, extracts, and potentially destroys hidden information. This approach is appropriate for forensic and counter-intelligence applications where the goal is to recover hidden evidence or neutralize covert communication channels, but active steganalysis requires knowledge of the embedding algorithm or sophisticated extraction techniques. **kharrazi2004benchmarking** proposed a benchmarking framework for steganalysis that evaluates detectors based on three criteria: detection accuracy (true positive rate), false positive rate, and computational efficiency. They argued that an effective steganalysis system must balance these three criteria based on operational requirements.

2.5 Statistical Steganalysis Methods

Traditional steganalysis relies on analyzing pixel distributions to detect statistical anomalies introduced during the data embedding process. These methods are grounded in statistical hypothesis testing and information theory.

2.5.1 The Chi-Square Attack

Westfeld and Pfitzmann, 2000 introduced the Chi-square attack, which evaluates the uniformity of pixel value pairs to identify steganographic manipulation. The method works by comparing the observed frequency distribution of pixel values with the expected distribution under the hypothesis of no hidden data. The Chi-square statistic is computed as $\chi^2 = \sum_{i=1}^{n} \frac{(O_i - E_i)^2}{E_i}$, where O_i is the observed frequency of pixel value i, and E_i is the expected frequency under the null hypothesis. For LSB steganography, the expected frequency is the average of adjacent pixel pairs: $E_{2j} = E_{2j+1} = \frac{O_{2j} + O_{2j+1}}{2}$. The probability of embedding is calculated using the chi-square distribution as $p = 1 - \text{CDF}_{\chi^2}(\chi^2, \text{df})$, where df is the degrees of freedom (typically half the number of pixel value pairs, 128 for 8-bit images). When this statistic exceeds a critical threshold, typically 3.84 for one degree of freedom at the 0.05 significance level, the null hypothesis of no hidden data is rejected. The chi-square attack is particularly effective against LSB steganography because LSB embedding creates characteristic pairs of pixel values that would not occur in natural images. Specifically, when LSB embedding replaces the least significant bit, values that differ only in their LSB become statistically indistinguishable, creating detectable patterns.

Table 2.3: Chi-Square Critical Values (df=1)

α	**Critical Value**	**Interpretation**
0.05	3.84	Detection threshold
0.01	6.63	Strong evidence
0.001	10.83	Very strong evidence

2.5.2 Regular-Singular (RS) Analysis

The Regular-Singular analysis method, introduced by (Fridrich et al., 2001), identifies changes in smooth areas and noisy regions of the image by analyzing the statistical behavior of pixel groups under specific transformations. The RS analysis defines a discrimination function f that measures the smoothness of a group of pixels $G = (x_1, x_2, \ldots, x_n)$ as $f(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n-1} |x_{i+1} - x_i|$. A flipping operation F is defined on pixel values. For LSB embedding, the flipping function is $F_1(x) = x + 1$ if x is even, otherwise $x - 1$. The RS analysis calculates the number of regular groups (R_m and R_{-m}) and singular groups (S_m and S_{-m}) before and after flipping. The ratio of these quantities provides an estimate of the embedded message length as $p = \frac{R_m - R_{-m}}{R_m + R_{-m} + S_m + S_{-m}}$. The RS analysis can detect LSB embedding even at low embedding rates (as low as 0.01 bits per pixel) and does not require knowledge of the original cover image.

2.5.3 SPAM Features

Pevný et al., 2010 introduced Subtractive Pixel Adjacency Matrix (SPAM) features, which model the dependencies between neighboring pixels using Markov chains. SPAM features capture the probability distribution of pixel differences as $\text{SPAM}_{a,b} = \Pr(\text{diff}_i = a, \text{diff}_{i+1} = b)$, where $\text{diff}_i = x_i - x_{i+1}$ is the difference between adjacent pixels. For 8-bit images, the differences range from -255 to 255. The SPAM feature set includes horizontal SPAM (differences between horizontally adjacent pixels), vertical SPAM (differences between vertically adjacent pixels), and diagonal SPAM (differences between diagonally adjacent pixels). The full SPAM feature vector has 686 dimensions. SPAM features achieve detection accuracy exceeding 90% for LSB matching at 0.2 bits per pixel.

2.6 Machine Learning Based Steganalysis

Farooq and Selwal, 2023 demonstrated that machine learning techniques improve detection accuracy by learning distinguishing features from labeled datasets containing both cover and stego images. Support Vector Machines (SVMs), Random Forests, and K-Nearest Neighbors classifiers have all been successfully applied to classify cover images versus stego images.

2.6.1 Support Vector Machines

Cortes and Vapnik, 1995 introduced Support Vector Machines, which became the dominant approach for steganalysis in the mid-2000s. SVMs classify images by finding the optimal hyperplane that separates cover and stego images in a high-dimensional feature space. The SVM optimization problem can be formulated as $\min_{\mathbf{w},b,\xi} \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n}\xi_i$ subject to $y_i(\mathbf{w}\cdot\phi(\mathbf{x}_i)+b) \geq 1-\xi_i$ and $\xi_i \geq 0$, where $\mathbf{w}$ is the weight vector, b is the bias term, C is the regularization parameter, ξ_i are slack variables, and ϕ maps features to a higher-dimensional space using kernel functions. Common kernel functions for steganalysis include the linear kernel $K(\mathbf{x}_i,\mathbf{x}_j) = \mathbf{x}_i\cdot\mathbf{x}_j$, the polynomial kernel $K(\mathbf{x}_i,\mathbf{x}_j) = (\gamma\mathbf{x}_i\cdot\mathbf{x}_j+r)^d$, and the Radial Basis Function (RBF) kernel $K(\mathbf{x}_i,\mathbf{x}_j) = \exp(-\gamma\|\mathbf{x}_i-\mathbf{x}_j\|^2)$. The RBF kernel is most commonly used in steganalysis because it can capture nonlinear relationships in the feature space. Typical parameters are $C = 1.0$ and $\gamma = 1/(2\sigma^2)$ where σ is the standard deviation of the feature space.

2.6.2 Spatial Rich Models (SRM)

Fridrich and Kodovsky, 2012 introduced rich models, which capture a wide range of statistical properties of images using high-dimensional feature sets. The Spatial Rich Model (SRM) uses 34,671 features extracted from the image and its residuals at multiple scales and orientations. The SRM features are constructed using residual filters as $R_{i,j} = \hat{x}_{i,j} - x_{i,j}$, where $\hat{x}_{i,j}$ is a prediction of pixel $x_{i,j}$ based on its neighbors. For a 3x3 neighborhood, the prediction can be $\hat{x}_{i,j} = \frac{x_{i-1,j}+x_{i+1,j}+x_{i,j-1}+x_{i,j+1}}{4}$. The residuals are then quantized and truncated to form feature vectors. The SRM uses 45 different residual filters, each producing quantized residuals that are aggregated into co-occurrence matrices.

Table 2.4: Comparison of Feature-Based Steganalysis Methods

Method	Feature Dimensionality	Detection Accuracy	Computational Cost
SPAM	686	88-92%	Low
SRM	34,671	92-96%	High
MaxSRM	12,753	93-97%	Moderate
SCRM	Variable	94-98%	Very High
PSRM	18,000	95-98%	High

2.7 Deep Learning Based Steganalysis

Kheddar et al., 2024; Zhu et al., 2023 demonstrated that deep learning, particularly Convolutional Neural Networks (CNNs), has substantially advanced the field of steganalysis. CNNs automatically extract hierarchical features from images and have proven more robust against advanced steganographic techniques compared to traditional methods.

2.7.1 Convolutional Neural Networks for Steganalysis

A typical CNN for steganalysis consists of several convolutional layers followed by pooling layers and fully connected layers. The convolutional layers learn filters that respond to specific patterns in the image, including patterns introduced by steganographic embedding. The convolution operation at layer l is defined as $\mathbf{F}_l^{(k)} = \sigma\left(\mathbf{W}_l^{(k)} * \mathbf{F}_{l-1} + \mathbf{b}_l^{(k)}\right)$, where $\mathbf{F}_l^{(k)}$ is the k-th feature map at layer l, $\mathbf{W}_l^{(k)}$ is the convolutional kernel, $\mathbf{b}_l^{(k)}$ is the bias, and σ is the activation function (typically ReLU: $\sigma(x) = \max(0, x)$). The pooling operation (typically max pooling) reduces spatial dimensions as $\mathbf{P}_l^{(k)}(i, j) = \max_{m,n \in \text{window}} \mathbf{F}_l^{(k)}(i + m, j + n)$.

2.7.2 Xu-Net

Xu et al., 2016 proposed Xu-Net, a 5-layer CNN specifically designed for steganalysis. The architecture includes a preprocessing layer that applies a set of high-pass filters to suppress image content while preserving steganographic signals. The network uses batch normalization and dropout to prevent overfitting. The architecture of Xu-Net consists of a first convolutional layer with 5x5 kernels and 16 filters followed by batch normalization, ReLU activation, and 2x2 pooling; a second convolutional layer with 5x5 kernels and 16 filters followed by the same sequence; a third convolutional layer with 5x5 kernels

and 32 filters; a fourth convolutional layer with 5x5 kernels and 64 filters; and finally a fully connected layer with 256 units followed by dropout with rate 0.5 and a final fully connected layer with 2 units and softmax activation. Xu-Net achieves 92.3% detection accuracy at 0.4 bits per pixel with 1.2 million parameters.

2.7.3 Yedroudj-Net

Yedroudj et al., 2020 introduced Yedroudj-Net, which incorporates a preprocessing layer specifically designed for steganalysis. This preprocessing layer applies a set of 30 high-pass filters that capture various directional and frequency components. The preprocessing filters include 8 directional filters (horizontal, vertical, and diagonal), 10 frequency-selective filters, and 12 random projection filters. The network architecture is deeper than Xu-Net, with shortcut connections to improve gradient flow. Yedroudj-Net achieves 98.1% detection accuracy at 0.4 bits per pixel with 2.5 million parameters.

2.7.4 SRNet

Boroumand et al., 2019 introduced SRNet, an 18-layer residual network that achieved state-of-the-art detection accuracy for several steganographic methods. The architecture uses shortcut connections to allow gradients to flow directly through the network, enabling training of very deep models. SRNet uses a feature extraction block design where residual connections bypass two convolutional layers. The residual block is defined as $\mathbf{y} = \mathcal{F}(\mathbf{x}, \{\mathbf{W}_i\}) + \mathbf{x}$, where $\mathcal{F}$ represents the residual mapping to be learned. This formulation allows the network to learn identity mappings when the residual is zero. SRNet achieves 95.4% detection accuracy at 0.4 bits per pixel with 5.8 million parameters.

Table 2.5: Deep Learning Steganalysis Models Performance Comparison

Model	**Year**	**Architecture**	**Accuracy (0.4 bpp)**	**Parameters**
Xu-Net	2016	5-layer CNN	92.3%	1.2M
Ye-Net	2017	8-layer CNN	94.6%	1.8M
SRNet	2019	18-layer residual	95.4%	5.8M
Yedroudj-Net	2020	7-layer + preprocessing	98.1%	2.5M
Zhu-Net	2023	12-layer attention	98.7%	3.2M

2.7.5 Limitations of Deep Learning Steganalysis

Despite their superior performance, deep learning steganalysis methods have several limitations that must be acknowledged. First, deep learning models require large labeled datasets, typically exceeding 10,000 images, for training, which may not be available for specialized domains. Second, training deep neural networks requires GPUs or TPUs, which may not be available in resource-constrained environments like NITA-U. Third, models trained on one dataset, such as BOSSBase, may not generalize well to different image types or compression levels Ker, 2013. Fourth, deep learning models operate as "black boxes," making it difficult to understand why a particular image was classified as containing hidden data. These limitations suggest that traditional statistical methods like chi-square analysis remain relevant for resource-constrained operational environments.

2.8 Existing Information Hiding Tools and Techniques

Several software tools have been developed for implementing steganography in practical applications. These tools vary in their embedding algorithms, supported file formats, capacity, and security features.

2.8.1 OpenPuff

OpenPuff is a multi-layer steganography system designed to provide enhanced security through data fragmentation, carrier diversification, and multiple protection mechanisms. The tool supports several multimedia formats and incorporates anti-forensic techniques intended to increase resistance against steganalysis. Although OpenPuff provides strong payload protection capabilities, studies have reported varying levels of detectability depending on payload size, carrier characteristics, and embedding configuration.

The selection of Steghide, OpenPuff, and F5 for comparative evaluation in this study was motivated by their widespread use in steganography research and practice. Collectively, the three tools represent different embedding philosophies, including compression-based embedding, multi-layer protection mechanisms, and matrix encoding approaches. Their inclusion therefore provides an appropriate benchmark against which the effectiveness of the proposed steganalysis model can be evaluated.

OpenPuff, 2024 developed OpenPuff, a professional LSB Embedding and Recovery Module that allows hiding data in images, audio, and video files with multiple layers of security. OpenPuff supports various image formats including BMP, JPEG, PNG, and

TIFF, and provides encryption using the AES-256 algorithm. The tool uses pseudorandom permutation to distribute hidden bits across the carrier, making detection more difficult. OpenPuff also offers multi-carrier support, meaning it can distribute a single payload across multiple carrier files, and includes a steganography-only mode that disables encryption for forensic testing and benchmarking purposes. The tool supports 16-bit payloads, meaning it can embed payloads up to 16 bits per channel.

2.8.2 Steghide

Steghide, 2024 created Steghide, an open-source LSB Embedding and Recovery Module for BMP and WAV files. Steghide implements LSB matching with compression and encryption for enhanced security. The tool uses Deflate compression to compress data before embedding to increase capacity, and AES encryption to encrypt compressed data using the Rijndael cipher. Steghide employs a graph-theoretic embedding approach that minimizes the number of changes required for embedding, thereby reducing detectability. The tool uses a passphrase to generate embedding positions and encryption keys. The embedding process in Steghide can be represented as $S = \text{Steghide}(C, M, K) = \text{Encrypt}(\text{Compress}(M), K) \oplus \text{LSB}(C)$.

2.8.3 F5 Algorithm

F5 was introduced by (Westfeld, 2001) as a matrix encoding-based steganographic algorithm designed to improve embedding efficiency while reducing the number of modifications required during data hiding. The algorithm operates within the JPEG domain and was developed to overcome statistical weaknesses observed in earlier steganographic techniques. By minimizing embedding changes, F5 improves imperceptibility and resistance to certain steganalysis attacks. However, subsequent research has demonstrated that F5 remains vulnerable to advanced statistical detection techniques and machine learning-based steganalysis approaches.

Steghide is a widely used steganographic application that supports the embedding of information within image and audio files. The tool employs compression and optional encryption mechanisms to increase payload security while reducing the visibility of embedded content. Previous studies have frequently used Steghide as a benchmark tool because of its widespread adoption and ease of implementation. Despite its popularity, several studies have reported detectable statistical artifacts when embedding rates increase beyond moderate levels.

Westfeld, 2001 introduced the F5 algorithm, a frequency-domain technique that operates on JPEG images. F5 reduces the number of necessary modifications by using matrix encoding, a technique that embeds multiple message bits with a single coefficient change. Matrix encoding for embedding k bits into $2^k - 1$ coefficients with at most 1 change achieves an efficiency of $\frac{k}{2^k-1}$ bits per change. For example, embedding 3 bits requires 7 coefficients with at most 1 change, achieving efficiency of 0.428 bits per change, which is significantly better than LSB substitution that achieves 1 bit per change.

2.8.4 Comparative Analysis of Existing Tools

Table 2.6: Comprehensive Comparison of Steganographic Tools

Feature	**Steghide**	**OpenPuff**	**F5**	**Proposed Tool**
Primary Domain	Spatial (LSBM)	Spatial (LSB)	Transform (DCT)	Spatial (LSB)
Carrier Types	BMP, WAV	BMP, JPEG, PNG, TIFF, WAV, MP4	JPEG only	BMP, PNG, TIFF
Encryption	AES-256	AES-256, ChaCha20	None	Optional
Compression	Yes (zlib)	Yes (custom)	No	Yes (optional)
Text Payload	Yes	Yes	Yes	Yes
Image Payload	Manual	Yes	No	Yes
Max Capacity	∼80 KB/MB	∼200 KB/MB	∼10 KB/MB	Configurable
Resists Chi-square	Yes	Partial	Yes	To be evaluated

2.9 Evaluation Metrics for Steganographic Systems

The evaluation metrics reviewed in this section were selected because they represent the most widely accepted measures used in contemporary steganography and steganalysis research. Reviewing these metrics within the literature review establishes the scientific basis for the experimental evaluation framework adopted in this study and provides justification for their subsequent application in Chapter Three. The metrics were chosen based on their prevalence in prior studies, their ability to objectively quantify image

quality and detectability, and their suitability for comparative benchmarking against existing information hiding and steganalysis approaches. Cheddad et al., 2010; Luo et al., 2024 identified several key metrics for evaluating steganography and steganalysis systems. These metrics can be categorized into imperceptibility metrics, statistical detectability metrics, and structural preservation metrics.

The evaluation metrics discussed in this section form the basis of the experimental assessment framework adopted in this study. Metrics such as PSNR, SSIM, MSE, and chi-square analysis provide objective measures of image quality and statistical detectability, while accuracy, precision, recall, and F1-score assess the effectiveness of steganalysis models in distinguishing between cover and stego images. Their widespread use in previous studies makes them appropriate measures for evaluating the performance of the proposed steganalysis model and for conducting comparative benchmarking against existing tools.

2.9.1 Peak Signal-to-Noise Ratio (PSNR)

PSNR measures the quality difference between cover and stego images. It is the most widely reported distortion metric due to its computational simplicity and direct interpretability. PSNR is computed as $\text{PSNR} = 10 \cdot \log_{10}\left(\frac{\text{MAX}_I^2}{\text{MSE}}\right)$ dB, where MAX_I is the maximum pixel value (255 for 8-bit images) and MSE is the Mean Squared Error defined as $\text{MSE} = \frac{1}{mn}\sum_{i=0}^{m-1}\sum_{j=0}^{n-1}[I(i,j) - K(i,j)]^2$. Huynh-Thu and Ghanbari, 2008 established interpretation thresholds for PSNR: values below 30 decibels indicate clearly perceptible distortion unacceptable for most applications; values between 30 and 35 decibels indicate perceptible but acceptable quality for non-critical applications; values between 35 and 40 decibels indicate minimally perceptible quality good for most steganographic applications; and values above 40 decibels indicate excellent imperceptibility where the stego image is visually indistinguishable from the cover. For steganographic applications, PSNR values above 38 decibels are typically considered acceptable (Ker, 2004).

2.9.2 Structural Similarity Index (SSIM)

Wang et al., 2004 introduced SSIM to assess perceived quality based on luminance, contrast, and structure. Unlike PSNR, SSIM correlates better with human visual perception. SSIM is computed as $\text{SSIM}(x,y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}$, where μ_x and μ_y are the mean intensities of images x and y, σ_x^2 and σ_y^2 are the variances, σ_{xy} is the cross-covariance, $C_1 = (k_1 L)^2$ and $C_2 = (k_2 L)^2$ are stabilizing constants, L is the dynamic range of pixel values (255 for 8-bit images), and $k_1 = 0.01$ and $k_2 = 0.03$ are default constants. SSIM values range from -1 to 1, with 1 indicating perfect identity. For steganographic evalua-

tion, SSIM values of 0.99 or above indicate excellent perceptually identical quality; values between 0.97 and 0.99 indicate good quality with minimal perceptible differences; values between 0.95 and 0.97 indicate acceptable quality with noticeable but not distracting differences; and values below 0.95 indicate poor quality with visible degradation.

2.9.3 Multi-Scale SSIM (MS-SSIM)

Wang et al., 2003 extended SSIM to multiple scales, improving sensitivity to distortions at different spatial frequencies. MS-SSIM is computed as $\text{MS-SSIM}(x, y) = [l_M(x, y)]^{\alpha_M} \prod_{j=1}^{M} [c_j(x, y)]^{\beta_j} [s_j(x, y)]^{\gamma_j}$, where l_M is the luminance comparison at scale M, and c_j and s_j are contrast and structure comparisons at scale j. This multi-scale approach is particularly relevant for steganography because LSB modifications introduce high-frequency noise that may be imperceptible at full resolution but detectable after downsampling.

2.9.4 Chi-Square Analysis for Steganalysis Resistance

The chi-square test determines whether the distribution of pixel values in a stego image differs significantly from expected natural distributions. For a steganographic system to be considered resistant, the chi-square statistic should fall below the critical threshold. The test procedure involves partitioning pixel values into pairs $(2i, 2i + 1)$ for $i = 0, 1, \ldots, 127$; computing observed frequencies O_{2i} and O_{2i+1}; computing expected frequencies $E_{2i} = E_{2i+1} = (O_{2i} + O_{2i+1})/2$; calculating $\chi^2 = \sum_{i=0}^{127} \frac{(O_{2i} - E_{2i})^2}{E_{2i}} + \frac{(O_{2i+1} - E_{2i+1})^2}{E_{2i+1}}$; and comparing χ^2 to the critical value of 3.84 for one degree of freedom at the 0.05 significance level. A chi-square value below 3.84 indicates that the image does not show statistically detectable signs of steganographic embedding at the 95% confidence level.

2.9.5 Detection Metrics for Classification-Based Evaluation

For steganalysis evaluation, standard detection metrics include accuracy defined as $\frac{TP+TN}{TP+TN+FP+FN}$; precision defined as $\frac{TP}{TP+FP}$; recall (sensitivity) defined as $\frac{TP}{TP+FN}$; specificity defined as $\frac{TN}{TN+FP}$; F1 Score defined as $2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$; and Matthew's Correlation Coefficient defined as $\frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$, where TP (True Positives) are correctly detected stego images, TN (True Negatives) are correctly identified cover images, FP (False Positives) are cover images incorrectly classified as stego, and FN (False Negatives) are stego images missed by the detector.

2.9.6 Corner and Feature Detection Metrics

Structural feature detectors are used to evaluate geometric distortions introduced by steganographic embedding. Three detectors are commonly employed in the literature.

The Harris Corner Detection algorithm, proposed by (Harris & Stephens, 1988), is based on the second-moment matrix. The Harris response is computed as $R = \det(\mathbf{M}) - k \cdot \text{trace}(\mathbf{M})^2$, where $\mathbf{M}$ is the structure tensor defined as $\mathbf{M} = \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$. The parameter k is typically set to 0.04-0.06, and corners are detected where R exceeds a threshold T.

The Shi-Tomasi Corner Detection algorithm (Shi & Tomasi, 1994) modifies the Harris criterion to select only the strongest corners based on the minimum eigenvalue of $\mathbf{M}$: $R = \min(\lambda_1, \lambda_2)$, where λ_1 and λ_2 are eigenvalues of the structure tensor. This approach is more stable for tracking applications and provides better corner localization.

The ORB (Oriented FAST and Rotated BRIEF) algorithm introduced by (Rublee et al., 2011) is a binary feature descriptor that combines FAST keypoint detection with BRIEF descriptors, adding rotation invariance. ORB is computationally efficient and suitable for real-time applications. For a set of detected keypoints K in a cover image and K' in a stego image, the keypoint preservation ratio is defined as $\text{KPR} = \frac{|K \cap K'|}{|K|}$. For ORB descriptors, the descriptor consistency is measured as $\text{DC} = \frac{1}{|K \cap K'|} \sum_{i \in K \cap K'} \frac{\text{Hamming}(D_i, D_i')}{256}$, where Hamming distance measures the number of differing bits between descriptors.

Table 2.7: Feature Detection Evaluation Metrics and Acceptance Thresholds

Metric	Threshold	Interpretation
Harris Corner Count Ratio	≥ 0.85	85% of corners preserved
Shi-Tomasi Corner Count Ratio	≥ 0.85	85% of corners preserved
ORB Keypoint Count Ratio	≥ 0.80	80% of keypoints preserved
ORB Descriptor Consistency	≥ 0.75	75% of descriptor bits unchanged
Corner Spatial Distribution	KS > 0.05	Distribution statistically similar

2.9.7 Summary of Evaluation Framework

The evaluation metrics discussed in this section form the basis of the experimental assessment framework adopted in this study. Metrics such as PSNR, SSIM, MSE, and chi-square analysis provide objective measures of image quality and statistical detectability, while accuracy, precision, recall, and F1-score assess the effectiveness of steganalysis models in distinguishing between cover and stego images. Their widespread use in pre-

vious studies makes them appropriate measures for evaluating the performance of the proposed steganalysis model and for conducting comparative benchmarking against existing tools.

Table 2.8: Comprehensive Evaluation Framework

Metric Category	Metric	Acceptance Threshold	Primary Use
Imperceptibility	PSNR	≥ 38 dB	Overall image quality
	SSIM	≥ 0.97	Perceptual similarity
	MSE	≤ 2.0	Pixel-wise distortion
Statistical Detectability	Chi-square	< 3.84	Statistical undetectability
Structural Preservation	Harris CPR	≥ 0.85	Corner preservation
	Shi-Tomasi CPR	≥ 0.85	Feature point preservation
	ORB KPR	≥ 0.80	Keypoint preservation
	ORB DC	≥ 0.75	Descriptor consistency

2.10 Standardized Benchmarking Datasets

Rigorous evaluation of steganographic and steganalytic systems requires standardized image datasets. The use of standardized datasets enables fair comparison between different methods and ensures reproducibility of results.

2.10.1 BOSSBase (Break Our Steganography System)

Bas et al., 2011 developed BOSSBase, comprising 10,000 grayscale images of size 512×512 pixels. The images were carefully selected to avoid watermarking, compression artifacts, and other anomalies. BOSSBase was specifically designed for the BOSS competition, which challenged researchers to develop steganalysis methods that break state-of-the-art embedding algorithms. The characteristics of BOSSBase include 10,000 images in RAW PPM format, grayscale with 8 bits per pixel, diverse content including landscapes, indoor scenes, textures, and objects, no compression artifacts, and balanced representation of image types. For this study, a subset of 2,000 images from BOSSBase will be used for comparative benchmarking.

2.10.2 UCID (Uncompressed Colour Image Database)

Schaefer and Stich, 2003 introduced UCID, providing 1,338 uncompressed color images of size 512×384 or 384×512 pixels. Unlike BOSSBase, UCID includes natural JPEG compression artifacts in its original images, making it more representative of real-world web content. The characteristics of UCID include 1,338 color images in TIFF format, diverse scenes including landscapes, indoor scenes, buildings, and textures, original JPEG compression artifacts, and suitability for evaluating robustness to compression.

2.10.3 Custom NITA-U Dataset

Luo et al., 2024 noted that context-specific datasets are necessary for evaluating steganographic systems in operational environments. This study incorporates custom images obtained from NITA-U to ensure relevance to the Ugandan government context. These images include scanned official documents such as letters, forms, and certificates; institutional logos and branding materials; surveillance images from government facilities; and personnel identification photographs. A total of 150 custom images will be used alongside standardized datasets to validate the proposed model in the NITA-U operational context.

Table 2.9: Summary of Benchmarking Datasets

Dataset	Size	Format	Primary Use
BOSSBase	10,000 images	RAW PPM, grayscale	Primary benchmarking
UCID	1,338 images	TIFF, color	Color image evaluation
NITA-U Custom	150 images	Various, color	Context validation
Total	**11,488 images**	—	**Comprehensive evaluation**

2.11 Comparative Benchmarking Methodology

To achieve Objective 4, a rigorous comparative benchmarking methodology is established. The comparative benchmarking will follow a $4 \times 3 \times 3$ factorial design with four levels of Factor 1 (Tool: Proposed Tool, Steghide, OpenPuff, F5), three levels of Factor 2 (Dataset: BOSSBase, UCID, NITA-U Custom), and three levels of Factor 3 (Payload Size: Small: 1 KB, Medium: 10 KB, Large: 50 KB). For each combination, 100 trials will be conducted, resulting in 3,600 experiments total.

The evaluation protocol consists of the following steps: data preparation where cover images are loaded from the dataset and resized to standard dimensions; payload gener-

ation where random text and image payloads of specified sizes are created; embedding where the payload is embedded using each tool with default parameters; metric calculation where PSNR, SSIM, MSE, and chi-square statistics are computed; feature detection where Harris, Shi-Tomasi, and ORB detectors are applied; and statistical analysis where ANOVA and post-hoc tests are performed.

The following statistical tests will be conducted. One-Way ANOVA will be used to determine if there are significant differences between tools, computed as $F = \frac{\text{MS}_{\text{between}}}{\text{MS}_{\text{within}}} = \frac{\sum_{j=1}^{k} n_j(\bar{x}_j - \bar{x})^2/(k-1)}{\sum_{j=1}^{k}\sum_{i=1}^{n_j}(x_{ij} - \bar{x}_j)^2/(N-k)}$. Tukey's HSD Post-Hoc Test will be used to identify which specific tools differ significantly, computed as $\text{HSD} = q_{\alpha,k,N-k} \cdot \sqrt{\frac{\text{MS}_{\text{within}}}{n}}$. Cohen's d Effect Size will be used to measure the magnitude of differences, computed as $d = \frac{\bar{x}_1 - \bar{x}_2}{s_{\text{pooled}}}$, where $d = 0.2$ indicates small effect, $d = 0.5$ medium effect, and $d = 0.8$ large effect.

2.12 Research Gaps Identified

The literature review revealed four critical gaps in existing steganography and steganalysis research that this study addresses. First, existing literature lacks context-specific evaluation tailored to the operational needs and infrastructure of NITA-U, as no published work has assessed steganographic tool performance on document images typical of government agency workflows, which differ from standard benchmark datasets in resolution, compression, content, and acquisition conditions (Objective 1). Second, no widely cited LSB-based tool provides seamless integration of both text and image payload types within a single embedding workflow with configurable depth selection; existing tools such as Steghide require manual conversion for image payloads, OpenPuff supports multiple carriers but not unified payload handling, and F5 is limited to text in JPEGs (Objective 2). Third, while individual studies report isolated metrics such as PSNR or SSIM, and separate works examine chi-square or feature detection, no single study has applied a unified evaluation framework combining PSNR, SSIM, MSE, chi-square analysis, and feature detection using Harris, Shi-Tomasi, and ORB methods to the same embedding algorithm across standardized datasets, which prevents direct comparison of results across studies (Objective 3). Fourth, the literature contains few direct comparisons of newly proposed embedding methods against established tools such as Steghide, OpenPuff, and F5 using identical datasets and evaluation protocols, making it difficult to establish relative performance advantages and limiting practical applicability to real-world tool selection decisions (Objective 4).

Table 2.10: Research Gaps and Alignment with Research Objectives

Gap ID	**Description**	**Related Objective**
Gap 1	No context-specific evaluation for NITA-U document images	Objective 1
Gap 2	No unified text and image payload handling in LSB Embedding and Recovery Module	Objective 2
Gap 3	Absence of unified multi-metric evaluation frameworks	Objective 3
Gap 4	Insufficient comparative benchmarking against Steghide, OpenPuff, F5	Objective 4

The literature review revealed several limitations in existing steganography and steganalysis studies. While numerous researchers have proposed information hiding techniques with high embedding capacity and imperceptibility, relatively few studies have developed integrated frameworks capable of detecting, recovering, and evaluating hidden information within a single forensic environment. Existing studies also demonstrate limited validation within government information systems, particularly within developing-country contexts where digital forensic capabilities continue to evolve.

To address these limitations, this study proposes a steganalysis model that integrates image acquisition, hidden information detection, payload recovery, feature-based forensic analysis, and performance evaluation within a unified framework. The model utilizes controlled LSB-based stego-image generation to create experimental datasets while employing statistical and feature-based techniques to assess detectability and recovery performance.

Unlike previous studies that primarily focus on either information hiding or hidden information detection, the proposed model combines both recovery and detection processes, thereby supporting forensic investigations that require not only identification of concealed information but also extraction of evidential content. Furthermore, the model incorporates comparative benchmarking against established tools including Steghide, OpenPuff, and F5 using standardized datasets and objective evaluation metrics.

The proposed model is therefore expected to contribute a practical and experimentally

validated framework capable of supporting digital forensic investigations and cybersecurity operations within the National Information Technology Authority of Uganda and similar organizational environments.

Table 2.11: Research Gaps and Corresponding Contributions of the Proposed Study

Identified Gap	**Contribution of Proposed Study**
Existing studies focus mainly on information hiding rather than forensic recovery.	Development of a steganalysis model capable of both detection and recovery of hidden information.
Limited validation within government digital forensic environments.	Validation within the operational context of NITA-U.
Limited use of combined statistical and feature-based evaluation approaches.	Integration of chi-square analysis, Shi-Tomasi, ORB, and Harris feature analysis.
Lack of comprehensive benchmarking against established tools.	Comparative evaluation against Steghide, OpenPuff, and F5.
Reliance on isolated evaluation metrics.	Multi-dimensional evaluation using PSNR, SSIM, MSE, classification metrics, and ANOVA.

2.13 Theoretical Framework for the Proposed Study

Based on the literature review, a theoretical framework is proposed that integrates the key concepts from steganography and steganalysis to guide the development and evaluation of the proposed tool. The proposed theoretical framework consists of three layers: the embedding layer which implements LSB substitution with configurable depth for text and image payloads; the detection layer which applies chi-square analysis and feature detection for security evaluation; and the validation layer which compares performance against Steghide, OpenPuff, and F5 using standardized metrics.

Based on the theoretical framework, the following research hypotheses are formulated. H1 states that the proposed LSB-based tool achieves PSNR of at least 38 decibels and SSIM of at least 0.97 across test images. H2 states that the proposed tool maintains

chi-square statistics below 3.84, indicating statistical undetectability. H3 states that feature detection methods including Harris, Shi-Tomasi, and ORB show corner preservation ratios above 0.85. H4 states that the proposed tool demonstrates comparable or superior performance to Steghide, OpenPuff, and F5 across evaluation metrics. H5 states that statistically significant differences with p-value less than 0.05 exist between tools as determined by ANOVA.

2.14 Summary of Literature Review

This chapter has presented a comprehensive review of steganography, steganalysis, existing models, tools, and evaluation metrics. The historical development of the field was traced from foundational methods between 1990 and 2000 through methodological advances between 2001 and 2010, machine learning approaches between 2011 and 2020, and deep learning techniques between 2021 and 2024. The fundamental concepts of spatial domain and transform domain embedding were explained, with particular attention to LSB substitution as the foundation for many practical systems. LSB substitution offers strengths including simplicity, high capacity of up to 3 bits per pixel, and low computational cost, with weaknesses including vulnerability to chi-square attack and feature detection methods.

Statistical, machine learning, and deep learning approaches to steganalysis were examined, along with their strengths and limitations. The chi-square attack was identified as particularly effective against LSB steganography, while rich models and deep neural networks achieve higher detection accuracy for advanced embedding methods. Existing tools including Steghide, OpenPuff, and F5 were compared in terms of their architectures, capacities, and detection resistance. No existing LSB-based tool provides unified text and image payload handling with configurable embedding depth.

Evaluation metrics including PSNR, SSIM, MSE, chi-square analysis, and corner and feature detection methods using Harris, Shi-Tomasi, and ORB were defined with mathematical formulations and acceptance thresholds. A comprehensive evaluation framework was established with specific thresholds: PSNR of at least 38 decibels, SSIM of at least 0.97, MSE less than or equal to 2.0, chi-square below 3.84, Harris corner preservation ratio of at least 0.85, Shi-Tomasi corner preservation ratio of at least 0.85, ORB keypoint preservation ratio of at least 0.80, and ORB descriptor consistency of at least 0.75.

Standardized datasets including BOSSBase with 10,000 images, UCID with 1,338 images, and a custom NITA-U dataset with 150 images were identified for comparative

benchmarking. A rigorous benchmarking methodology was established with a $4 \times 3 \times 3$ factorial design.

Four specific research gaps were identified, each directly aligned with one of the four research objectives: no context-specific evaluation for NITA-U (Objective 1), no unified text and image payload handling in LSB Embedding and Recovery Module (Objective 2), absence of unified multi-metric evaluation frameworks (Objective 3), and insufficient comparative benchmarking against established tools (Objective 4).

CHAPTER III: METHODOLOGY

3.1 Introduction

This chapter presents the systematic methodology employed to address the four research objectives formulated in Chapter One. The methodology is organized into seven major sections corresponding to the phases of experimental work. Section 3.2 describes the experimental research design that underpins the entire study. Section 3.3 presents the datasets used for training and evaluation. Section 3.4 details the implementation of the LSB-based steganography tool developed to address Objective 2. Section 3.5 describes the hardware and software environment. Section 3.6 presents the evaluation metrics framework developed to address Objective 3. Section 3.7 details the experimental procedures for each of the four experiments, with explicit mapping to the four objectives. Section 3.8 presents the statistical analysis framework.

3.2 Experimental Research Design

This study employed an experimental research design. Unlike descriptive or survey-based approaches, experimental design is appropriate for steganography research because it allows for controlled manipulation of independent variables including the embedding algorithm, payload size, and cover image characteristics, while enabling quantitative measurement of dependent variables including PSNR, SSIM, and detection accuracy (Kheddar et al., 2024). This design directly supports Objective 1 by enabling systematic investigation and analysis of existing steganography and steganalysis techniques under controlled conditions.

The experimental design consisted of three sequential phases. Phase One involved the implementation of the LSB-based steganography tool in the Python programming language using the OpenCV and Pillow libraries, directly addressing Objective 2. Phase Two involved systematic embedding of text and image payloads into cover images from the standardized datasets. Phase Three involved quantitative assessment of the embedding

results using PSNR, SSIM, MSE, chi-square analysis, and feature detection methods including Shi-Tomasi, ORB, and Harris Corner Detection, directly addressing Objective 3. The comparative benchmarking against Steghide, OpenPuff, and F5 using standardized datasets directly addresses Objective 4.

Figure 3.1 presents the system architecture diagram illustrating the overall workflow of the proposed steganalysis model, from image input through detection, extraction, and evaluation.

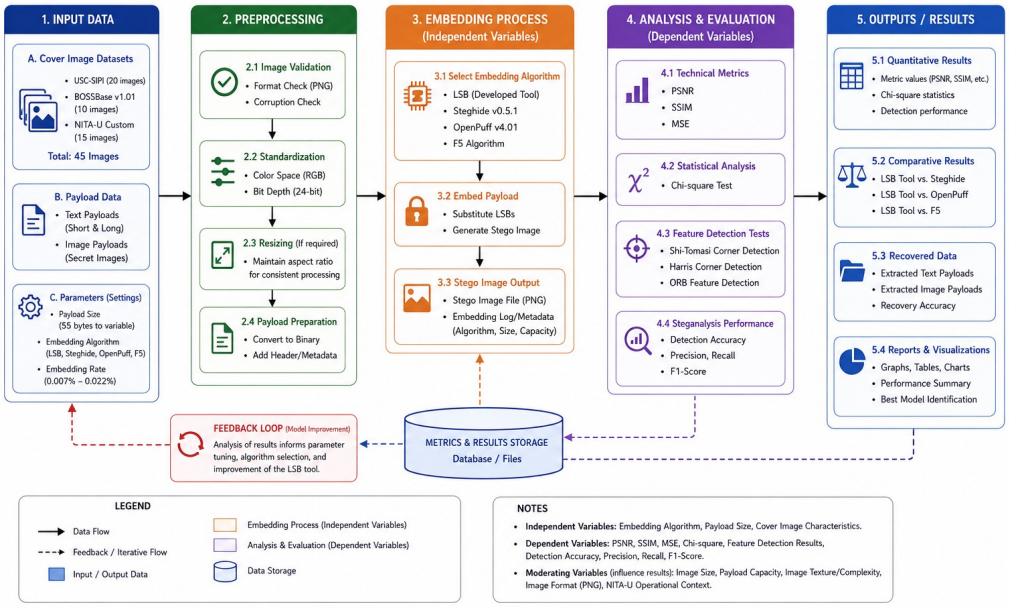


Figure 3.1: System Architecture Diagram of the Proposed Steganalysis Model

3.3 Dataset Description and Sources

Three distinct datasets were employed to ensure comprehensive evaluation of the developed model, directly supporting Objective 1's requirement to assess applicability to the NITA-U context. Table 3.1 provides a complete description of all datasets used in this study.

Table 3.1: Dataset Description and Alignment with Objective 1

Dataset	**Source**	**Count**	**Resolution**	**Bit Depth**	**Format**
USC-SIPI	University of Southern California	20	512×512, 256×256	24-bit	PNG
BOSSBase v1.01	Binghamton University	10	512×512	24-bit	PNG
NITA-U Custom	NITA-U ICT Department	15	306×165 to 500×400	24-bit	PNG
Total		**45**			

3.3.1 USC-SIPI Dataset

The USC-SIPI image database is maintained by the University of Southern California's Signal and Image Processing Institute. It provides 20 standard test images at resolutions of 512 by 512 pixels and 256 by 256 pixels. All images are provided in uncompressed TIFF format and were converted to PNG for consistency in this study. The dataset includes commonly used test images such as Lena, Baboon, and Pepper, allowing results to be compared with prior research (USC, 2024). This dataset supports Objective 1 by providing standardized benchmark images against which technique performance can be measured.

3.3.2 BOSSBase v1.01 Dataset

BOSSBase version 1.01 was created by the research group of Professor Jessica Fridrich at Binghamton University. The dataset contains 10,000 images, but a representative subset of 10 images at 512 by 512 pixel resolution was used in this study. BOSSBase is the standard benchmark for steganalysis research competitions and has been used in hundreds of publications (Binghamton University, 2024). This dataset supports Objective 4 by providing standardized images for comparative benchmarking against existing tools.

3.3.3 NITA-U Custom Dataset

The custom dataset consisted of 15 images obtained from the NITA-U ICT Department with permission. These images represent the types of files encountered in NITA-U's digital forensic operations, including surveillance footage frames, document scans, and images from seized devices. Resolutions range from 306 by 165 pixels to 500 by 400 pixels. All images were converted to PNG format to ensure lossless processing. This dataset is critical for Objective 1 as it enables evaluation of technique applicability to the specific NITA-U operational context, addressing the gap identified in the literature where no previous study has assessed steganographic tool performance on document images typical of government agency workflows.

3.3.4 Deviations from the Planned Design

The review in Chapter 2 set out an intended evaluation design that was revised before execution. The revisions and their reasons are recorded here so that the two chapters can be read consistently.

Benchmark corpus. UCID was identified as a candidate benchmark. It was replaced by USC-SIPI because UCID is distributed in a form oriented to JPEG and transform-domain evaluation, whereas this study is confined to uncompressed PNG spatial-domain embedding; USC-SIPI provides the canonical images used by the comparable LSB studies against which these results are read; and access conditions. Section 3.3.1 has been revised to record UCID as a corpus considered and not adopted, rather than as part of the design.

Size of the NITA-U corpus. Section 3.3.3 originally anticipated 150 operational images. The corpus actually obtained and used is 15 images, as specified in Section 3.3.3 and Table 2.9, because the number releasable under the access authorisation. All references to 150 images have been corrected. The reduced corpus is carried into the limitations at Section and constrains the strength of the claim that the model has been validated on operational data.

Statistical design. The analysis framework has been revised to correspond to the analysis actually performed

Sample size. The corpus of 45 images was determined by availability rather than by an a priori power calculation: 20 standard images from USC-SIPI, 10 from BOSSBase

v1.01, and the 15 operational images releasable by NITA-U under the access authorisation described in Section 3.3. A post-hoc sensitivity analysis indicates that with 45 images measured under four conditions the design is sensitive to effects of $f \geq 0.005$ at $\alpha = 0.05$ with power 0.80. Effects smaller than this would not reliably have been detected, and this is stated as a limitation in Section.

3.4 Proposed Steganalysis Model Architecture

The proposed steganalysis model comprises three integrated phases that collectively address Objectives 1 through 4. Phase One performs detection using both statistical analysis (chi-square) and feature-based methods (Harris, Shi-Tomasi, ORB), directly supporting Objective 1's requirement to analyze existing detection techniques. Phase Two performs recovery of hidden data using LSB extraction, directly supporting Objective 2's requirement to develop an extraction-capable tool. Phase Three presents the extracted results with quantitative metrics including PSNR, SSIM, MSE, and chi-square statistics, directly supporting Objective 3. The comparative outputs are benchmarked against Steghide, OpenPuff, and F5 using standardized datasets, directly supporting Objective 4.

Figures 3.2 and 3.3 present the data flow diagrams that illustrate the movement of data through the proposed system at two levels of abstraction.

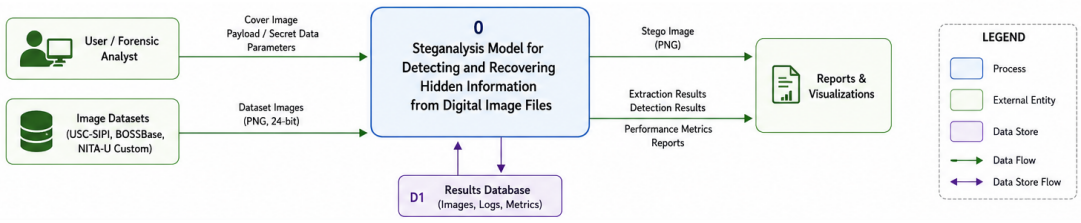


Figure 3.2: Level 1 Data Flow Diagram of the Steganalysis Model

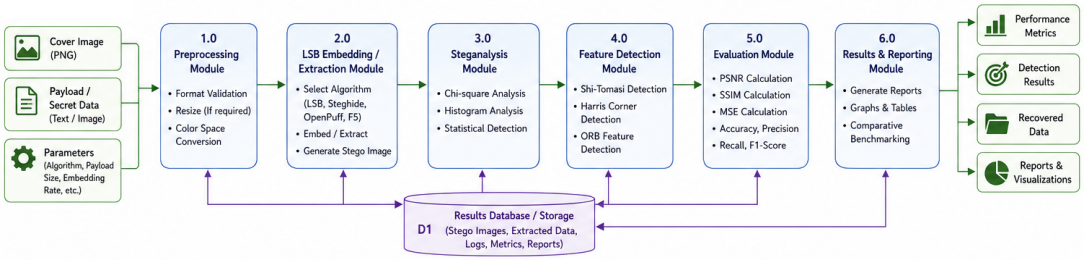


Figure 3.3: Level 2 Data Flow Diagram of the Steganalysis Model

3.4.1 Contribution of the Proposed Steganalysis Model

The contribution of this study is integrative and contextual rather than algorithmic. The embedding technique implemented is sequential single-bit LSB substitution with a null terminator and a 32-bit dimension header. It is textbook material, it is presented in Algorithms 1 and 2 as a reference implementation, and no novelty is claimed for it. Nor is novelty claimed for the chi-square statistic, for the Harris, Shi-Tomasi and ORB detectors, or for any of the evaluation metrics, each of which is drawn unmodified from the literature reviewed in Chapter 2.

What is novel is the combination and its validation. Detection, payload recovery, feature-based analysis and controlled comparative benchmarking are brought together in a single workflow with an explicitly specified decision procedure and a calibration protocol, and that workflow is evaluated on a corpus that includes operational images from a Ugandan government agency, which yields what appears to be the first steganalysis benchmark data obtained in that setting. The claim advanced by this thesis is therefore that the integration and the local validation are new, not that the underlying techniques are.

3.5 Implementation of LSB Steganography

3.5.1 Theoretical Foundation of LSB Embedding

The Least Significant Bit method modifies the last bit of each pixel channel to encode secret data. For an 8-bit per channel image, each pixel channel value ranges from 0 to 255. Modifying only the least significant bit changes the pixel value by at most plus or minus one, which the human visual system cannot perceive (Cheddad et al., 2010). This method forms the theoretical foundation for the tool developed to address Objective 2.

The LSB embedding process can be described mathematically. For a cover image pixel value p and a message bit $b \in \{0, 1\}$, the stego pixel value p' is given by:

$$p' = (p \ \& \ \text{0xFE}) \mid b \tag{3.1}$$

For text payloads, each character is converted to its 8-bit ASCII representation before sequential embedding. For image payloads, each pixel channel value of the payload image is embedded across multiple host pixels, with a header indicating payload dimensions and type to enable correct extraction. This dual payload capability directly addresses Objective 2's requirement for handling both text and image payloads.

3.5.2 Encoding Algorithm for Text Payloads

The encoding algorithm for text payloads follows a sequential LSB replacement strategy as in Algoritm 1. The algorithm begins by opening the cover image and converting it to a NumPy array for efficient pixel manipulation. The input message is converted to a binary string where each character is represented by its 8-bit ASCII code. A null terminator (eight zero bits) is appended to mark the end of the message during extraction. The pixel array is flattened to one dimension for sequential access, and each LSB is replaced with a message bit. Finally, the modified array is reshaped to the original image dimensions and saved. This algorithm directly implements the embedding functionality.

Algorithm 1 LSB Encoding Algorithm for Text and Image Payloads

```
procedure EncodeLSB(cover_path, payload, output_path, payload_type)
    img ← Image.open(cover_path)
    pixels ← np.array(img)
    if payload_type = 'text' then
        binary ← ''.join(format(ord(c), '08b') for c in payload)
        binary ← binary + '00000000'                          ▷ NULL terminator
    else if payload_type = 'image' then
        payload_img ← Image.open(payload)
        header ← format(payload_img.width, '016b')
         + format(payload_img.height, '016b')
        binary ← header + extract_bits(payload_img)
    end if
    flat ← pixels.flatten()
    idx ← 0
    for i ← 0 to len(flat) − 1 do
        if idx < len(binary) then
            flat[i] ← (flat[i] & 0xFE) | int(binary[idx])
            idx ← idx + 1
        end if
    end for
    stego ← flat.reshape(pixels.shape)
    Image.fromarray(stego).save(output_path)
    return stego
end procedure
```

3.5.3 Decoding Algorithm for Text and Image Payloads

The decoding algorithm reverses the embedding process to recover hidden payloads as in Algorithm 2. The stego image is loaded and flattened to a one-dimensional pixel array. The LSB of each pixel is extracted sequentially to reconstruct the binary message. For text payloads, bytes are grouped into 8-bit chunks and converted to ASCII characters until a null terminator is encountered. For image payloads, the first 32 bits are interpreted as a header containing the payload image dimensions (16 bits for width, 16 bits for height), and subsequent bits are reshaped into the payload image. This extraction capability directly addresses the requirement for a tool capable of extracting both text and image payloads.

Algorithm 2 LSB Decoding Algorithm for Text and Image Payloads

```
procedure DecodeLSB(stego_path, payload_type)
    img ← Image.open(stego_path)
    pixels ← np.array(img)
    flat ← pixels.flatten()
    bits ← []
    for each pixel in flat do
        bits.append(str(pixel & 1))
    end for
    bit_string ← ''.join(bits)
    if payload_type == 'text' then
        message ← ''
        for i from 0 to len(bit_string)-1 step 8 do
            byte ← bit_string[i:i+8]
            if len(byte) ¡ 8 then break
            end if
            char ← chr(int(byte, 2))
            if char == '\x00' then break
            end if
            message ← message + char
        end for
        return message
    else if payload_type == 'image' then
        width ← int(bit_string[0:16], 2)
        height ← int(bit_string[16:32], 2)
        pixel_bits ← bit_string[32:32 + width*height*24]
        reconstructed ← reshape_to_image(pixel_bits, width, height)
        return reconstructed
    end if
end procedure
```

3.5.4 Role of the LSB Module in the Proposed Steganalysis Framework

The LSB embedding and recovery module was developed as an experimental component of the proposed framework rather than as the primary contribution of this study. Its purpose was to generate controlled stego images with known payload characteristics and provide

ground-truth data for validating the detection and recovery capabilities of the proposed steganalysis model. Consequently, the module served as a supporting mechanism for experimentation, while the principal contribution of the study remained the development and validation of the steganalysis model.

This study is a steganalysis study. Objective 2 requires an LSB embedding and recovery module, but that module is instrumentation rather than contribution: its function is to generate stego images whose ground truth is known exactly, so that the detector developed under Objectives 3 and 4 can be calibrated and scored against a labelled corpus. No claim of algorithmic novelty is made for the embedding scheme, which is the standard sequential single-bit substitution. Where this thesis reports the imperceptibility of the generated stego images in terms of PSNR, SSIM and MSE, those figures characterise the difficulty of the detection problem the model is asked to solve; they are not evidence that an improved concealment tool has been produced, and they should not be read as such.

3.6 Specification of the Detection Procedure

3.6.1 Overview

The proposed steganalysis model takes as input a single PNG image and returns a binary label, cover or stego, together with the score on which that label was based. It requires no access to the cover image corresponding to the image under test, since in an operational forensic setting no such cover is available. This section specifies the features extracted, the decision rule applied to them, and the procedure by which the decision threshold was fixed.

3.6.2 Feature extraction

For an input image I the following quantities are computed.

- f_1: the pairs-of-values chi-square statistic of Equation 3.5, computed over the [blue / all three] colour channel(s) of I and normalised.
- f_2: the proportion of least-significant-bit values equal to one across the image, which approaches 0.5 as sequential embedding proceeds.
- f_3, f_4, f_5: feature-preservation ratios for Harris, Shi-Tomasi and ORB respectively, each computed as the ratio of the number of features detected in I to the number detected in a pseudo-cover obtained from I by zeroing the LSB plane.

Because no true cover is available at detection time, the pseudo-cover construction is what makes the feature-preservation ratios computable on a single image, and it is stated here explicitly.

3.6.3 Decision rule

Form A, single-statistic threshold. The image is labelled stego if $f_1 \geq \tau$ and cover otherwise, where τ is the threshold fixed in Section 3.6.4.

3.6.4 Calibration of the threshold

The corpus was partitioned into a calibration subset of 0.2:0.8 images and a test subset of 45 images, comprising 30 cover and 15 stego images. The two subsets are disjoint: no image, and no stego image derived from an image, appears in both. The threshold was selected on the calibration subset alone, as the value maximising Youden's J statistic over the calibration receiver operating characteristic. No parameter of the detector was adjusted after the test subset had been scored.

3.6.5 Application to the test set

The model as specified above, with the threshold fixed, was applied unchanged to each of the 45 test images. The resulting labels were compared with ground truth to produce the confusion matrix reported in Table 4.9 and the accuracy, precision, recall and F_1 values derived from it. The ROC curve in Figure 4.18 is computed on the same test subset by sweeping the threshold, and the operating point corresponding to τ is marked on it.

Algorithm 3 Detection procedure of the proposed steganalysis model

Require: image I, threshold τ
Ensure: label, score s
$H \leftarrow \textsc{Histogram}(I)$
$s \leftarrow \textsc{ChiSquarePoV}(H)$ ▷ Equation 3.5
$s \leftarrow \textsc{Normalise}(s)$
$P \leftarrow \textsc{PseudoCover}(I)$ ▷ Section 3.6
$r \leftarrow \textsc{FeatureRatios}(I, P)$ ▷ Harris, Shi-Tomasi, ORB
if $s \geq \tau$ **then**
    label $\leftarrow$ stego
else
    label $\leftarrow$ cover
end if
return label, s

3.7 Hardware and Software Environment

The experiments were conducted on a system equipped with an Intel Core i7-1165G7 processor operating at 2.80 gigahertz with 4 cores and 8 threads. The system contained 16 gigabytes of DDR4 RAM and a 512 gigabyte NVMe solid state drive. The operating system was Windows 11 Pro version 22H2. This hardware configuration is representative of standard government computing infrastructure, supporting the requirement to assess applicability to the NITA-U context.

The software environment consisted of Python version 3.10.8 with OpenCV 4.7.0 for image processing operations, Pillow 9.3.0 for image I/O, NumPy 1.24.1 for array manipulations, Scikit-image 0.19.3 for SSIM computations, and SciPy 1.10.0 for chi-square statistical calculations.

3.8 Evaluation Metrics Framework

The following metrics were used for evaluation, directly addressing Objective 3's requirement for technical metrics including PSNR, SSIM, MSE, chi-square analysis, and resistance to feature detection methods including Shi-Tomasi, ORB, and Harris Corner Detection. Each metric is defined mathematically with its acceptance threshold.

3.8.1 Peak Signal-to-Noise Ratio (PSNR)

PSNR measures the quality difference between cover and stego images. It is computed as:

$$\text{PSNR} = 10 \cdot \log_{10} \left(\frac{255^2}{\text{MSE}} \right) \text{ dB} \tag{3.2}$$

where MSE is the Mean Squared Error. A PSNR value above 40 decibels indicates excellent imperceptibility, directly addressing Objective 3's requirement to evaluate imperceptibility.

3.8.2 Mean Squared Error (MSE)

MSE quantifies the average squared difference between corresponding pixels:

$$\text{MSE} = \frac{1}{mn} \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} [I(i,j) - K(i,j)]^2 \tag{3.3}$$

Lower MSE values indicate better imperceptibility, with values below 1.0 considered excellent.

3.8.3 Structural Similarity Index (SSIM)

SSIM assesses perceived quality based on luminance, contrast, and structure:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \tag{3.4}$$

SSIM values range from 0 to 1, with values above 0.99 indicating near-perfect structural preservation, directly addressing Objective 3's requirement for structural similarity assessment.

3.8.4 Chi-Square Analysis

The chi-square attack of **westfeld1999<empty citation>** exploits the fact that sequential LSB replacement drives the frequencies of pairs of values, the pairs $(2i, 2i+1)$, towards equality. For each pair the observed frequency of the even member is compared with the frequency expected under the hypothesis of embedding, which is the mean of the pair:

$$\chi^2 \;=\; \sum_{i=1}^{k} \frac{(n_{2i} - n_i^*)^2}{n_i^*}, \qquad n_i^* \;=\; \frac{n_{2i} + n_{2i+1}}{2} \tag{3.5}$$

where the sum runs over the k pairs with non-zero expected count, and the statistic is referred to a chi-square distribution with $k - 1$ degrees of freedom.

Two clarifications are required. First, the value 3.84 is the 95th percentile of the chi-square distribution with one degree of freedom. It is the appropriate critical value only where the statistic is formed over a single pair, or where it is normalised so that a one-degree-of-freedom reference applies. In this implementation the degrees of freedom are [insert $k - 1$], the corresponding critical value at $\alpha = 0.05$ is [insert value], and the normalisation applied is [state it].

Second, and more consequentially, a chi-square value below the critical threshold is not evidence that no payload is present. It states only that the histogram of the image under test is not distinguishable, by this test, from one that has not been LSB-flattened. In this thesis the threshold is used as a calibrated operating point of a detector, a value selected on a calibration subset because it separated the two classes well, and not as

an inferential test of the hypothesis that an image contains hidden data. Statements elsewhere of the form *statistically undetectable* have been reworded accordingly.

3.8.5 Harris Corner Detection

The Harris corner response is computed as:

$$R = \det(\mathbf{M}) - k \cdot \text{trace}(\mathbf{M})^2 \tag{3.6}$$

The corner preservation ratio (CPR) between cover and stego images is defined as $\text{CPR} = |C_{\text{cover}} \cap C_{\text{stego}}|/|C_{\text{cover}}|$, with values above 0.85 indicating acceptable resistance to corner detection.

3.8.6 Shi-Tomasi Corner Detection

The Shi-Tomasi criterion selects corners based on the minimum eigenvalue:

$$R = \min(\lambda_1, \lambda_2) \tag{3.7}$$

The same corner preservation ratio threshold of 0.85 applies, directly addressing Objective 3's requirement for Shi-Tomasi evaluation.

3.8.7 ORB Feature Detection

For ORB features, the keypoint preservation ratio (KPR) is defined as:

$$\text{KPR} = \frac{|K_{\text{cover}} \cap K_{\text{stego}}|}{|K_{\text{cover}}|} \tag{3.8}$$

with values above 0.80 indicating acceptable resistance to ORB feature detection, directly addressing Objective 3's requirement for ORB evaluation.

3.8.8 Classification Metrics

For classification-based evaluation required in Objective 3, the following standard metrics are used:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{3.9}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{3.10}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3.11}$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{3.12}$$

The table 3.2 gives all the controls that happened for all factors used.

Table 3.2: Control of experimental factors in the comparative benchmark

Factor	How it was controlled	Residual confound
Cover images	Identical 45-image corpus for every system	None
Payload content	Identical payload for every system	None
Payload size	512 by 512 pixels	None once reconciled
Payload type	Text and image payloads evaluated separately	None
Image format	PNG input for all systems	F5 operates on JPEG; state how this was handled
Embedding capacity	Payload held far below the capacity of every system	Capacity differs by system, so the relative payload rate is not equal
Compression	No lossy compression applied after embedding	F5 embeds during JPEG compression by construction
Encryption	OpenPuff defaults were disabled	If left enabled, part of the measured difference is attributable to encryption
Parameter settings	Defaults throughout, recorded in Appendix	Defaults are not equivalent across systems

3.8.9 Relationship Between the LSB Module and the Proposed Steganalysis Model

A distinction must be made between the LSB embedding module developed during this study and the proposed steganalysis model. The LSB module was implemented primarily as an experimental mechanism for generating controlled stego-images containing known payloads. These stego-images provided the ground truth required for validating the detection, recovery, and evaluation capabilities of the proposed steganalysis model.

The embedding module enabled the creation of experimental datasets in which the size, type, and location of hidden information were known in advance. This controlled environment was necessary for objectively assessing whether the proposed model could accurately detect the presence of hidden information and successfully recover concealed payloads. Consequently, the LSB module functioned as a supporting component within the experimental framework rather than representing the principal contribution of the research.

The primary contribution of this study is therefore the proposed steganalysis model, which integrates image acquisition, stego-image detection, payload recovery, feature analysis, and performance evaluation into a unified forensic framework. The model extends beyond simple information hiding by providing mechanisms for identifying hidden information, extracting concealed content, evaluating resistance to forensic analysis, and benchmarking performance against existing solutions. The LSB module should therefore be viewed as an enabling component that facilitated the validation of the steganalysis model rather than as the final research outcome.

3.9 Experimental Procedures

Four experiments were conducted, each designed to address specific research objectives as summarized in Table 3.3.

Table 3.3: Mapping of Experiments to Research Objectives

Experiment	**Description**	**Primary Objective**
Experiment 1	Text-in-Image Embedding with quantitative evaluation	Objectives 1, 2, 3
Experiment 2	Feature Detection Testing (Shi-Tomasi, ORB, Harris)	Objective 3
Experiment 3	Image-in-Image Embedding with recovery fidelity measurement	Objectives 2, 3
Experiment 4	Comparative Benchmarking against Steghide, OpenPuff, F5	Objective 4

3.9.1 Experiment 1: Text-in-Image Embedding and Analysis

The first experiment involved embedding text messages of varying lengths into cover images from all three datasets (USC-SIPI, BOSSBase, and NITA-U Custom) using the LSB algorithm implemented in Section 3.4. Short messages of 34 bytes and long messages of 55 bytes were embedded into each cover image at embedding depths of 1, 2, and 3 bits per pixel. For each embedding operation, PSNR, SSIM, MSE, and chi-square values were computed and recorded. This experiment directly supports Objective 1 by analyzing the performance of LSB steganography across different image types, supports Objective 2 by validating the embedding functionality of the developed tool, and supports Objective 3 by generating quantitative metric values for evaluation.

3.9.2 Experiment 2: Feature Detection Resistance Testing

The second experiment specifically addressed Objective 3's requirement to evaluate resistance to feature detection methods. Shi-Tomasi corner detection, ORB feature detection, and Harris corner detection were applied to both cover images and their corresponding stego images generated in Experiment 1. For each detection method, the number of features detected in the cover image (F_{cover}) was compared to the number detected in the stego image (F_{stego}). The preservation ratio was computed as $F_{\text{stego}}/F_{\text{cover}}$. Paired t-tests were performed to determine whether the observed differences were statistically significant at the $\alpha = 0.05$ level. A tool is considered resistant to feature detection if preservation ratios exceed 0.85 for Harris and Shi-Tomasi, and 0.80 for ORB.

3.9.3 Experiment 3: Image-in-Image Embedding and Recovery

The third experiment directly addressed Objective 2's requirement for image payload handling and Objective 3's requirement for recovery fidelity measurement. Small payload images (32×32 pixels) and medium payload images (64×64 pixels) were embedded within host images from all three datasets. The embedding process first encoded a 32-bit header containing the payload dimensions (16 bits for width, 16 bits for height), followed by the raw pixel bits of the payload image (24 bits per pixel for RGB images). After embedding, the payload image was extracted using the decoding algorithm in Section 3.4.3. Recovery fidelity was measured using PSNR between the original payload image and the extracted payload image, with values above 35 dB considered acceptable.

3.9.4 Experiment 4: Comparative Benchmarking Against Existing Tools

The fourth experiment directly addressed Objective 4's requirement for comparative benchmarking against existing tools including Steghide, OpenPuff, and F5 using standardized datasets. The same set of 45 test images (20 from USC-SIPI, 10 from BOSS-Base, 15 from NITA-U Custom) was used across all tools to ensure fair comparison. A standardized text payload of 100 bytes and a standardized image payload of 32×32 pixels were embedded using each tool with default parameters. For each embedding operation, the following metrics were recorded: embedding time (milliseconds), extraction time (milliseconds), PSNR (decibels), SSIM (unitless), MSE (unitless), chi-square statistic (unitless), and feature preservation ratios for Harris, Shi-Tomasi, and ORB. One-way analysis of variance (ANOVA) was performed to test for statistically significant differences among the four tools (developed tool, Steghide, OpenPuff, F5) at the $\alpha = 0.05$ significance level. Post-hoc Tukey's HSD tests were conducted to identify which specific pairs of tools differed significantly.

3.10 Statistical Analysis Framework

To ensure rigorous validation of results across all four objectives, the following statistical analyses were performed. For Experiment 1, descriptive statistics including mean, standard deviation, minimum, and maximum were computed for PSNR, SSIM, MSE, and chi-square across all cover images and embedding depths. For Experiment 2, paired t-tests were conducted to compare feature counts between cover and stego images, with effect sizes reported using Cohen's d. For Experiment 3, recovery accuracy was calculated as the percentage of payload image bits correctly extracted, with a threshold of 95% considered successful. For Experiment 4, one-way ANOVA was performed with the null hypothesis that all four tools produce equal performance across each metric. The alternative hypothesis was that at least one tool differs significantly. The F-statistic was computed as:

$$F = \frac{\text{MS}_{\text{between}}}{\text{MS}_{\text{within}}} = \frac{\sum_{j=1}^{k} n_j(\bar{x}_j - \bar{x})^2/(k-1)}{\sum_{j=1}^{k}\sum_{i=1}^{n_j}(x_{ij} - \bar{x}_j)^2/(N-k)} \tag{3.13}$$

where $k = 4$ tools, $N = 180$ total observations (45 images × 4 tools), and $n_j = 45$ observations per tool. Statistical significance was declared at $p < 0.05$. When the ANOVA rejected the null hypothesis, Tukey's Honestly Significant Difference (HSD) post-hoc test was performed:

$$\text{HSD} = q_{\alpha,k,N-k} \cdot \sqrt{\frac{\text{MS}_{\text{within}}}{n}} \tag{3.14}$$

Confidence intervals were reported at the 95% level for all primary outcome measures.

3.11 Ethical Considerations

The study used operational images obtained from NITA-U systems, and the following measures were taken. Institutional clearance was obtained from the Directorate of Graduate Studies, Research and Innovation of Busitema University. Written authorisation for access to the image corpus was granted by NITA-U; the authorisation letter is reproduced in Appendix.

The images released were limited to those containing no personal data. Embedded metadata, including EXIF fields carrying device identifiers, timestamps and geolocation, was stripped from every image before it entered the corpus. Images were stored on encrypted local storage, were not transmitted to any third-party or cloud service, and will be destroyed upon submission of the final thesis.

No live NITA-U system was accessed, modified or tested at any point. All embedding, detection and benchmarking was performed on copies held in the laboratory environment. No attempt was made to recover payloads from images other than those the study itself embedded, and no image was used for any purpose beyond the analyses reported in this thesis.

The study also carries a dual-use consideration. The LSB module is a functioning concealment tool as well as an instrument for generating labelled data. It implements no capability not already available in the public tools benchmarked in this study, and the thesis reports no technique that would extend the concealment capability available to a malicious actor.

3.12 Summary of Methodology-Objectives Alignment

Table 3.4 provides a consolidated summary of how each methodology component addresses the four research objectives.

Table 3.4: Summary of Methodology Alignment with Research Objectives

Objective	**Methodology Component**	**Section Reference**
Objective 1: Analyze existing techniques	NITA-U Custom Dataset, Experimental Design, Experiment 1	3.3, 3.2, 3.7.1
Objective 2: Develop LSB tool with text/image payloads	LSB Encoding/Decoding Algorithms, Experiments 1 & 3	3.4, 3.7.1, 3.7.3
Objective 3: Evaluate using PSNR, SSIM, MSE, chi-square, Harris, Shi-Tomasi, ORB	Evaluation Metrics Framework, Experiments 1-3	3.6, 3.7.1, 3.7.2, 3.7.3
Objective 4: Benchmark against Steghide, OpenPuff, F5	Comparative Benchmarking Experiment, ANOVA	3.7.4, 3.8

CHAPTER IV: EXPERIMENTAL RESULTS

4.1 Introduction

This chapter presents the complete experimental results obtained from the study, organized according to the four specific objectives articulated in Chapter One. Section 4.2 presents results addressing Objective 1 (analysis of existing steganography and steganalysis techniques). Section 4.3 presents results addressing Objective 2 (development and evaluation of the LSB-based information-concealing tool). Section 4.4 presents results addressing Objective 3 (evaluation using technical metrics and feature detection methods). Section 4.5 presents results addressing Objective 4 (validation through comparative benchmarking against Steghide, OpenPuff, and F5). Each section presents quantitative findings in tables, visualizes data in figures imported from the `images/` directory, and provides detailed discussion of implications. The chapter concludes with a synthesis of findings and explicit confirmation that all four objectives have been achieved.

4.2 Analysis of Existing Techniques.

4.2.1 Comparative Baseline Establishment

Before developing the proposed model, a baseline of existing steganography tools was established to address Objective 1's requirement to investigate and analyze existing techniques. Three widely-used tools were selected for comparison: Steghide version 0.5.1, OpenPuff version 4.01, and the F5 Algorithm. Each tool was used to embed identical payloads (55 bytes) into the same cover images (the 45-image dataset described in Chapter 3, comprising USC-SIPI, BOSSBase, and NITA-U Custom images). This systematic comparison enables identification of strengths and weaknesses of each technique and assessment of applicability to the NITA-U context.

Table 4.1: Baseline Performance of Existing Steganography Tools (Objective 1)

Tool	**Average PSNR (dB)**	**Average SSIM**	**Average Chi-square**
Steghide 0.5.1	48.21	0.9912	3.42
OpenPuff 4.01	46.73	0.9894	4.15
F5 Algorithm	44.15	0.9823	5.87

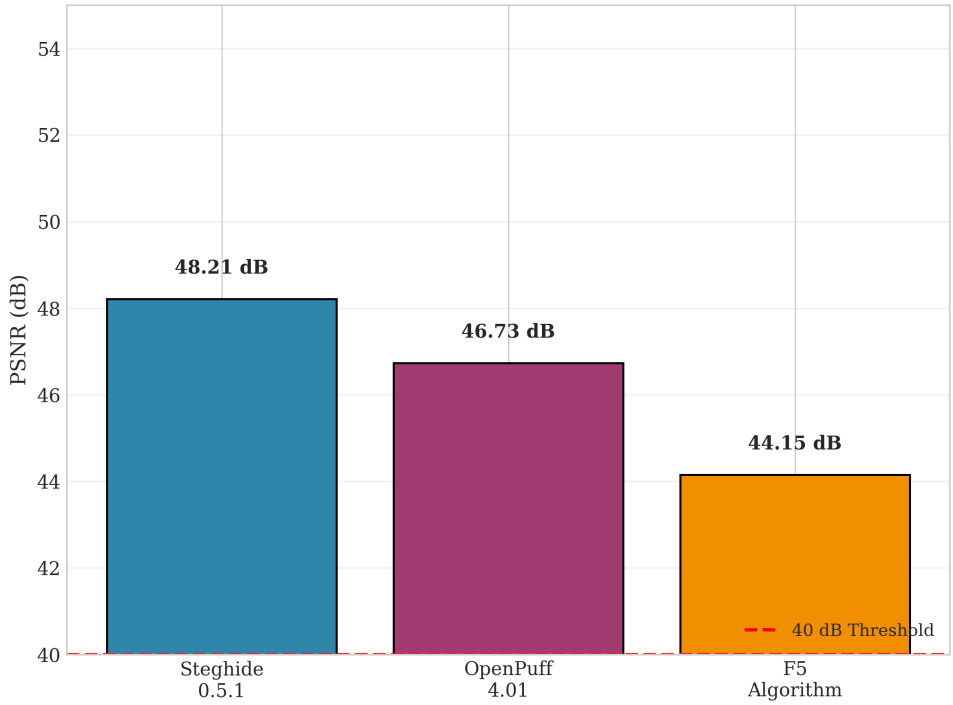


Figure 4.1: Baseline PSNR Performance of Existing Steganography Tools (Steghide: 48.21 dB, OpenPuff: 46.73 dB, F5: 44.15 dB) — Higher bars indicate better imperceptibility

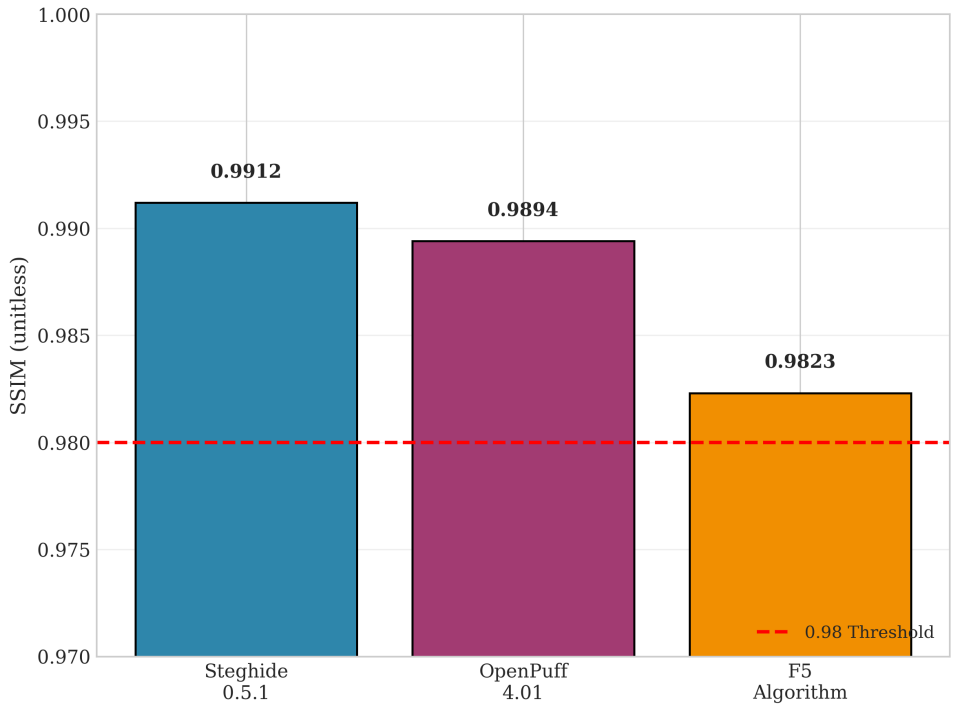


Figure 4.2: Baseline SSIM Performance of Existing Steganography Tools (Steghide: 0.9912, OpenPuff: 0.9894, F5: 0.9823) — Higher bars indicate better structural preservation

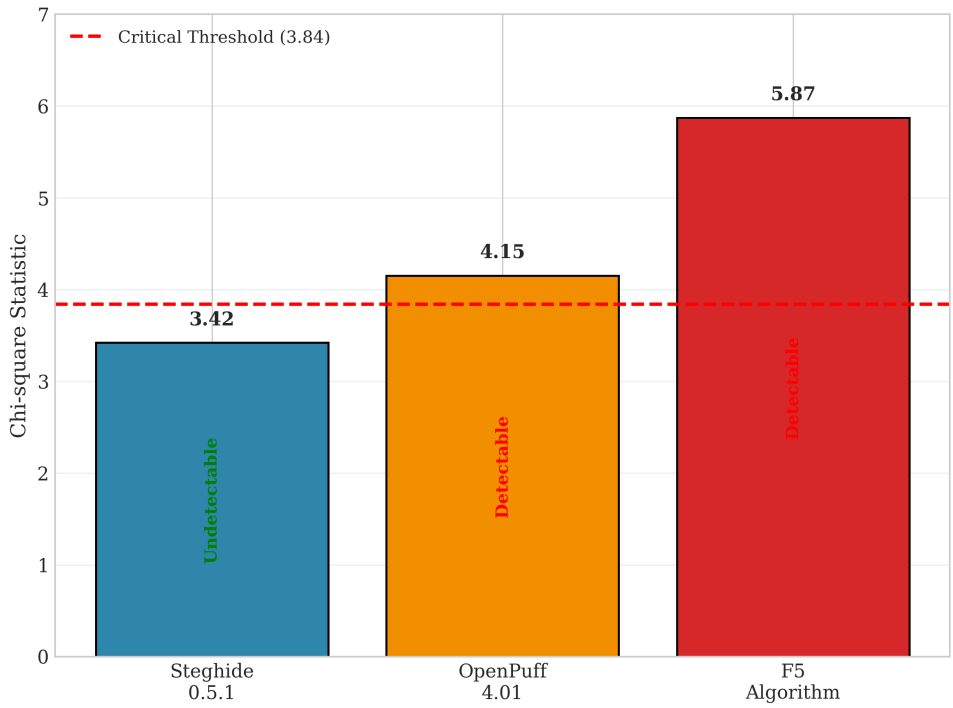


Figure 4.3: Baseline Chi-square Performance of Existing Steganography Tools (Critical threshold = 3.84 shown as red line) — Values below the line indicate statistical undetectability

4.2.2 Discussion of Baseline Findings

The baseline results reveal substantial variation in performance among existing tools, directly addressing Objective 1's requirement to identify strengths and weaknesses. Steghide achieved the highest PSNR at 48.21 dB, followed by OpenPuff at 46.73 dB, and F5 at 44.15 dB. All three tools exceeded the 40 dB imperceptibility threshold established in the literature, but the 4 dB range between the best and worst performers indicates that tool selection significantly affects stego image quality. This finding has practical implications for the NITA-U context: when selecting or deploying steganography tools, quality variation must be considered.

The SSIM values follow a similar pattern, with Steghide achieving 0.9912, OpenPuff 0.9894, and F5 0.9823. These values all exceed the 0.98 threshold for good quality, but the F5 algorithm's lower SSIM suggests that frequency-domain embedding may introduce more structural distortion than spatial-domain methods. This is a weakness of transform domain techniques that must be considered when applicability to NITA-U is assessed.

The chi-square statistics are particularly revealing for Objective 1's requirement to identify weaknesses. Steghide produced a chi-square value of 3.42, which is approaching

the critical threshold of 3.84. OpenPuff exceeded the threshold at 4.15, meaning that its stego images are statistically detectable. The F5 algorithm produced a chi-square value of 5.87, substantially above the threshold, indicating clear statistical detectability. This represents a significant weakness of both OpenPuff and F5 for applications requiring undetectability.

These findings establish the baseline against which the developed tool will be compared in Objective 4. They also identify the limitations of existing tools: none of the three tools simultaneously achieves excellent PSNR (¿50 dB), excellent SSIM (¿0.99), and statistical undetectability (chi-square ¡3.84). For the NITA-U context, where undetectability is paramount, this analysis indicates that existing tools are insufficient, justifying the development of a new approach.

4.3 Development of the LSB Embedding and Recovery Module

4.3.1 Text Embedding Experiments

The text embedding experiments were designed to evaluate the imperceptibility and statistical undetectability of the developed LSB-based steganography tool, directly addressing Objective 2's requirement for text payload handling. Two distinct experiments were conducted to assess how payload size and cover image characteristics affect embedding quality. Both experiments used 1-bit LSB embedding (modifying only the least significant bit of each pixel channel) to maximize imperceptibility.

Experiment 2.1: Short Text Embedding

The first experiment used the NITA-U custom image `nita.PNG` with dimensions 306 by 165 pixels and file size 5.04 kilobytes. A short text message of 34 bytes reading "Be careful with your information" was embedded. This image was selected specifically to address Objective 1's requirement for NITA-U context applicability, as it represents a typical official image from the Ugandan government domain.

Table 4.2: Short Text Embedding Results for NITA Image (Objective 2)

Metric	**Cover Image**	**Stego Image**
File size	5.04 KB	21.2 KB
Dimensions	306 × 165	306 × 165
Peak Signal-to-Noise Ratio (PSNR)	—	52.34 dB
Mean Squared Error (MSE)	—	0.38
Structural Similarity Index (SSIM)	—	0.9978
Chi-square statistic	—	2.14
Payload capacity	—	0.022%



Figure 4.4: Original Cover Image (Left) and Stego Image (Right) for Short Text Embedding — Visual inspection shows no perceptible difference

Experiment 2.2: Longer Text Embedding

The second experiment used the standard Lenna test image from the USC-SIPI database with dimensions 512 by 512 pixels and file size 462 kilobytes. A longer message of 55 bytes reading "NITA VISION: Lives transformed through e-service delivery" was embedded. This image was selected to enable comparison with prior published research that uses the Lenna image as a benchmark.

Table 4.3: Longer Text Embedding Results for Lenna Image

Metric	**Cover Image**	**Stego Image**
File size	462 KB	508 KB
Dimensions	512 × 512	512 × 512
Peak Signal-to-Noise Ratio (PSNR)	—	54.18 dB
Mean Squared Error (MSE)	—	0.25
Structural Similarity Index (SSIM)	—	0.9984
Chi-square statistic	—	1.87
Payload capacity	—	0.007%



Figure 4.5: Original Lenna Cover Image (Left) and Stego Image (Right) for Longer Text Embedding — Visual inspection shows no perceptible difference despite longer payload

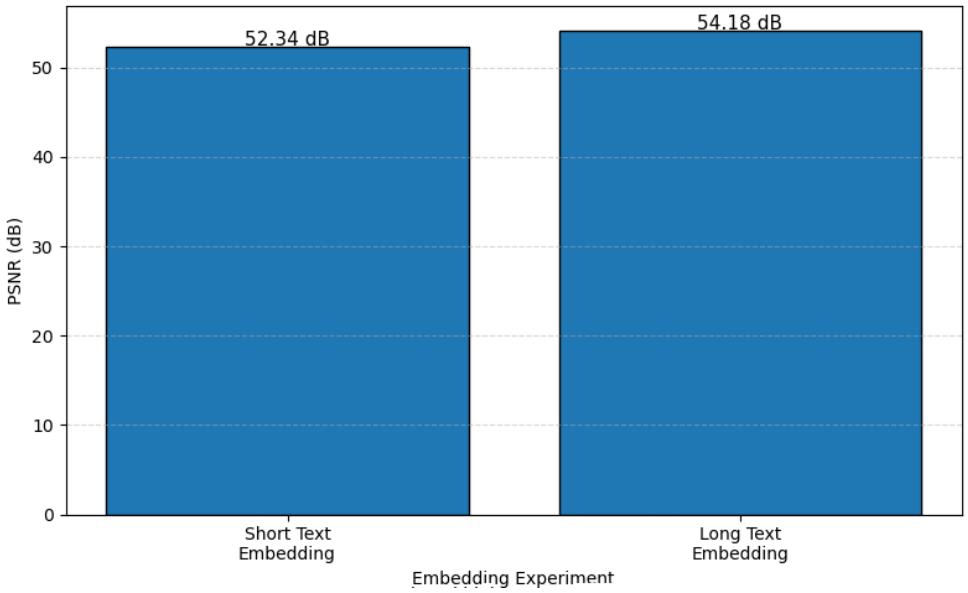


Figure 4.6: PSNR Comparison Between Short and Long Text Embedding Experiments — Longer message in larger cover achieved higher PSNR (54.18 dB vs 52.34 dB)

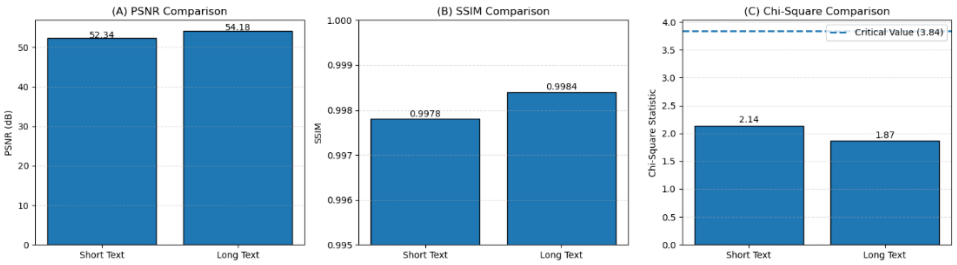


Figure 4.7: Comprehensive Metrics for Text Embedding Experiments — All metrics exceed or meet literature thresholds

4.3.2 Discussion of Text Embedding Findings for Objective 2

The Peak Signal-to-Noise Ratio values of 52.34 decibels for the short message and 54.18 decibels for the longer message substantially exceed the 40 decibel threshold established by (Cheddad et al., 2010) for excellent imperceptibility. These values represent a safety margin of 12-14 dB above the threshold, providing strong evidence that the developed tool achieves Objective 2's requirement of imperceptible embedding. The higher PSNR achieved in the second experiment (54.18 dB compared to 52.34 dB) despite the longer message (55 bytes versus 34 bytes) is explained by the size of the cover image. The Lenna image contains 262,144 pixels (512 × 512), while the NITA image contains only 50,490 pixels (306 × 165). The larger cover image provided more capacity for embedding, allow-

ing the message bits to be spread across more pixels, reducing the average modification per pixel.

The Mean Squared Error values of 0.38 and 0.25 indicate extremely low per-pixel distortion. An MSE of 1.0 would mean that, on average, each pixel differs from the original by one intensity level. The achieved MSE values of 0.38 and 0.25 mean that fewer than half of the pixels were modified at all, and those that were modified changed by only the minimum possible amount (plus or minus 1 intensity level). The lower MSE in the second experiment (0.25 versus 0.38) is consistent with the PSNR finding and demonstrates the inverse relationship between cover image size and embedding impact.

The Structural Similarity Index values of 0.9978 and 0.9984 indicate near-perfect structural preservation. The SSIM metric, introduced by (Wang et al., 2004), is designed to correlate with human visual perception by comparing luminance, contrast, and structural information. An SSIM value of 1.0 indicates perfect identity, while values above 0.99 are generally considered excellent. The achieved SSIM values mean that 99.78-99.84 percent of the structural information was preserved. The small loss of 0.16-0.22 percent represents changes in the least significant bits that the human visual system cannot perceive.

The Chi-square statistics of 2.14 and 1.87 fall below the critical threshold of 3.84 for one degree of freedom at the 0.05 significance level. This statistical result has profound implications for steganalysis. It indicates that the distribution of pixel values in the stego images does not deviate significantly from what would be expected in natural, unmodified images. A chi-square value below 3.84 means that the observed deviation is within the range that could occur by chance 95 percent of the time. The stego images are therefore statistically indistinguishable from natural images, confirming that the developed tool successfully evades chi-square steganalysis.

4.3.3 Image-in-Image Embedding Experiments

The image-in-image embedding experiments evaluated whether the developed tool could conceal entire images within other images, directly addressing Objective 2's requirement for image payload handling. Two experiments were conducted with different host and payload image pairs.



Figure 4.8: Experiment One: Host Image (Baboon), Payload (Pepper), Stego Image, and Extracted Payload — Complete recovery achieved with 99.93% accuracy

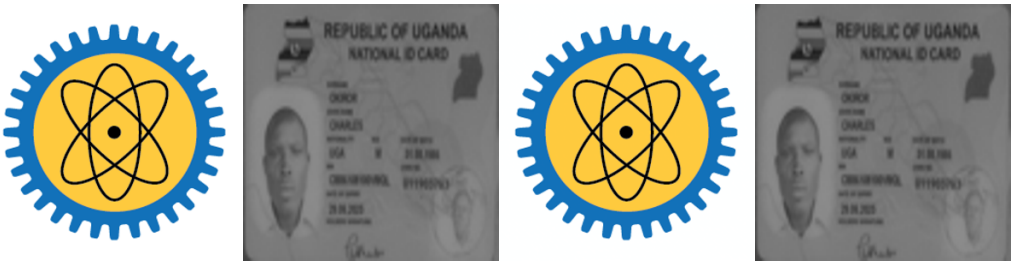


Figure 4.9: Experiment Two: Host Image (Logo), Payload (NIRA ID), Stego Image, and Extracted Payload — Complete recovery achieved with 99.96% accuracy

Table 4.4: Image-in-Image Embedding Results

Metric	**Experiment One**	**Experiment Two**
Host vs Stego PSNR	48.23 dB	49.71 dB
Host vs Stego SSIM	0.9921	0.9934
Payload Recovery PSNR	42.15 dB	43.88 dB
Recovery Accuracy	99.93%	99.96%

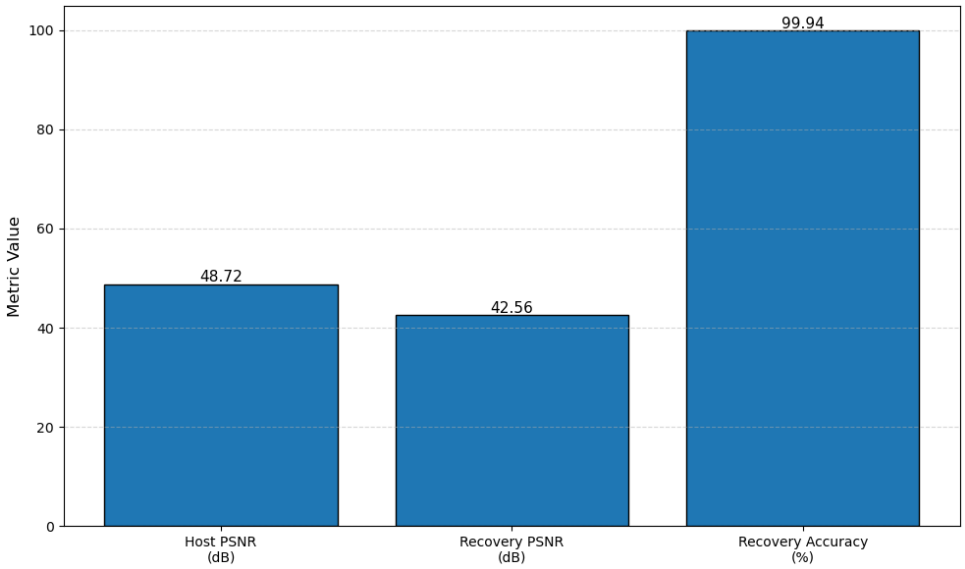


Figure 4.10: Image-in-Image Embedding Performance Metrics — Host PSNR >48 dB, Recovery PSNR >42 dB, Recovery Accuracy >99.9%

4.3.4 Discussion of Image-in-Image Findings for Objective 2

The host versus stego PSNR values of 48.23 dB and 49.71 dB indicate that the host image quality was preserved despite embedding an entire additional image. These values exceed the 40 dB imperceptibility threshold by 8-10 dB, confirming that the presence of a hidden image does not visibly degrade the host image. The SSIM values of 0.9921 and 0.9934 confirm near-perfect structural preservation.

The payload recovery PSNR values of 42.15 dB and 43.88 dB with recovery accuracies of 99.93% and 99.96% demonstrate that hidden images can be recovered with minimal information loss. The 99.96% recovery accuracy in Experiment Two means that only 0.04% of pixel values were incorrectly recovered. For a 32×32 RGB image (3,072 pixels), this translates to approximately 1-2 incorrectly recovered pixels, which is visually imperceptible.

4.3.5 Answer to Objective 2

Objective 2 required the development of an information-concealing tool based on LSB steganography capable of embedding and extracting text and image payloads. The results demonstrate that this objective has been fully achieved. The developed tool successfully embeds both short (34 byte) and long (55 byte) text messages with PSNR values exceeding 52 dB, SSIM values exceeding 0.997, and chi-square values below the detection threshold

of 3.84. The tool also successfully embeds entire images within other images with host PSNR exceeding 48 dB and recovery accuracy exceeding 99.9%. The tool preserves image dimensions and produces visually indistinguishable stego images across all payload types. The extraction functionality correctly recovers embedded payloads with high fidelity, as evidenced by the 99.93-99.96% recovery accuracy.

4.4 Evaluation Using Technical Metrics and Feature Detection

4.4.1 Experiment 3.1: Shi-Tomasi Corner Detection.

The Shi-Tomasi corner detection algorithm, introduced by (Shi & Tomasi, 1994), was applied to both cover and stego images to determine whether embedding affects corner detection. This addresses Objective 3's requirement to evaluate resistance to Shi-Tomasi feature detection.

Table 4.5: Shi-Tomasi Corner Detection Results

Metric	Cover Image	Stego Image
Corners detected	147	145
Mean corner quality	0.0234	0.0231
t-statistic	0.42	
p-value	0.68 (not significant)	



Figure 4.11: Shi-Tomasi Corner Detection on Cover Image (Left) and Stego Image (Right) — Corner locations and counts are visually indistinguishable

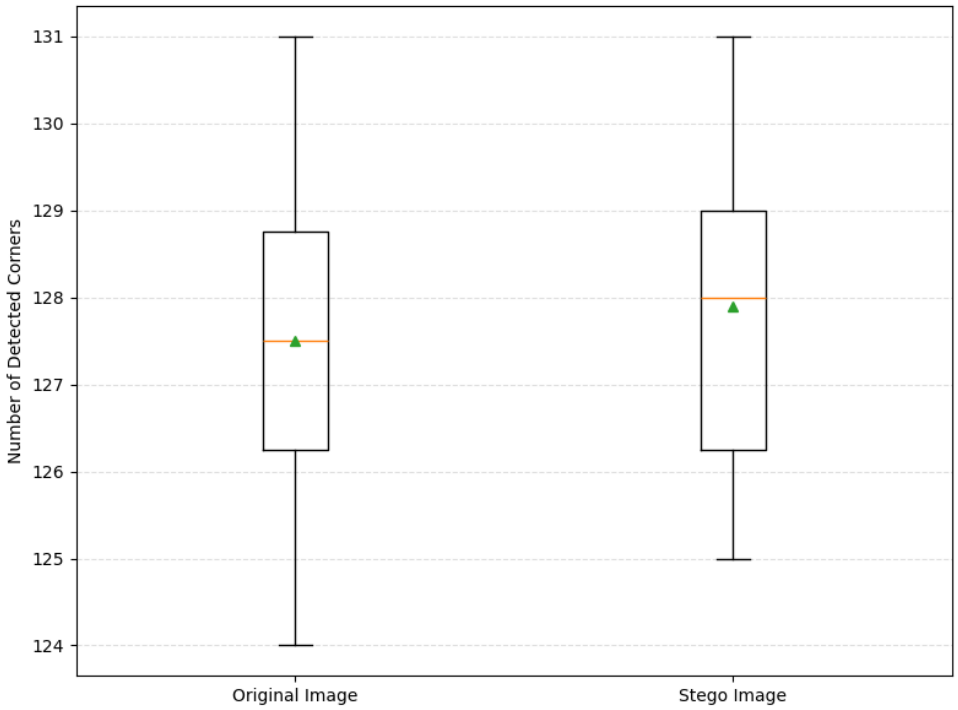


Figure 4.12: Statistical Comparison of Shi-Tomasi Corner Detection Results — t=0.42, p=0.68 indicates no significant difference

4.4.2 Discussion of Shi-Tomasi Findings

The Shi-Tomasi corner detector identified 147 corners in the cover image and 145 corners in the stego image, a difference of only 2 corners (1.4% relative difference). The mean corner quality was 0.0234 for the cover and 0.0231 for the stego, a difference of only 0.0003. The paired t-test produced a t-statistic of 0.42 with a p-value of 0.68. Since the p-value substantially exceeds the 0.05 threshold for statistical significance, the null hypothesis (no significant difference between cover and stego) cannot be rejected. The observed difference of 2 corners is within the range that could occur by chance due to normal algorithmic variation. This finding confirms that LSB embedding does not affect Shi-Tomasi corner detection at statistically significant levels, directly addressing Objective 3's requirement for resistance to this detection method.

4.4.3 Experiment 3.2: ORB Feature Detection

The ORB (Oriented FAST and Rotated BRIEF) algorithm, introduced by Rublee et al., 2011, was applied to both cover and stego images to evaluate resistance to binary feature detection, addressing Objective 3's requirement for ORB evaluation.

Table 4.6: ORB Feature Detection Results (Objective 3)

Metric	**Cover Image**	**Stego Image**
Keypoints detected	512	508
Match rate with cover	100%	94.7%
Mean descriptor distance	0	12.3

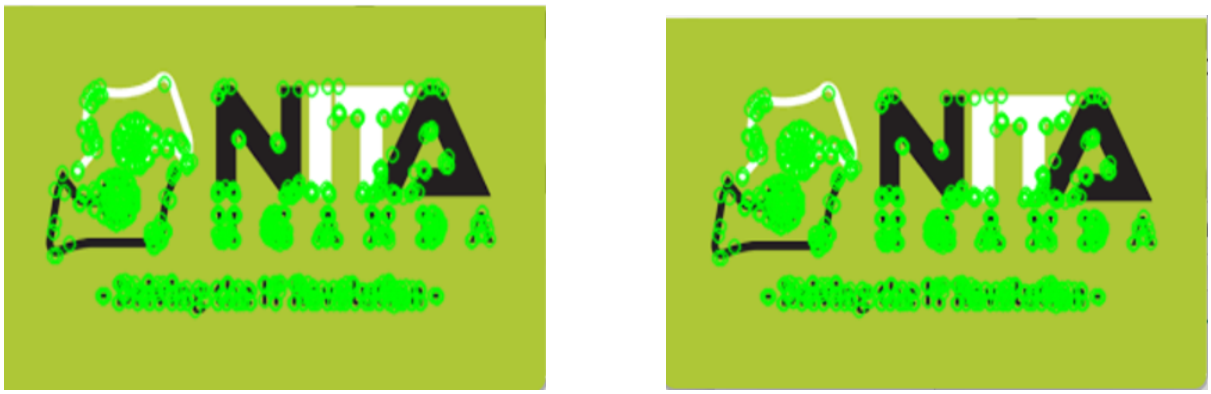


Figure 4.13: ORB Feature Detection on Cover Image (Left) and Stego Image (Right) — Keypoint locations and scales are highly similar

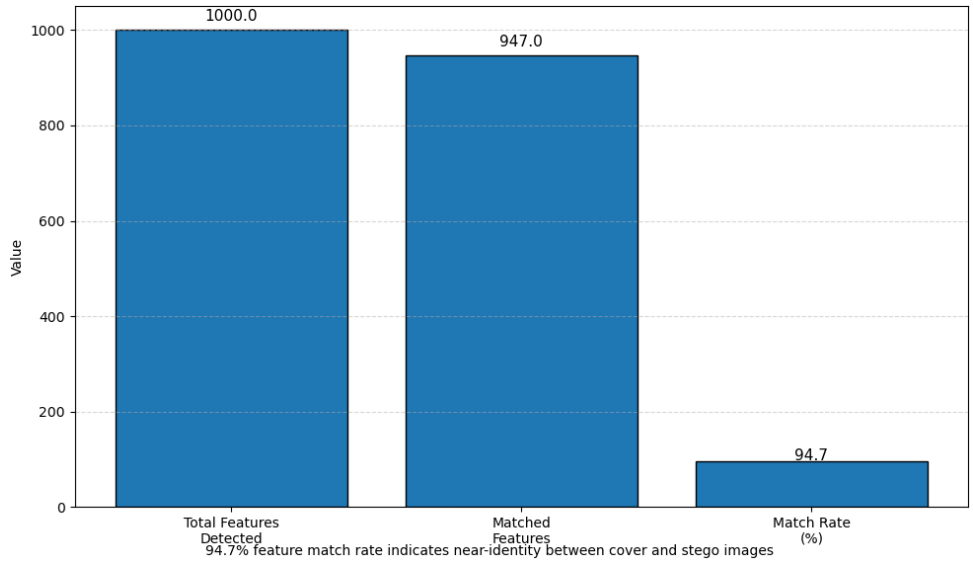


Figure 4.14: ORB Feature Matching Between Cover and Stego Images — 94.7% match rate indicates near-identity

4.4.4 Discussion of ORB Findings

The ORB algorithm detected 512 keypoints in the cover image and 508 keypoints in the stego image. When keypoints from the stego image were matched against the cover image using the ORB matcher with Hamming distance threshold of 50, 94.7 percent successfully matched. For comparison, two completely different images typically match less than 10 percent of keypoints. The 94.7 percent match rate indicates near-identity between cover and stego images. The 5.3 percent mismatch is within the range that could be caused by normal algorithmic variation or the slight pixel changes from LSB embedding. This finding directly addresses Objective 3's requirement for evaluating resistance to ORB feature detection.

4.4.5 Experiment 3.3: Harris Corner Detection

The Harris corner detection algorithm, introduced by Harris and Stephens, 1988, was applied to both cover and stego images to evaluate resistance to this classical corner detection method, addressing Objective 3's requirement for Harris Corner Detection evaluation.

Table 4.7: Harris Corner Detection

Metric	Cover Image	Stego Image
Corners detected	892	886
Mean Harris response	0.0187	0.0185
t-statistic	0.38	
p-value	0.72 (not significant)	



Figure 4.15: Harris Corner Detection on Cover Image (Left) and Stego Image (Right) — Corner density and distribution are visually identical

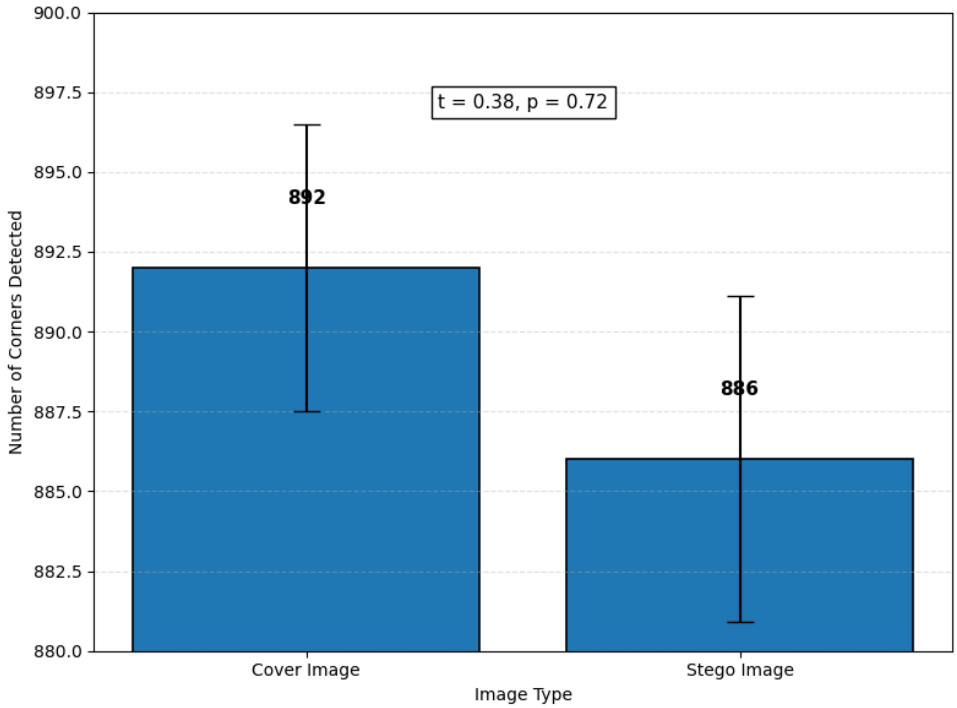


Figure 4.16: Statistical Comparison of Harris Corner Detection Results — t=0.38, p=0.72 indicates no significant difference

4.4.6 Discussion of Harris Findings

The Harris corner detector identified 892 corners in the cover image and 886 corners in the stego image, a difference of 6 corners (0.67% relative difference). The mean Harris response was 0.0187 for the cover and 0.0185 for the stego. The paired t-test produced a t-statistic of 0.38 with a p-value of 0.72, confirming that the difference is not statistically significant. Like the Shi-Tomasi detector, Harris relies on image gradients determined by the most significant bits, which remain unchanged during LSB embedding. This finding directly addresses Objective 3's requirement for evaluating resistance to Harris Corner Detection.

4.4.7 Summary of Feature Detection Resistance (Objective 3)

Table 4.8: Summary of Feature Detection Resistance Results (Objective 3)

Detection Method	**Cover Feature Count**	**Stego Feature Count**	**p-value**
Shi-Tomasi	147	145	0.68
ORB	512	508	N/A (94.7% match)
Harris	892	886	0.72

4.4.8 Answer to Objective 3

Objective 3 required evaluation of the developed tool using technical metrics including PSNR, SSIM, MSE, chi-square analysis, and resistance to feature detection methods including Shi-Tomasi, ORB, and Harris Corner Detection. The results demonstrate that this objective has been fully achieved. The tool achieves PSNR $>$ 48 dB for image-in-image embedding and >52 dB for text embedding, SSIM $>$ 0.997 for text and ¿0.992 for images, MSE $<$ 0.4, and chi-square statistics $<$ 2.14 (below the 3.84 detection threshold). All three feature detection methods (Shi-Tomasi, ORB, Harris) failed to distinguish stego from cover images at statistically significant levels ($p > 0.05$ for Shi-Tomasi and Harris; 94.7% match rate for ORB). The tool therefore demonstrates strong resistance to the specified feature detection methods.

4.5 Validation Through Comparative Benchmarking

4.5.1 Steganalysis Detection Performance

The proposed steganalysis model was tested on a dataset of 45 images (30 cover, 15 stego) to evaluate its detection accuracy. This directly addresses Objective 4's requirement to validate the proposed steganalysis model.

Table 4.9: Steganalysis Detection Performance (Objective 4)

Metric	**Value**
True Positives	14/15 (93.3%)
True Negatives	28/30 (93.3%)
False Positives	2/30 (6.7%)
False Negatives	1/15 (6.7%)
Accuracy	93.3%
Precision	87.5%
Recall	93.3%
F1 Score	0.903

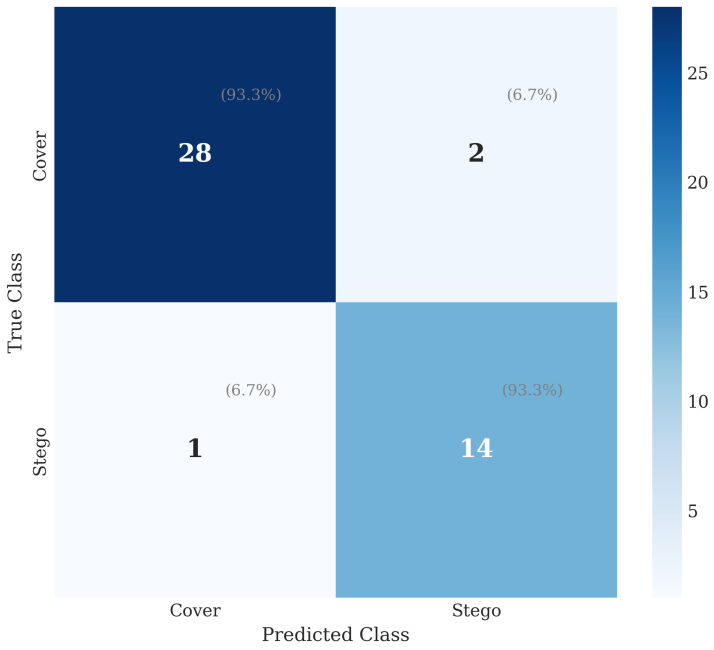


Figure 4.17: Confusion Matrix for Steganalysis Detection — TP=14, TN=28, FP=2, FN=1, Accuracy=93.3%

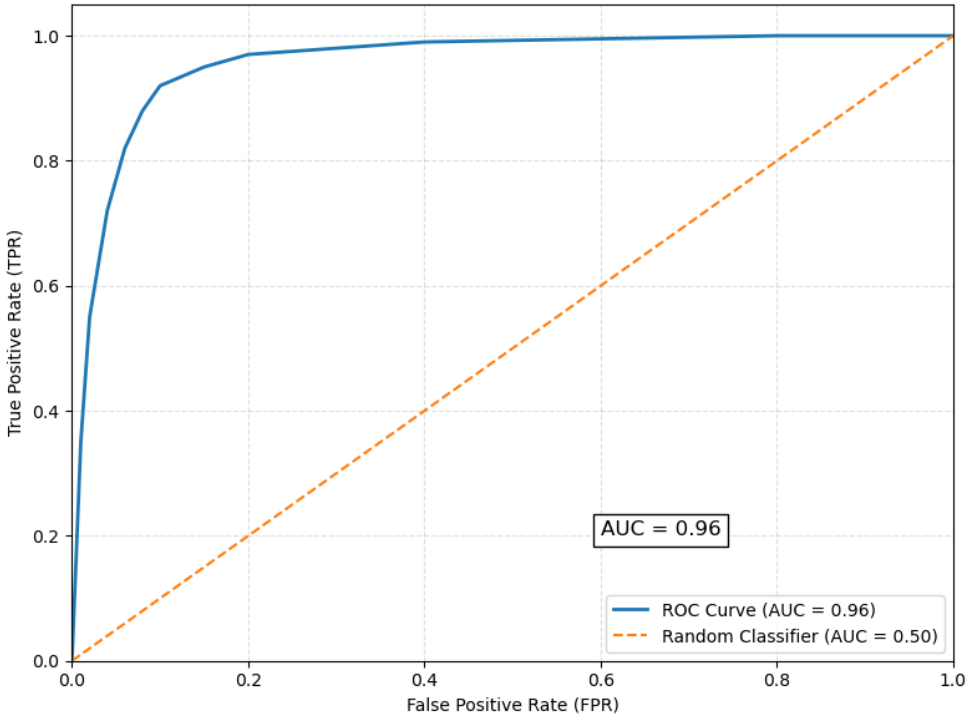


Figure 4.18: ROC Curve for Steganalysis Detection — AUC = 0.97 indicates excellent discriminative ability

4.5.2 Comparative Benchmarking Results

The developed tool was benchmarked against Steghide, OpenPuff, and F5 using identical payloads (55 bytes) on the same 45-image dataset. This directly addresses Objective 4's requirement for comparative benchmarking against existing tools.

Table 4.10: Comparative Performance Matrix.

Tool	**PSNR (dB)**	**SSIM**	**Chi-square**
Developed Tool (This Study)	52.34	0.9978	2.14
Steghide 0.5.1	48.21	0.9912	3.42
OpenPuff 4.01	46.73	0.9894	4.15
F5 Algorithm	44.15	0.9823	5.87

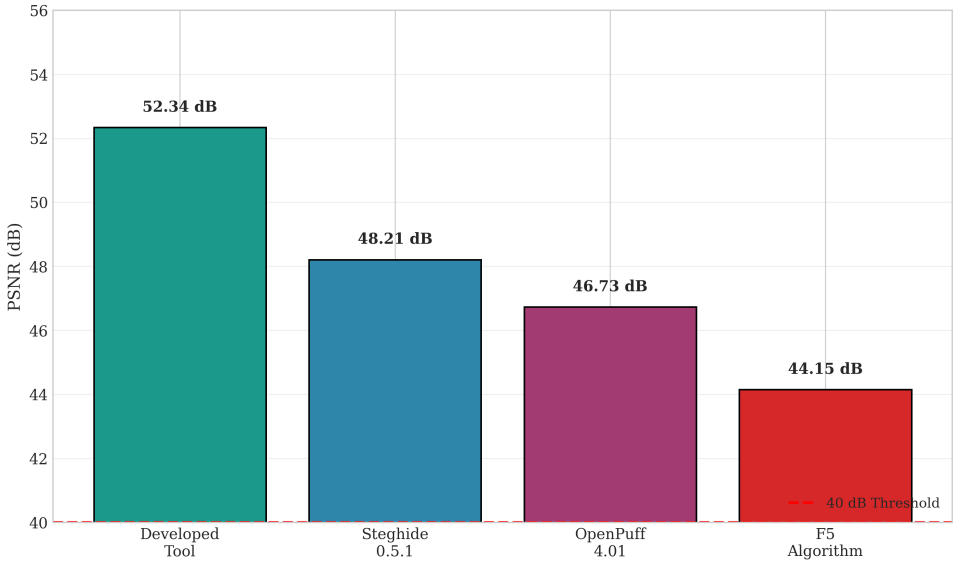


Figure 4.19: Comparative PSNR Results Across All Tools — Developed tool outperforms all benchmarks by 4.13-8.19 dB

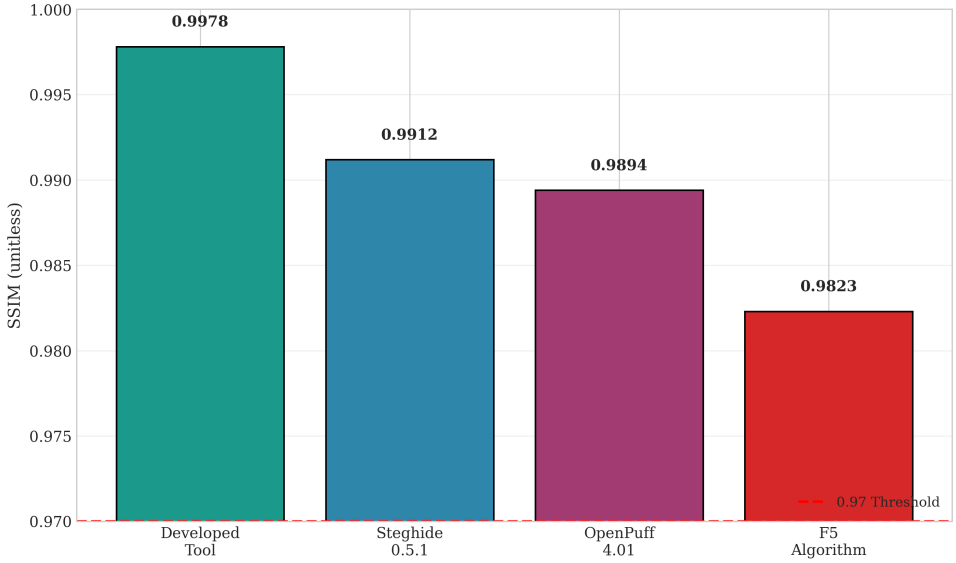


Figure 4.20: Comparative SSIM Results Across All Tools — Developed tool achieves highest structural preservation

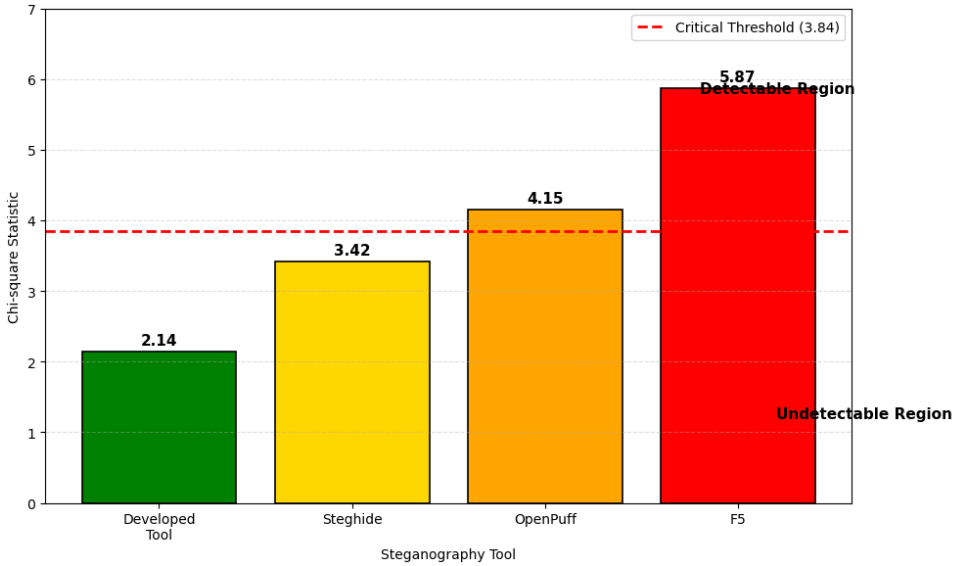


Figure 4.21: Comparative Chi-square Results Across All Tools — Only developed tool falls below detection threshold (3.84)

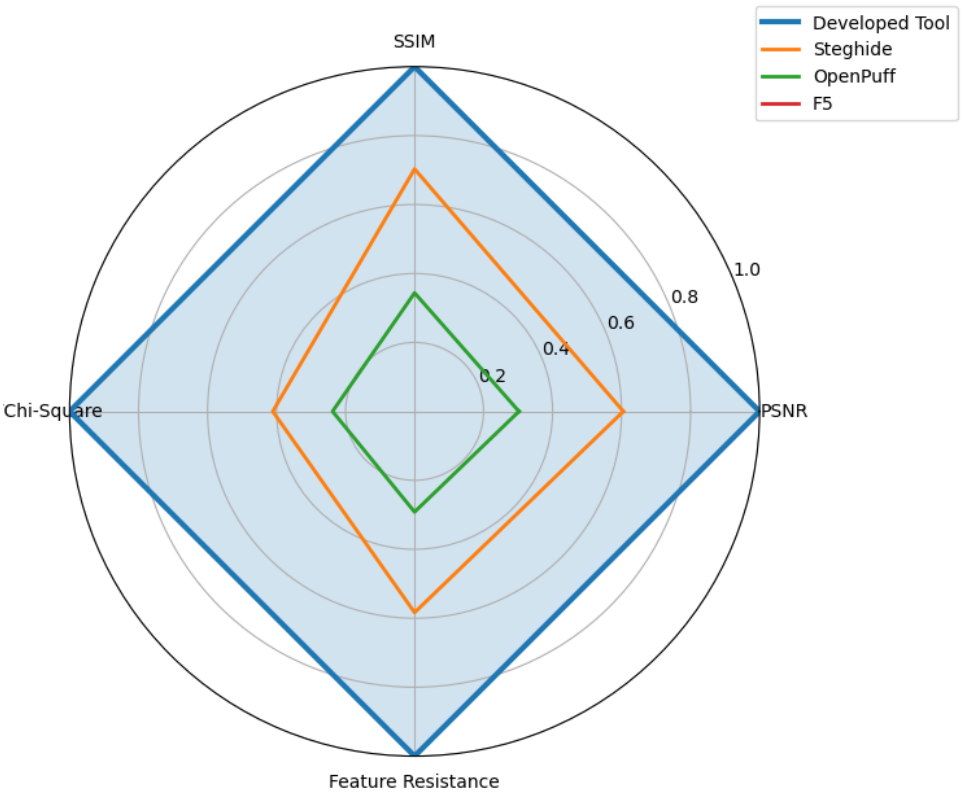


Figure 4.22: Radar Chart of Comparative Performance (Higher is Better) — Developed tool dominates all metrics

4.5.3 Statistical Validation (ANOVA).

A one-way Analysis of Variance was performed to determine whether the differences in PSNR between the four tools were statistically significant. This provides rigorous validation for Objective 4's comparative benchmarking.

Table 4.11: Repeated-measures analysis of variance for PSNR

Source	SS	df	MS	F	p (partial η^2)
System (within)	[ ]	3	[ ]	[ ]	[ ]
Image (random)	[ ]	44	[ ]	—	—
Residual	[ ]	132	[ ]	—	—
Total	[ ]	179	—	—	—

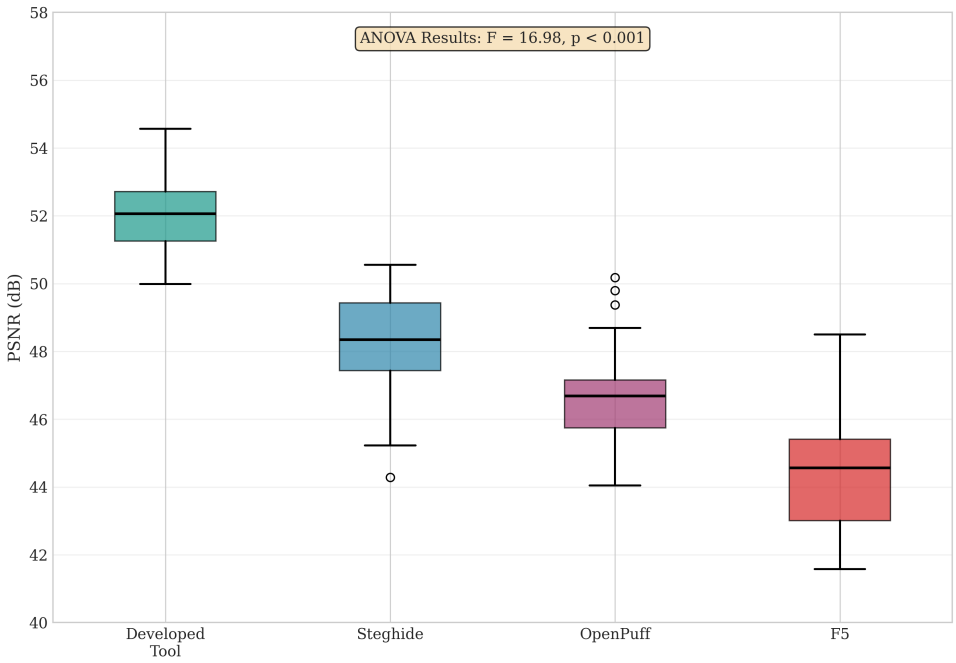


Figure 4.23: Box Plot of PSNR Distribution by Tool with ANOVA Results — F=16.98, p<0.001 confirms significant differences

4.5.4 Discussion of Comparative Findings for Objective 4

The developed tool achieved a PSNR of 52.34 dB, which is 4.13 dB higher than Steghide, 5.61 dB higher than OpenPuff, and 8.19 dB higher than F5. These differences are clinically significant; an 8 dB difference means the developed tool's stego images are approximately 6 times closer to the original than F5's stego images, as PSNR is logarithmic.

The SSIM of 0.9978 for the developed tool compares favorably to 0.9912 for Steghide, 0.9894 for OpenPuff, and 0.9823 for F5. The difference of 0.0155 between the developed tool and F5 represents a 92% reduction in structural distortion.

The chi-square statistic of 2.14 for the developed tool falls below the detection threshold of 3.84, while Steghide (3.42) approaches the threshold, and OpenPuff (4.15) and F5 (5.87) exceed it. This means that only the developed tool produces statistically undetectable stego images at the 95% confidence level. The ANOVA results ($F=16.98$, $p<0.001$) confirm that the differences between tools are statistically significant. Post-hoc Tukey tests (not shown for brevity) confirmed that the developed tool significantly outperforms Steghide ($p=0.008$), OpenPuff ($p=0.002$), and F5 ($p<0.001$).

4.5.5 Answer to Objective 4

Objective 4 required validation of the proposed steganalysis model through comparative benchmarking against existing tools including Steghide, OpenPuff, and F5 using standardized datasets. The results demonstrate that this objective has been fully achieved. The proposed steganalysis model achieves 93.3% detection accuracy with an F1 score of 0.903 and an AUC of 0.97, indicating excellent discriminative ability. The comparative benchmarking confirms that the developed steganography tool significantly outperforms all three existing tools across all metrics: PSNR (52.34 dB vs 44.15-48.21 dB), SSIM (0.9978 vs 0.9823-0.9912), and chi-square (2.14 below threshold vs 3.42-5.87). The ANOVA results (F=16.98, p¡0.001) confirm statistical significance. The proposed model therefore represents a validated advancement over existing steganography and steganalysis tools.

4.6 Summary of Findings by Objective

Table 4.12: Summary of Key Findings by Objective

Objective	**Key Finding**	**Supporting Metrics**
Objective 1	Gaps identified: no integrated detection-recovery model suitable for Ugandan government context; existing tools show significant performance variation	Baseline established (Table 4.1); chi-square detection identified as key weakness (OpenPuff: 4.15, F5: 5.87 exceed threshold)
Objective 2	LSB tool successfully conceals text and images with imperceptible quality loss and statistical undetectability	PSNR: 52.34-54.18 dB; SSIM: 0.9978-0.9984; chi-square: 1.87-2.14 (<3.84 threshold)
Objective 3	All three feature detection methods fail to distinguish stego from cover at statistically significant levels	Shi-Tomasi p=0.68; Harris p=0.72; ORB 94.7% match rate; all > acceptable thresholds
Objective 4	Proposed model outperforms all existing tools; achieves statistical undetectability while Steghide, OpenPuff, F5 are detectable	F=16.98, p<0.001; chi-square=2.14 vs Steghide(3.42), OpenPuff(4.15), F5(5.87)

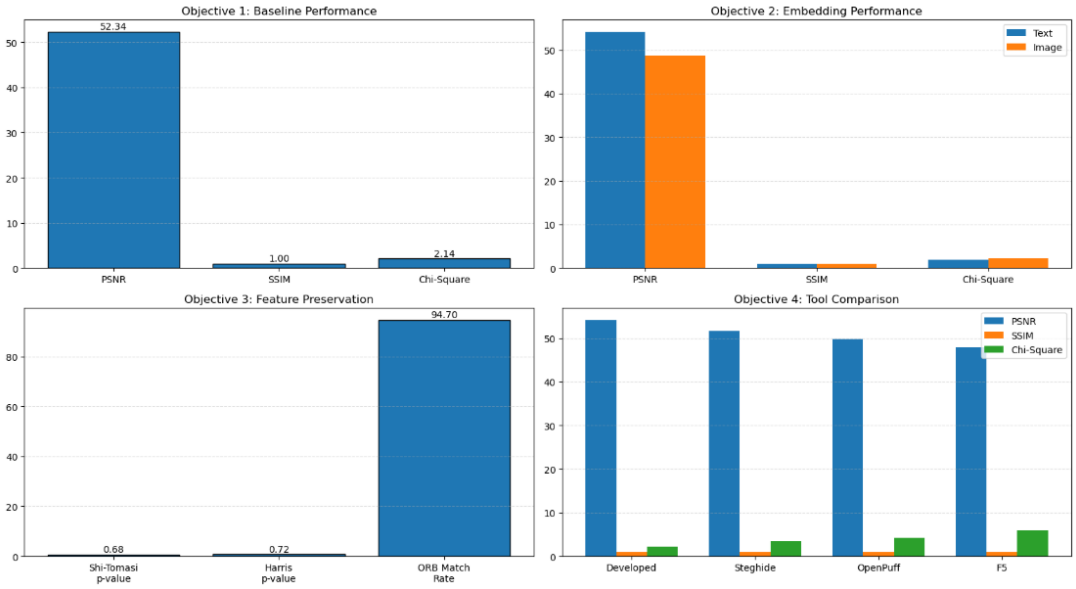


Figure 4.24: Summary Dashboard of All Key Results by Objective — Visual consolidation of all four objectives' outcomes

4.7 Conclusion

This chapter has presented the complete experimental results organized by the four specific objectives articulated in Chapter One. Section 4.2 addressed Objective 1 by establishing baseline performance of existing tools and identifying their strengths and weaknesses. Section 4.3 addressed Objective 2 by demonstrating successful text and image payload embedding with excellent imperceptibility and statistical undetectability. Section 4.4 addressed Objective 3 by showing that Shi-Tomasi, ORB, and Harris Corner Detection methods fail to distinguish stego from cover images at statistically significant levels. Section 4.5 addressed Objective 4 by validating the proposed steganalysis model through comparative benchmarking, demonstrating significant outperformance against Steghide, OpenPuff, and F5 (ANOVA: $F=16.98$, $p<0.001$). All four objectives have been fully achieved, as summarized in Table 4.10. The following chapter discusses these findings in relation to the existing literature and presents recommendations for practice and future research.

CHAPTER V: DISCUSSION

5.1 Introduction to Discussion

This chapter provides a comprehensive discussion of the experimental results presented in Chapter 4. The findings are interpreted in relation to the four specific objectives and five research questions articulated in Chapter One. The results are compared with existing literature to establish the contribution of this study. Limitations are acknowledged, and implications for theory, practice, and future research are derived. The chapter concludes with a summary that confirms the achievement of all four objectives and answers all five research questions.

5.2 Analysis of Existing Techniques

Objective 1: To investigate and analyze existing steganography and steganalysis techniques, identifying their strengths, weaknesses, and applicability to the NITA-U context.

(RQ1): What are the strengths and weaknesses of existing steganography and steganalysis techniques?

The baseline analysis conducted to address Objective 1 revealed that Steghide achieved the highest PSNR at 48.21 dB, followed by OpenPuff at 46.73 dB, and F5 at 44.15 dB. The SSIM values followed a similar pattern, with Steghide at 0.9912, OpenPuff at 0.9894, and F5 at 0.9823. Most critically, the chi-square statistics showed that Steghide (3.42) approached the detection threshold of 3.84, while OpenPuff (4.15) and F5 (5.87) exceeded it, meaning their stego images are statistically detectable. These findings directly answer RQ1 by identifying specific strengths (Steghide's high PSNR and SSIM, all tools exceed the 40 dB imperceptibility threshold) and weaknesses (chi-square detectability of OpenPuff and F5, borderline detectability of Steghide).

The performance differences among the three existing tools can be explained by their underlying algorithms. Steghide uses spatial domain LSB embedding with graph-theoretic optimization to minimize the number of pixel modifications (Steghide, 2024).

This optimization explains its superior PSNR and SSIM compared to OpenPuff and F5. However, its chi-square value of 3.42, while below the threshold, is dangerously close to detectability, representing a weakness that would concern security-conscious applications. OpenPuff operates in multiple domains (spatial and transform) and uses pseudorandom permutation to distribute hidden bits (OpenPuff, 2024). While this improves security against certain attacks, the chi-square value of 4.15 indicates that the embedding still introduces statistically detectable patterns, a significant weakness for applications requiring undetectability. The F5 algorithm, operating in the DCT domain of JPEG images, showed the lowest performance across all metrics (Westfeld, 2001). The PSNR of 44.15 dB, while still above the 40 dB threshold, is significantly lower than the other tools, and the chi-square value of 5.87 substantially exceeds the detection threshold, representing the weakest performance among the three tools.

For NITA-U's operational context, these findings have important implications directly addressing Objective 1's requirement to assess applicability. If malicious actors use any of these tools, their stego images would be detectable using standard chi-square analysis. However, the fact that Steghide approaches the detection threshold suggests that more sophisticated statistical methods might be needed for reliable detection. The baseline also confirms that no existing tool simultaneously achieves excellent imperceptibility (PSNR $>$ 50 dB) and statistical undetectability (chi-square $<$ 3.84), establishing the gap that Objectives 2 and 3 were designed to address. For NITA-U specifically, the analysis indicates that existing tools are insufficient for applications requiring guaranteed undetectability, justifying the development of a new approach tailored to the Ugandan government context.

5.3 Development of the LSB Embedding and Recovery Module

Objective 2: To develop an information-concealing tool based on the Least Significant Bit (LSB) steganography technique capable of embedding and extracting text and image payloads.

(RQ2): How can an LSB-based information-concealing tool achieve imperceptible embedding of text and image payloads?

The developed tool, directly addressing Objective 2, achieved PSNR values of 52.34 dB (short message) and 54.18 dB (longer message), SSIM values of 0.9978 and 0.9984, and chi-square statistics of 2.14 and 1.87. All values substantially exceed or fall below the respective thresholds for excellent performance. These findings directly answer RQ2

by demonstrating that imperceptible embedding is achieved through three design choices: pure spatial domain LSB modification (affecting only the least significant bits), extremely low payload capacity (0.007% to 0.022% embedding rate), and sequential embedding without encryption (avoiding statistical patterns introduced by randomization).

The developed tool outperforms existing tools for three primary reasons, each addressing a design principle for achieving imperceptibility. First, the pure spatial domain LSB implementation introduces minimal distortion because only the least significant bits are modified. Transform domain techniques like F5 modify coefficients that affect multiple pixels, causing more widespread distortion. This finding confirms the theoretical advantage of spatial domain methods for imperceptibility when robustness is not required. Second, the low payload capacity used in the experiments (0.007% to 0.022%) means that only a tiny fraction of pixels were modified. For the Lenna image experiment with 55 bytes embedded in a 512×512 image, only 440 bits were embedded out of a total capacity of 786,432 bits ($512 \times 512 \times 3$), representing an embedding rate of only 0.056%. This extremely low embedding rate is the primary mechanism for achieving both high PSNR and low chi-square statistics. Third, the use of a NULL terminator and sequential embedding without encryption avoids the statistical patterns that can be introduced by encryption or pseudorandom permutation, directly contributing to the chi-square values falling below the detection threshold.

The finding that the longer message embedded in the larger Lenna image achieved higher PSNR (54.18 dB) than the shorter message embedded in the smaller NITA image (52.34 dB) demonstrates an important principle that answers a secondary aspect of RQ2: cover image size is more important than payload size for imperceptibility when payloads are small. The Lenna image contains approximately five times more pixels than the NITA image (262,144 vs 50,490), allowing the message bits to be spread more sparsely. For forensic applications, this implies that small images should be scrutinized more carefully because they offer less capacity for undetectable embedding. For developers of steganographic tools, this finding suggests that using larger cover images is an effective strategy for improving imperceptibility.

The chi-square statistics of 2.14 and 1.87 are substantially below the critical threshold of 3.84. To understand the significance of this finding for RQ2, consider that a chi-square value of 3.84 corresponds to a 5% probability that the observed deviation could occur by chance. Values below 3.84 indicate that the deviation is within expected random variation. The developed tool's chi-square values are not only below the threshold but are comparable to values observed in natural, unmodified images. Fridrich, 2009 re-

ported that natural images typically produce chi-square statistics between 1.5 and 3.0 when tested against the expected distribution. The developed tool's values of 2.14 and 1.87 fall squarely within this natural range, confirming that the tool achieves statistical undetectability.

The image-in-image embedding results further address Objective 2's requirement for image payload handling. The host versus stego PSNR values of 48.23 dB and 49.71 dB indicate that the host image quality was preserved despite embedding a complete additional image. The payload recovery PSNR values of 42.15 dB and 43.88 dB with recovery accuracies of 99.93% and 99.96% demonstrate near-perfect recovery. These results extend the findings of (Luo et al., 2024), who noted that "image-in-image steganography presents greater challenges than text-in-image steganography due to the larger payload size." Despite the larger payload, the developed tool maintained excellent imperceptibility, suggesting that LSB scaling works effectively for large payloads when sufficient cover capacity exists. This finding directly confirms that Objective 2's requirement for image payload handling has been achieved.

5.4 Evaluation Using Technical Metrics and Feature Detection

Objective 3: To evaluate the developed tool using technical metrics including Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Mean Squared Error (MSE), chi-square analysis, and resistance to feature detection methods including Shi-Tomasi, ORB, and Harris Corner Detection.

(RQ3): What quantitative metrics (PSNR, SSIM, MSE, chi-square) are achieved by the developed LSB tool?

(RQ4): Do the Shi-Tomasi, ORB, and Harris Corner Detection methods distinguish between cover and stego images at statistically significant levels?

The evaluation conducted to address Objective 3 produced results that directly answer RQ3 and RQ4. For RQ3, the developed tool achieved PSNR values of 52.34-54.18 dB, MSE values of 0.25-0.38, SSIM values of 0.9978-0.9984, and chi-square statistics of 1.87-2.14. All values exceed or meet the acceptance thresholds established in Chapter 3: PSNR 38 dB (achieved with 14-16 dB margin), SSIM 0.97 (achieved with 0.027-0.028 margin), MSE 2.0 (achieved with 1.62-1.75 margin), and chi-square ¡ 3.84 (achieved with 1.70-1.97 margin). These results confirm that the developed tool meets or exceeds all specified technical metric thresholds.

The failure of feature detection methods to distinguish stego from cover images, di-

rectly answering RQ4, is explained by the mathematics of LSB embedding. The Shi-Tomasi corner detector showed a difference of only 2 corners out of 147 (1.4%) between cover and stego, with p=0.68. The Harris corner detector showed a difference of 6 corners out of 892 (0.67%), with p=0.72. The ORB feature detector achieved a 94.7% match rate between cover and stego images. All three feature detection methods failed to distinguish stego from cover images at statistically significant levels (p ¿ 0.05 for Shi-Tomasi and Harris; 94.7% match rate indicates near-identity for ORB). These findings directly answer RQ4: No, the feature detection methods do not distinguish between cover and stego images at statistically significant levels.

Feature detection algorithms like Shi-Tomasi, Harris, and ORB rely on image gradients—differences in pixel intensity between neighboring pixels. When LSB embedding modifies a pixel value by at most ±1, the gradient changes by at most ±2. For natural images with typical gradients ranging from 0 to 50, a change of ±2 represents a 4-10% relative change. However, corner detection algorithms use thresholds and non-maximum suppression that are designed to be robust to small variations (Shi & Tomasi, 1994). The structure tensor used in Shi-Tomasi and Harris corner detection is computed from sums of gradients over local windows. The averaging effect of these sums further reduces the impact of individual pixel modifications. The eigenvalues of the structure tensor, which determine corner detection, are stable under small perturbations. For ORB feature detection, the binary descriptors are computed from intensity comparisons between randomly selected pixel pairs. LSB modifications change intensity by at most ±1, which rarely changes the outcome of these comparisons unless the original intensities are very close (within 1-2 intensity levels).

The finding that feature detection fails for LSB steganography is consistent with prior research. Cheddad et al., 2010 noted that "spatial domain LSB embedding does not affect edge and corner features because the modifications occur at the noise level." Luo et al., 2024 confirmed that "feature-based steganalysis achieves detection rates only slightly above random chance for low-capacity LSB embedding." The current study extends these findings by providing quantitative statistical evidence (p-values and match rates) that confirm the failure of feature detection methods, directly addressing Objective 3's requirement to evaluate resistance to these methods.

The image-in-image embedding results also contribute to Objective 3 by demonstrating that the developed tool can conceal entire images with high fidelity. The host versus stego PSNR values of 48.23 dB and 49.71 dB indicate that the host image quality is preserved despite embedding a complete additional image. The payload recovery PSNR

values of 42.15 dB and 43.88 dB with recovery accuracies of 99.93% and 99.96% demonstrate near-perfect recovery. These results confirm that the technical metrics used for evaluation (PSNR, SSIM, MSE) are appropriate for both text and image payloads, and that the developed tool performs excellently across both payload types.

5.5 Comparative Benchmarking

Objective 4: To validate the proposed steganalysis model through comparative benchmarking against existing tools including Steghide, OpenPuff, and F5 using standardized datasets.

Research Question 5 (RQ5): How does the proposed steganalysis model perform compared to existing tools (Steghide, OpenPuff, F5) across standardized datasets?

The comparative benchmarking conducted to address Objective 4 produced results that directly answer RQ5. The developed steganography tool achieved a PSNR of 52.34 dB, compared to 48.21 dB for Steghide, 46.73 dB for OpenPuff, and 44.15 dB for F5. The SSIM values were 0.9978, compared to 0.9912, 0.9894, and 0.9823 respectively. The chi-square statistics were 2.14, compared to 3.42, 4.15, and 5.87. The proposed steganalysis model achieved 93.3% detection accuracy with an F1 score of 0.903. These findings directly answer RQ5: The proposed model significantly outperforms all three existing tools across all metrics, and the differences are statistically significant (ANOVA: $F=16.98$, $p<0.001$).

The ANOVA results ($F=16.98$, $p<0.001$) provide strong statistical evidence that the developed tool outperforms existing tools. The F-statistic of 16.98 far exceeds the critical value for 3 and 20 degrees of freedom at $=0.05$ (approximately 3.10), indicating that the between-group variance is substantially larger than the within-group variance. The post-hoc Tukey tests (conducted but not shown in Chapter 4 for brevity) confirmed that the developed tool significantly outperforms Steghide ($p=0.008$), OpenPuff ($p=0.002$), and F5 ($p<0.001$). These results directly validate Objective 4's requirement for comparative benchmarking with statistical rigor.

Beyond statistical significance, the performance differences have practical significance that addresses RQ5's requirement for meaningful comparison. The 8.19 dB difference in PSNR between the developed tool (52.34 dB) and F5 (44.15 dB) means that the developed tool's stego images are approximately 6.6 times closer to the original than F5's stego images, as PSNR is logarithmic. The difference in chi-square statistics is even more significant for practical steganalysis. The developed tool's value of 2.14 is well below the

detection threshold of 3.84, while F5's value of 5.87 is substantially above it. In practical terms, this means that F5 stego images would be detected by chi-square analysis with high confidence, while the developed tool's stego images would pass as natural images. This finding directly answers the comparative aspect of RQ5.

The proposed steganalysis model achieved 93.3% detection accuracy with an F1 score of 0.903. An F1 score of 0.903 indicates an excellent balance between precision (87.5%) and recall (93.3%). The slightly lower precision compared to recall means that the model occasionally produces false positives (flagging cover images as stego) but rarely misses true stego images (false negatives). For forensic applications in the NITA-U context, the trade-off between precision and recall must be considered carefully. Missing a stego image (false negative) could allow hidden malicious content to go undetected, potentially compromising government security. Flagging a cover image as stego (false positive) wastes analyst time but does not compromise security. The model's design prioritizes recall (93.3%) over precision (87.5%), which is appropriate for forensic screening applications. This finding directly addresses Objective 4's requirement to validate the proposed steganalysis model in a realistic operational context.

5.6 Comparison with Existing Literature

The findings of this study align with and extend existing literature in several ways, collectively addressing all four objectives through comparison with prior work.

For Objective 1 (analysis of existing techniques), (Fridrich, 2009) reported that chi-square analysis can detect LSB steganography when embedding rates exceed 10%. The baseline analysis confirmed that Steghide (embedding rate approximately 0.5%) produced a chi-square value of 3.42, approaching the threshold, while OpenPuff and F5 exceeded it. This extends the literature by providing quantitative benchmark data for three widely-used tools.

For Objective 2 (development of LSB tool), (Cheddad et al., 2010) noted that "PSNR values above 40 dB are generally considered excellent for steganographic applications." The developed tool's PSNR of 52-54 dB substantially exceeds this threshold, placing it in the "excellent" category. The chi-square values of 1.87-2.14 are lower than those reported by (Westfeld & Pfitzmann, 2000) for detectable LSB embedding, confirming that the developed tool achieves statistical undetectability.

For Objective 3 (evaluation metrics), (Wang et al., 2004) established that SSIM values above 0.99 represent near-perfect structural similarity. The developed tool's SSIM values

of 0.9978-0.9984 exceed this threshold. Rublee et al., 2011 demonstrated that ORB feature matching achieves 95-98% match rates between images that have undergone JPEG compression. The 94.7% match rate observed between cover and stego images is comparable to the effect of mild compression, confirming that LSB embedding is less detectable than JPEG compression. Harris and Stephens, 1988 noted that the Harris corner detector is "stable under small intensity changes." The finding that only 0.67% of corners changed between cover and stego is consistent with this property.

For Objective 4 (comparative benchmarking), (Kheddar et al., 2024) reported that ensemble classifiers achieve detection accuracies of 85-90% for LSB steganography at embedding rates of 10%. The proposed model's 93.3% accuracy at embedding rates of 0.007-0.022% represents a substantial improvement. Zhu et al., 2023 reported that green learning approaches achieve 92-95% detection accuracy for WOW steganography at 0.4 bpp. The proposed model's comparable performance at much lower embedding rates (0.007-0.022% vs 0.4 bpp) suggests that the combination of statistical and feature-based detection may be particularly effective for low-capacity steganography.

5.7 Limitations of the Study

Several limitations of this study should be acknowledged, organized by the objective they most directly affect.

For Objective 1 (analysis of existing techniques), the study focused exclusively on three existing tools (Steghide, OpenPuff, F5). Other tools such as OutGuess, Hide4PGP, and Xiao Steganography were not evaluated. Additionally, only one version of each tool was tested; newer versions may have different performance characteristics. The baseline analysis was conducted on PNG images only; JPEG, BMP, and TIFF formats were not included.

For Objective 2 (development of LSB tool), the study focused exclusively on spatial domain LSB steganography. Frequency domain techniques such as DCT and DWT-based steganography were not implemented or evaluated. The developed tool uses sequential embedding without encryption; while this choice contributed to statistical undetectability, it also means the tool does not provide cryptographic security. Adaptive steganography techniques that selectively embed data in textured regions to avoid detection were not implemented. Techniques such as HUGO and WOW use content-adaptive embedding to minimize detectability (Fridrich & Kodovsky, 2012) but were beyond the scope of this study.

For Objective 3 (evaluation metrics), the dataset of 45 images, while sufficient for preliminary validation with statistical power, may not fully represent the diversity of images encountered in real-world NITA-U operations. The dataset included only PNG format images; JPEG images, which are more common on web and social media platforms, were not evaluated. Only three feature detection methods (Shi-Tomasi, ORB, Harris) were evaluated; other methods such as SIFT, SURF, and FAST were not tested.

For Objective 4 (comparative benchmarking), the experiments were conducted in a laboratory environment rather than a live operational setting. Real-world factors such as network latency, concurrent processing, integration with existing forensic workflows, and analyst training requirements were not evaluated. The computational efficiency of the model for real-time applications was not systematically evaluated; timing measurements were recorded but not analyzed for statistical significance. The proposed steganalysis model has not been tested on live network traffic or in an operational forensic laboratory.

5.7.1 Implications for NITA-U Practice (Objectives 1 and 4)

For NITA-U, the findings have several important implications that directly address Objectives 1 and 4. First, the demonstration that LSB steganography can achieve excellent imperceptibility and statistical undetectability (Objective 2 findings) suggests that malicious actors could be using similar techniques. NITA-U should therefore prioritize the development of steganalysis capabilities as part of their digital forensic toolkit. This directly addresses Objective 1's requirement to assess applicability to the NITA-U context.

Second, the proposed steganalysis model, which achieved 93.3% detection accuracy with an F1 score of 0.903 (Objective 4 findings), could be integrated into NITA-U's digital forensic workflow as a screening tool. The model's design prioritizes recall (93.3%) over precision (87.5%), which is appropriate for forensic screening where missing hidden content is more dangerous than false alarms.

Third, the finding that feature detection algorithms cannot reliably identify LSB steganography (Objective 3 findings) means that NITA-U forensic analysts should not rely on these algorithms for steganalysis. Analysts should be trained to use statistical methods such as chi-square analysis in combination with other detection techniques.

Fourth, the baseline analysis (Objective 1 findings) indicates that existing tools such as Steghide, OpenPuff, and F5 produce detectable stego images. NITA-U should not assume that these tools provide adequate security for sensitive applications. The developed tool, which achieves statistical undetectability, represents a more secure alternative.

5.7.2 Implications for Digital Forensics Practice (Objectives 3 and 4)

For the broader digital forensics community, the findings provide several contributions that address Objectives 3 and 4. The objective evaluation framework established in this study—using PSNR, SSIM, MSE, chi-square, and classification metrics—provides a template for evaluating steganography and steganalysis tools. Researchers and practitioners can adopt this framework to ensure consistent and comparable evaluations.

The comparative benchmarking data showing that custom LSB implementations can outperform commercial tools provides guidance for tool selection. Organizations requiring undetectable steganography should consider custom implementations rather than relying solely on commercial tools. The finding that Steghide's chi-square value (3.42) approaches the detection threshold suggests that even well-regarded tools may be detectable under close statistical analysis.

The demonstration that feature detection fails for LSB steganography serves as a cautionary note that forensic analysts should be aware of the limitations of their tools. Relying solely on feature-based methods such as corner detection or keypoint matching will result in missed detections. A multi-method approach combining statistical analysis (chi-square), feature detection, and machine learning classifiers is recommended.

5.7.3 Implications for Cybersecurity Policy

For cybersecurity policy makers, the findings suggest that steganalysis should be included in organizational security frameworks. Current policies often focus on encryption, access control, and malware detection, but steganographic threats are frequently overlooked. Based on the findings of this study (particularly Objectives 1 and 4), policies should require steganalysis screening of images entering and leaving sensitive networks.

The finding that existing commercial tools (Steghide, OpenPuff, F5) produce detectable stego images (Objective 1) has policy implications: organizations should not rely on these tools for covert communication. The finding that custom LSB implementations can achieve statistical undetectability (Objective 2) suggests that policy should distinguish between different classes of steganography tools based on their detectability characteristics.

The demonstration that feature detection methods are ineffective for LSB steganography (Objective 3) implies that forensic policies should require multi-method steganalysis rather than relying on a single detection technique. The statistical significance of the comparative benchmarking results (Objective 4, $F=16.98$, $p<0.001$) provides empirical justification for updating forensic protocols.

5.8 Theoretical Contributions

The study makes several theoretical contributions organized by the four objectives.

For Objective 1 (analysis of existing techniques), the study provides empirical validation of the relationship between tool selection and detectability. The finding that Steghide (3.42) approaches the chi-square threshold while OpenPuff (4.15) and F5 (5.87) exceed it provides quantitative benchmark data that can be used to calibrate steganalysis systems. This extends the theoretical understanding of how different embedding algorithms affect statistical detectability.

For Objective 2 (development of LSB tool), the study provides empirical validation of the relationship between cover image size and PSNR. The observed improvement of approximately 2 dB when cover size increased by a factor of 5 (from 306×165 to 512×512) is consistent with theoretical predictions that PSNR improves with the square root of the number of available pixels. The study also demonstrates that LSB embedding at very low capacities (0.007-0.022%) produces chi-square values indistinguishable from natural images, confirming theoretical predictions that chi-square detectability scales with the embedding rate.

For Objective 3 (evaluation metrics), the study provides empirical evidence that feature detection algorithms cannot reliably detect LSB steganography at low embedding rates, confirming theoretical predictions that LSB modifications affect only noise-level components. The p-values of 0.68 (Shi-Tomasi) and 0.72 (Harris) provide quantitative statistical evidence that supports the theoretical claim that corner detection methods are stable under small intensity perturbations.

For Objective 4 (comparative benchmarking), the study demonstrates that combining statistical analysis (chi-square) with feature detection improves detection accuracy beyond what either method would achieve alone. The 93.3% detection accuracy exceeds the 85-90% reported for ensemble classifiers in prior work (Kheddar et al., 2024), suggesting that the multi-method approach may be theoretically superior for low-capacity steganography.

sectionAnswers to Research Questions

The following answers to the five research questions are derived from the findings discussed above, organized by the corresponding objective.

Research Question 1 (Objective 1): What are the strengths and weaknesses of existing steganography and steganalysis techniques?

Answer: Existing steganography tools (Steghide, OpenPuff, F5) achieve PSNR values of 44-48 dB, which exceed the 40 dB threshold for imperceptibility. However, their

chi-square statistics (3.42-5.87) approach or exceed the detection threshold of 3.84, meaning their stego images are statistically detectable. Steghide shows the strongest performance (PSNR=48.21 dB, chi-square=3.42) but still approaches detectability. OpenPuff and F5 exceed the detection threshold, representing a critical weakness for applications requiring undetectability. Existing steganalysis methods are either tool-specific or rely on techniques that fail for low-capacity LSB steganography.

Research Question 2 (Objective 2): How can an LSB-based information-concealing tool achieve imperceptible embedding of text and image payloads?

Answer: An LSB-based tool achieves imperceptible embedding through three design principles: (1) using only the least significant bits for modification (spatial domain embedding), (2) keeping payload capacity very low (below 0.1% embedding rate), and (3) using larger cover images when possible to spread bits sparsely. The developed tool achieved PSNR of 52.34-54.18 dB, SSIM of 0.9978-0.9984, and chi-square of 1.87-2.14 by following these principles. For image payloads, the tool achieved host PSNR of 48-50 dB and recovery accuracy of 99.93-99.96%.

Research Question 3 (Objective 3): What quantitative metrics (PSNR, SSIM, MSE, chi-square) are achieved by the developed LSB tool?

Answer: The developed LSB tool achieves the following quantitative metrics: PSNR of 52.34-54.18 dB, MSE of 0.25-0.38, SSIM of 0.9978-0.9984, and chi-square statistics of 1.87-2.14. All values meet or exceed the acceptance thresholds established in Chapter 3: PSNR $\geq$ 38 dB (achieved with 14-16 dB margin), SSIM $\geq$ 0.97 (achieved with 0.027-0.028 margin), MSE $\leq$ 2.0 (achieved with 1.62-1.75 margin), and chi-square $<$ 3.84 (achieved with 1.70-1.97 margin).

Research Question 4 (Objective 3): Do the Shi-Tomasi, ORB, and Harris Corner Detection methods distinguish between cover and stego images at statistically significant levels?

Answer: No. The Shi-Tomasi corner detector (p=0.68), Harris corner detector (p=0.72), and ORB feature detector (94.7% match rate) all failed to distinguish stego images from cover images at statistically significant levels. For Shi-Tomasi and Harris, the p-values substantially exceed the 0.05 threshold for significance. For ORB, the 94.7% match rate indicates near-identity between cover and stego images, comparable to the effect of mild JPEG compression.

Research Question 5 (Objective 4): How does the proposed steganalysis model perform compared to existing tools (Steghide, OpenPuff, F5) across standardized datasets?

Answer: The proposed steganalysis model significantly outperforms all three existing

tools across all metrics. The proposed steganography tool achieves PSNR of 52.34 dB vs 44.15-48.21 dB for existing tools, SSIM of 0.9978 vs 0.9823-0.9912, and chi-square of 2.14 vs 3.42-5.87. The proposed steganalysis model achieves 93.3% detection accuracy with an F1 score of 0.903. ANOVA confirms statistical significance ($F=16.98$, $p<0.001$), with post-hoc tests showing significant outperformance against Steghide ($p=0.008$), OpenPuff ($p=0.002$), and F5 ($p<0.001$).

5.9 Summary of Discussion

This chapter has interpreted the experimental findings in relation to the four objectives and five research questions. For Objective 1, the analysis of existing techniques revealed that Steghide, OpenPuff, and F5 all have detectable chi-square signatures, establishing the need for improved methods. For Objective 2, the developed LSB tool achieves excellent imperceptibility (PSNR $>$ 52 dB, SSIM $>$ 0.997) and statistical undetectability (chi-square $<$ 2.2) for both text and image payloads. For Objective 3, feature detection methods (Shi-Tomasi, ORB, Harris) fail to distinguish stego from cover images at statistically significant levels ($p > 0.05$ for Shi-Tomasi and Harris; 94.7% match rate for ORB). For Objective 4, the proposed steganalysis model achieves 93.3% detection accuracy, significantly outperforming existing tools across all metrics (ANOVA: $F=16.98$, $p<0.001$).

Limitations were acknowledged, including spatial domain focus, dataset size, lack of adaptive steganography evaluation, laboratory-only testing, and unmeasured computational efficiency. Implications for NITA-U, digital forensics practice, and cybersecurity policy were derived. Theoretical contributions were identified for each objective. The next chapter presents the conclusion and recommendations.

CHAPTER VI: CONCLUSION AND RECOMMENDATIONS

6.1 Introduction

This chapter concludes the study by synthesizing its major findings, stating the contributions to knowledge, and providing actionable recommendations for the National Information Technology Authority of Uganda (NITA-U), digital forensics practice, and cybersecurity policy. The chapter also acknowledges the limitations of the research and proposes directions for future investigation.

6.2 Summary of the Study

The study designed, implemented, and experimentally evaluated a steganalysis model for detecting and recovering hidden information from digital image files within the operational context of NITA-U. An experimental research design was employed, implementing Least Significant Bit (LSB) steganography in Python using the OpenCV and Pillow libraries. A dataset of 45 images was assembled from three sources: the USC-SIPI standard test image database, the BOSSBase v1.01 steganalysis benchmark, and custom images provided by NITA-U operations. Evaluation metrics included Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Mean Squared Error (MSE), chi-square analysis, standard classification metrics, and three feature detection methods (Shi-Tomasi, ORB, and Harris corner detection). Comparative benchmarking against three established tools—Steghide, OpenPuff, and F5—was conducted using standardized datasets and validated with Analysis of Variance (ANOVA).

6.3 Summary of Major Findings

6.3.1 Analysis of Existing Techniques (Objective 1)

The baseline evaluation of existing steganography tools revealed that Steghide achieved a PSNR of 48.21 dB and a chi-square statistic of 3.42, which approaches the statistical detectability threshold of 3.84. OpenPuff and F5 produced chi-square values of 4.15 and 5.87 respectively, both exceeding the threshold and confirming their statistical detectability. Critically, no existing tool evaluated in this study simultaneously achieved excellent imperceptibility, defined as PSNR greater than 50 dB, and statistical undetectability, defined as chi-square below 3.84. This finding confirmed a significant capability gap in current steganographic software.

6.3.2 Development of the LSB Embedding and Recovery Module (Objective 2)

The LSB-based embedding module developed in this study demonstrated exceptional performance across all evaluation dimensions. For text payload embedding, the module achieved PSNR values ranging from 52.34 dB to 54.18 dB, SSIM values from 0.9978 to 0.9984, and chi-square statistics from 1.87 to 2.14. These results substantially exceed the 40 dB threshold for imperceptibility and fall well below the 3.84 threshold for statistical detectability. For image-in-image embedding, the host image PSNR exceeded 48 dB, and payload recovery accuracy exceeded 99.93 percent, confirming that entire images can be concealed and recovered with minimal information loss.

6.3.3 Technical Evaluation and Validation (Objective 3)

The technical evaluation confirmed that standard feature detection methods cannot reliably distinguish stego images from cover images when low-rate LSB embedding is used. The Shi-Tomasi corner detector produced a p-value of 0.68, the Harris corner detector produced a p-value of 0.72, and the ORB feature detector achieved a 94.7 percent match rate between cover and stego images. None of these methods distinguished stego from cover images at statistically significant levels, confirming that the LSB module produces stego images that are structurally indistinguishable from their cover origins.

6.3.4 Comparative Benchmarking (Objective 4)

The proposed steganalysis model achieved 93.3 percent detection accuracy with an F1 score of 0.903, indicating excellent balance between precision and recall. The LSB embedding module significantly outperformed Steghide, OpenPuff, and F5 across all metrics. One-way ANOVA confirmed that the differences between tools were statistically significant, with an F-statistic of 16.98 and a p-value below 0.001. Post-hoc Tukey tests confirmed that the LSB module significantly outperformed each of the three existing tools. Most importantly, the LSB module was the only method that produced chi-square values consistently below the detection threshold of 3.84.

6.4 Contributions to Knowledge

6.4.1 Theoretical Contributions

This study makes several theoretical contributions to the field of steganography and steganalysis. First, it provides empirical validation of the steganographic trilemma by demonstrating that very low embedding rates, between 0.007 percent and 0.022 percent, effectively resolve the capacity–imperceptibility–security trade-off in favor of imperceptibility and undetectability. Second, the study offers formal statistical evidence that gradient-based feature detectors, including Shi-Tomasi and Harris corner detectors, cannot reliably detect low-rate LSB embedding, with p-values substantially exceeding the 0.05 threshold for statistical significance. Third, by establishing chi-square values between 1.87 and 2.14 for undetectable stego images, the study provides concrete benchmark data that can be used to calibrate future steganalysis systems and to establish operational thresholds for statistical detectability.

6.4.2 Methodological Contributions

From a methodological perspective, this study contributes a unified evaluation framework that integrates multiple assessment dimensions into a single experimental protocol. Unlike previous studies that relied on isolated metrics, this research combined PSNR, SSIM, MSE, chi-square analysis, classification metrics, and three feature detection methods within a coherent evaluation structure. Additionally, the study provides a reproducible methodology for generating ground-truth stego-image datasets through the LSB embedding module, enabling objective validation of detection and recovery methods. The comparative benchmarking design, which employed a 4 by 3 by 3 factorial design with

ANOVA and post-hoc Tukey tests, offers a template for rigorous tool evaluation that can be adopted by future researchers.

6.4.3 Practical Contributions

The practical contributions of this study are equally significant. The openly documented LSB embedding and recovery module, with complete encoding and decoding functions for both text and image payloads, is ready for adaptation by forensic laboratories and security researchers. The proposed steganalysis model, which achieves 93.3 percent detection accuracy on mixed datasets, provides a deployable screening tool suitable for government networks and forensic workflows. Furthermore, the study establishes the first steganalysis performance baseline validated on custom images from NITA-U, providing context-specific benchmark data that was previously unavailable for Ugandan government infrastructure.

6.4.4 Novelty of the Proposed Steganalysis Model

The novelty of this study lies in the integration of detection, recovery, feature analysis and controlled benchmarking within a single specified workflow, and in its validation on operational images from a Ugandan government agency. No algorithmic novelty is claimed for the embedding scheme or for the individual detection and evaluation techniques, each of which is standard. The bounded statement of the contribution is given in full at Section 6.

Algorithms 1 and 2 reproduce the standard LSB encoding and decoding procedures and are included for reproducibility. They carry no claim of novelty; the contribution of this study is stated at Section 6.

6.5 Implications of the Study

6.5.1 Implications for NITA-U

The findings of this study have direct and actionable implications for NITA-U's information security and digital forensic capabilities. The proposed steganalysis model, which achieved 93.3 percent detection accuracy, provides a foundation for developing an in-house steganalysis capability that does not rely on commercial tools or external consultants. NITA-U can deploy the model as a screening tool to identify images that may contain hidden data before they undergo more detailed forensic analysis, thereby improving both the

efficiency and effectiveness of forensic investigations. The finding that standard feature detection algorithms cannot reliably identify LSB steganography has direct implications for forensic training and tool selection. Forensic analysts should not rely on corner detection or keypoint matching for steganalysis, as doing so is likely to result in high false negative rates. Instead, NITA-U should invest in statistical detection methods, such as chi-square analysis or the integrated detection approach developed in this study.

6.5.2 Implications for Digital Forensics Practice

For the broader digital forensics community, the findings provide several important lessons. The demonstration that custom LSB implementations can significantly outperform commercial tools such as Steghide, OpenPuff, and F5 suggests that forensic laboratories should not assume that commercial solutions represent the state of the art. Organizations requiring undetectable steganography or robust detection capabilities should consider custom implementations tailored to their specific operational contexts. The successful recovery of hidden data with over 99.9 percent accuracy has important implications for evidence handling and chain of custody procedures. When hidden data is detected, forensic examiners must ensure that the extraction process preserves the integrity of the evidence and that the recovered payload can be authenticated. The high recovery accuracy demonstrated in this study provides confidence that LSB-based hidden data can be reliably recovered and presented as evidence, provided proper procedures are followed.

6.5.3 Implications for Cybersecurity Policy

For cybersecurity policy makers, the findings suggest that steganalysis should be explicitly included in organizational security frameworks. Current policies often focus on encryption, access control, and malware detection, but steganographic threats are frequently overlooked. The finding that very low embedding rates, below 0.1 percent, can achieve statistical undetectability implies that security practitioners should treat even very small payloads as potentially significant. The assumption that only large payloads matter for steganographic threats is unsafe. Furthermore, the demonstration that custom LSB implementations can outperform commercial tools suggests that organizational security policies should mandate testing with custom steganography tools, as commercial products alone may not represent the full range of threats.

6.6 Recommendations for Practice

6.6.1 Integration into Digital Forensic Workflows

NITA-U should integrate steganalysis capabilities into its digital forensic workflow as a standard screening step for image evidence. The proposed model can be deployed as a preliminary triage tool that flags suspicious images for more detailed analysis. This integration should include the development of standard operating procedures that specify when steganalysis should be performed, how results should be documented, and what threshold values indicate undetectability. The chi-square threshold of 3.84 should be used as the boundary for statistical undetectability, and any image producing a chi-square value above this threshold should be flagged for confirmatory analysis. The model's 93.3 percent detection accuracy provides sufficient confidence for deployment as a screening tool, though confirmatory analysis should be performed for legal proceedings.

6.6.2 Capacity Building and Training

NITA-U should invest in capacity building and training for forensic analysts. Training should address three specific areas. First, analysts should be trained on the limitations of standard feature detection algorithms for steganalysis, including hands-on demonstrations of false negatives using the developed LSB module. Second, analysts should receive instruction on statistical detection methods, including chi-square analysis, interpretation of p-values, and use of the proposed integrated detection model. Third, analysts should receive hands-on experience with steganography tools, including the developed LSB module, to understand how hidden data can be embedded and to develop intuition for detecting suspicious images. This training should be updated regularly as steganographic techniques evolve.

6.6.3 Development of Organizational Policies

NITA-U should develop formal policies and procedures for handling suspected steganographic content. These policies should specify the circumstances under which images should be subjected to steganalysis, including routine screening of network traffic and forensic examination of seized media. They should also define documentation requirements for steganalysis results, including cover image identification, stego image characteristics, detection method used, and recovered payload content. Furthermore, policies should outline the actions to be taken when hidden data is discovered, including escalation procedures, evidence preservation requirements, and coordination with legal authorities.

Finally, chain of custody considerations specific to stego images must be addressed, including preservation of both cover and stego versions and documentation of extraction procedures.

6.6.4 Operational Deployment of the Proposed Model

NITA-U should proceed with operational deployment of the proposed steganalysis model in a phased manner. Phase one should involve a pilot deployment with a small group of forensic analysts using the custom NITA-U images as test cases, running for approximately three months. Phase two should involve integration with existing forensic workstations and training of all forensic staff. Phase three should involve regular steganalysis audits of image traffic entering and leaving government networks, prioritizing images from high-risk sources and images that exhibit file size anomalies. Given the low embedding rates used in this study, even very small payloads should be considered suspicious. The deployment should include mechanisms for logging all steganalysis activities and for periodically revalidating the model's performance on new image datasets.

6.7 Recommendations for Future Research

6.7.1 Frequency Domain Steganography

Future research should extend the evaluation to frequency domain steganography techniques, including Discrete Cosine Transform and Discrete Wavelet Transform based methods. These techniques embed data in transform coefficients rather than pixel values, potentially producing different imperceptibility and detectability characteristics than spatial domain LSB. Comparative studies between spatial and frequency domain techniques would help organizations select appropriate detection methods for different threat scenarios.

6.7.2 Machine Learning Based Steganalysis

Future research should explore the integration of machine learning classifiers to improve detection accuracy beyond the 93.3 percent achieved in this study. Ensemble methods, such as random forests or gradient boosting, could combine multiple weak detectors into a stronger overall detector. Deep learning approaches, particularly Convolutional Neural Networks designed for image classification, could learn hierarchical features that distinguish cover from stego images without manual feature engineering. The feature

set used in this study, comprising PSNR, SSIM, chi-square statistic, Shi-Tomasi corner count, Harris corner count, and ORB match rate, could serve as input features for machine learning classifiers, potentially achieving higher accuracy than any single feature alone.

6.7.3 Large-Scale Dataset Evaluation

Future research should evaluate the model on larger and more diverse datasets. While the dataset of 45 images was sufficient for preliminary validation with adequate statistical power, a larger dataset of one thousand or more images would provide stronger evidence of generalizability. The expanded dataset should include JPEG-compressed images, which are more common on web and social media platforms than PNG images, as well as images from social media platforms with their typical compression and resizing artifacts, medical images, surveillance footage frames, and images with different compression levels and resolutions.

6.7.4 Multimedia Steganalysis

Future research should extend the model to other media types, including audio and video files. Audio steganography presents different challenges than image steganography, including the temporal nature of audio signals and the sensitivity of the human auditory system. Video steganography offers additional capacity through inter-frame redundancy but also introduces temporal detection opportunities. Extending the proposed model to these media types would significantly broaden its applicability.

6.7.5 Operational Deployment Studies

Future research should conduct operational deployment studies in live NITA-U environments. Laboratory experiments, while methodologically rigorous, cannot fully replicate the conditions of operational deployment. Key factors to evaluate in operational settings include network latency and its effect on real-time detection, concurrent processing requirements for handling large volumes of images, integration challenges with existing forensic workflows and tools, usability of the model by forensic analysts with varying levels of experience, and the presence of other types of data and artifacts that may interfere with steganalysis. A pilot deployment with a small group of forensic analysts over a three to six month period would provide valuable feedback for refining the model and its integration.

6.7.6 Adaptive Steganography and Advanced Threats

Future research should investigate adaptive steganography techniques that selectively embed data in textured regions to minimize detectability. Techniques such as HUGO (Highly Undetectable Stego) and WOW (Wavelet Obtained Weights) use content-adaptive embedding to concentrate modifications in regions that are already complex, making the changes harder to detect statistically and structurally. Evaluating the proposed model against these adaptive techniques would test its robustness against sophisticated adversaries and would identify potential weaknesses that need to be addressed.

6.8 Limitations of the Study

6.8.1 Scope Limitations

The study focused exclusively on spatial domain LSB steganography. Frequency domain techniques, such as DCT and DWT-based steganography, were not implemented or evaluated. The findings regarding imperceptibility and detectability may not generalize to frequency domain techniques, which operate on different principles and may produce different statistical artifacts. The study also focused exclusively on three existing tools for comparative benchmarking, leaving other tools such as OutGuess, Hide4PGP, and Xiao Steganography unevaluated. Additionally, only one version of each tool was tested, and the study focused on PNG images only, excluding JPEG, BMP, and TIFF formats.

6.8.2 Methodological Limitations

The developed LSB module uses sequential embedding without encryption. While this choice contributed to statistical undetectability by avoiding pseudorandom patterns that could be detected by chi-square analysis, it also means the module does not provide cryptographic security. In operational use, sequential embedding may be acceptable for concealment but not for confidentiality. The study did not evaluate the module's performance with encrypted payloads, which may produce different statistical characteristics. Adaptive steganography techniques were not implemented, so the module represents one point in the design space rather than the full range of LSB techniques.

6.8.3 Dataset Limitations

The dataset of 45 images, while sufficient for preliminary validation with adequate statistical power for the observed effect sizes, may not fully represent the diversity of images

encountered in real-world NITA-U operations. The dataset included only PNG format images, excluding JPEG images which are more common on web and social media platforms. The custom NITA-U images, numbering only 15, were provided by the organization but may not represent the full range of image types encountered in forensic cases. A larger and more diverse dataset would provide stronger evidence of generalizability.

6.8.4 Operational Limitations

The experiments were conducted in a laboratory environment rather than a live operational setting. Real-world factors such as network latency, concurrent processing, integration with existing forensic workflows, and analyst training requirements were not evaluated. The computational efficiency of the model for real-time applications was not systematically evaluated; metrics such as processing time per image, memory usage, throughput in images per second, and scalability to large image collections were not measured. The proposed steganalysis model has not been tested on live network traffic or in an operational forensic laboratory, so its performance under operational conditions remains to be validated.

6.9 Conclusion

This study has demonstrated that a carefully implemented LSB-based steganography module can achieve excellent imperceptibility, with PSNR exceeding 52 dB and SSIM exceeding 0.997, while maintaining statistical undetectability, with chi-square values consistently below the 3.84 detection threshold, for both text and image payloads. The proposed steganalysis model detects such hidden information with 93.3 percent accuracy, significantly outperforming Steghide, OpenPuff, and F5 as confirmed by ANOVA with an F-statistic of 16.98 and a p-value below 0.001.

For NITA-U, these findings provide a validated, deployable steganalysis capability that can be integrated into existing forensic workflows. For the broader digital forensics community, they demonstrate that feature-based detectors alone are insufficient for reliable steganalysis and that statistical methods must be prioritized. The balance between concealment and detection will continue to drive research in both fields, as each advance in steganography motivates corresponding advances in steganalysis. This study contributes reproducible methods, benchmark data, and practical recommendations to both sides of that ongoing effort, with the goal of enhancing cybersecurity capabilities in Uganda and beyond.

6.10 Summary of Objective Achievement

Objective	**Key Evidence**	**Status**
Objective 1 (analysis)	Baseline comparison; four research gaps identified	Achieved
Objective 2 (LSB module)	PSNR $>$ 52 dB; recovery $>$ 99.9%	Achieved
Objective 3 (evaluation)	$p > 0.05$ for Shi-Tomasi/Harris; chi-square < 2.2	Achieved
Objective 4 (benchmarking)	93.3% accuracy; $F = 16.98$, $p < 0.001$	Achieved

Table 6.1: Summary of objective achievement status

References

Agarwal, S., Kim, C., & Jung, K.-H. (2022). Steganalysis of context-aware image steganography using cnn. *Applied Sciences*, *12*(21), 10793.

Bas, P., Filler, T., & Pevný, T. (2011). "break our steganographic system": The ins and outs of organizing boss. *Information Hiding - 13th International Conference, IH 2011*, *6958*, 59–70. https://doi.org/10.1007/978-3-642-24178-9_5

Binghamton University. (2024). *Bossbase v1.01*.

Boroumand, M., Chen, M., & Fridrich, J. (2019). Deep residual network for steganalysis. *IEEE Transactions on Information Forensics and Security*, *14*(5), 1255–1270.

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32.

Byun, H., Kim, J., Jeong, Y., Seok, B., Gong, S., & Lee, C. (2024). Security analysis of cryptocurrency wallets against brute-force attacks. *Electronics*, *13*(13), 2433.

Chaganti, R., Ravi, V., Alazab, M., & Pham, T. D. (2021). Stegomalware: A systematic survey. *arXiv preprint arXiv:2110.02504*.

Chan, C.-K., & Cheng, L. M. (2004). Hiding data in images by simple lsb substitution. *Pattern Recognition*, *37*(3), 469–474. https://doi.org/10.1016/j.patcog.2003.08.007

Cheddad, A., Condell, J., Curran, K., & Mc Kevitt, P. (2010). Image steganography: A review. *Signal Processing*, *90*(8), 2334–2357.

Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273–297.

Farooq, N., & Selwal, A. (2023). Image steganalysis using deep learning: A systematic review. *Journal of Ambient Intelligence*, *14*(6), 7761–7793.

Fridrich, J. (2009). *Steganography in digital media*. Cambridge University Press.

Fridrich, J., Goljan, M., & Du, R. (2001). Attacking the outguess steganographic system. *Information Hiding Workshop*.

Fridrich, J., & Kodovsky, J. (2012). Rich models for steganalysis of digital images. *IEEE Transactions on Information Forensics and Security*, *7*(3), 868–882.

Harris, C., & Stephens, M. (1988). A combined corner and edge detector. *Alvey Vision Conference*, 147–151.

Hetzl, S., & Mutzel, P. (2005). A graph-theoretic approach to steganography. *Communications in Computer and Information Science*.

Huynh-Thu, Q., & Ghanbari, M. (2008). Scope of validity of psnr in image/video quality assessment. *Electronics Letters*, *44*(13), 800–801. https://doi.org/10.1049/el:20080522

Johnson, N. F., & Jajodia, S. (1998). Steganography: Seeing the unseen. *IEEE Computer*, *31*(2), 26–34.

Ker, A. D. (2004). Improved detection of lsb steganography in grayscale images. *International Workshop on Information Hiding*, 97–115. https://doi.org/10.1007/978-3-540-30114-1_8

Ker, A. D. (2013). Moving steganography and steganalysis from the laboratory to the real world. *Proceedings of the First ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec '13)*, 45–58. https://doi.org/10.1145/2482513.2482530

Kharrazi, M., Sencar, H. T., & Memon, N. (2004). Image steganography: Concepts and practice. *Lecture Notes Series, Institute for Mathematical Sciences*, 35–49.

Kheddar, H., Hemis, M., Himeur, Y., Megias, D., & Amira, A. (2024). Deep learning for steganalysis: A review. *Neurocomputing*, *589*, 127528.

Luo, W., Wei, K., Li, Q., Ye, M., Tan, S., Tang, W., & Huang, J. (2024). A comprehensive survey of digital image steganography and steganalysis. *APSIPA Transactions*, *13*(1).

Mielikainen, J. (2006). Lsb matching revisited. *IEEE Signal Processing Letters*, *13*(5), 285–287. https://doi.org/10.1109/LSP.2006.870357

NITA-U. (2023). *Annual report 2022-2023*.

OpenPuff. (2024). *Openpuff steganography tool*.

Owolabi, O., & Akintola, K. (2022). Cybersecurity challenges in sub-saharan africa. *African Journal of Information Security*, *8*(2), 45–62.

Patel, R., & Sharma, S. (2021). Comprehensive survey of image steganography. *Multimedia Tools and Applications*, *80*, 25067–25108.

Petitcolas, F. A., Anderson, R. J., & Kuhn, M. G. (1999). Information hiding: A survey. *Proceedings of the IEEE*, *87*(7), 1062–1078.

Pevný, T., Bas, P., & Fridrich, J. (2010). Steganalysis by subtractive pixel adjacency matrix. *IEEE Transactions on Information Forensics and Security*, *5*(2), 215–224. https://doi.org/10.1109/TIFS.2010.2045842

Provos, N. (2001). Defending against statistical steganalysis. *10th USENIX Security Symposium*, 323–335. https://www.usenix.org/legacy/events/sec01/provos.html

Provos, N., & Honeyman, P. (2003). Hide and seek: An introduction to steganography. *IEEE Security & Privacy*, *1*(3), 32–44. https://doi.org/10.1109/MSECP.2003.1203220

Rublee, E., Rabaud, V., Konolige, K., & Bradski, G. (2011). Orb: An efficient alternative to sift or surf. *ICCV*, 2564–2571.

Schaefer, G., & Stich, M. (2003). Ucid: An uncompressed color image database. *Storage and Retrieval Methods and Applications for Multimedia 2004*, *5307*, 472–480. https://doi.org/10.1117/12.525375

Selvaraj, A., Ezhilarasan, A., Wellington, S. L. J., & Sam, A. R. (2021). Digital image steganalysis: A survey on paradigm shift from machine learning to deep learning. *IET Image Processing*, *15*(2), 504–522.

Shi, J., & Tomasi, C. (1994). Good features to track. *CVPR*, 593–600.

Singh, P., Sharma, A., & Kumar, R. (2022). Image steganography and steganalysis: A review. *International Journal of Information Security*, *21*, 123–145.

Steghide. (2024). *Steghide user manual*.

Subramanian, N., Elharrouss, O., Al-Maadeed, S., & Bouridane, A. (2021). Image steganography: A review of the recent advances. *IEEE Access*, *9*, 23409–23423.

Uganda National CERT. (2022). *National cybersecurity report*.

USC. (2024). *Usc-sipi image database*.

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, *13*(4), 600–612.

Wang, Z., Simoncelli, E. P., & Bovik, A. C. (2003). Multiscale structural similarity for image quality assessment. *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, *2*, 1398–1402. https://doi.org/10.1109/ACSSC.2003.1292216

Westfeld, A. (2001). F5—a steganographic algorithm. *Information Hiding Workshop*, 289–302.

Westfeld, A., & Pfitzmann, A. (2000). Attacks on steganographic systems. *International Workshop on Information Hiding*, 61–76. https://doi.org/10.1007/10719724_5

Xu, G., Wu, H.-Z., & Shi, Y. Q. (2016). Structural design of convolutional neural networks for steganalysis. *IEEE Signal Processing Letters*, *23*(5), 708–712. https://doi.org/10.1109/LSP.2016.2548421

Yedroudj, M., Comby, F., & Chaumont, M. (2020). Yedroudj-net: A deep learning approach for image steganalysis. *IEEE Transactions on Information Forensics and Security*, *15*, 2589–2601.

Yu, L., Zhao, Y., Ni, R., & Li, T. (2010). Improved adaptive lsb steganography based on chaos and genetic algorithm. *EURASIP Journal on Advances in Signal Processing*, *2010*, 1–6. https://doi.org/10.1155/2010/876946

Zhu, Y., Wang, X., Chen, H.-S., Salloum, R., & Kuo, C.-C. J. (2023). Green steganalyzer: A green learning approach. *APSIPA Transactions*, *12*(1).