SLSWCS
Student Workshop on Computer Science

8th Sri Lanka Student Workshop on Computer Science

SL-SWCS'25

In recognition with

IEEE
Sri Lanka Section

In collaboration with

IEEE Computational Intelligence Society
University of Jaffna Student Branch Chapter

Department of Computer Science, University of Jaffna

13 December 2025



BOOK OF ABSTRACTS

Sri Lanka Student Workshop on Computer Science

SL–SWCS'25

13 December 2025

Department of Computer Science, University of Jaffna

In recognition with



In collaboration with



Platinum Sponsor



Gold Sponsor



Silver Sponsors





Bronze Sponsors





About SL-SWCS

The Sri Lanka Student Workshop on Computer Science (SL-SWCS) is a biennial symposium established in 2011 and organised by the Department of Computer Science, Faculty of Science, University of Jaffna. The workshop is dedicated to final-year undergraduates and postgraduate research students across Sri Lanka, providing a structured forum for presenting early-stage research, receiving expert feedback, and engaging with peers and senior academics. Through its focus on foundational research training, SL–SWCS has become a nationally recognised platform for cultivating emerging talent in computer science.

The workshop format is designed to maximise scholarly interaction and support students in communicating their work effectively. All accepted submissions are featured in the poster session, which serves as the central component of the technical programme. To facilitate engagement, each participant delivers a concise one-minute spotlight presentation prior to the poster session, enabling attendees to gain an overview of the research and encouraging deeper discussion. Based on reviewer evaluations, a selected number of high-quality submissions are invited for oral presentation.

Over the years, the workshop's publication and review model has evolved to meet higher academic standards. Prior to 2023, the review process focused on the posters themselves, which were then printed in the proceedings. Beginning with SL-SWCS'23, the workshop adopted a more rigorous structure requiring the submission and review of extended abstracts. For SL-SWCS'25, this volume captures the full breadth of the event: it includes both the peer-reviewed abstracts and the corresponding posters, thereby enhancing the completeness and archival value of the workshop's published record.

Across its editions, SL–SWCS has served as an incubator for early-career researchers, a conduit for mentorship, and a catalyst for sustained collaboration within Sri Lanka's computer science community. The workshop continues to provide a vital platform for students to strengthen their research capabilities and contribute meaningfully to the country's evolving scientific landscape.



Message from the Vice Chancellor



Prof. Sivakolundu Srisatkunarajah
Vice Chancellor, University of Jaffna

It is with great pleasure that I extend my warm greetings to all participants of the SL-SWCS'25. I sincerely appreciate the proud initiative taken by the Department of Computer Science in establishing and organizing the SL-SWCS as a biennial workshop since 2011. This longstanding commitment reflects the department's dedication to advancing computer science education and fostering research excellence among students.

Over the years, this workshop has evolved into a respected forum that promotes intellectual curiosity and academic engagement. SL-SWCS provides a valuable opportunity for students to broaden their perspectives, refine their research skills, and engage with academics and professionals from both local and international institutions/ organisations. Such interactions play a crucial role in strengthening the research ecosystem of the country, helping to cultivate future leaders in science, technology, and innovation.

I commend the Department of Computer Science for its dedication to organising this event on a regular basis and for its ongoing efforts to strengthen research culture within the University. The recognition of the IEEE Sri Lanka Section and the collaboration with IEEE Computational Intelligence Society Student Branch Chapter, University of Jaffna and support received from various sponsors further enhance the stature and impact of SL-SWCS'25. Their collective contributions underscore the importance of partnerships in advancing national research initiatives.

I extend my sincere appreciation to the organising committee, reviewers, speakers, and participants whose commitment and hardwork have contributed to the successful conduct of this year's workshop. I am confident that SL-SWCS'25 will inspire new ideas, encourage meaningful collaborations, and motivate young scholars to pursue impactful research in the years ahead.

I wish all participants a productive, insightful, and enriching experience at SL-SWCS'25.

Message from the Dean



Prof. Punnaimoorthy Ravirajan
Senior Professor of Physics,
Dean, Faculty of Science, University of Jaffna

It gives me great pleasure to extend my warmest greetings to all participants of the SL-SWCS'25. Established in 2011 as the first student-focused national workshop in the field of Computer Science in Sri Lanka, SL-SWCS remains a proud and visionary initiative of the Department of Computer Science, University of Jaffna.

Since the inaugural workshop held on 8 December 2011, this event has grown steadily in both stature and impact. Over the years, SL-SWCS has provided an invaluable platform for young research students from across the country to present their work, exchange ideas, and engage with fellow researchers as well as eminent computer scientists from local and international institutions. Its interactive and collaborative format continues to inspire and nurture the next generation of scholars in Computer Science.

This year's workshop marks the eighth edition of the series, and I am delighted to note that SL-SWCS'25 has attracted more than 45 abstract submissions for review. Such enthusiasm reflects the workshop's strong reputation as a hub for intellectual inquiry, innovation, and academic networking. Through poster presentations and opportunities for constructive dialogue, SL-SWCS serves not only as an academic forum but also as a space for fostering meaningful professional relationships and future collaborations.

I commend the Department of Computer Science for its sustained commitment to organising this workshop since 2011 and for its dedication to cultivating a vibrant research culture among students. I also extend my sincere appreciation to the organising committee, reviewers, speakers, and participants whose contributions have made SL-SWCS'25 possible.

I wish all participants an enriching and inspiring experience, and I look forward to the continued growth of this important national initiative.

Message from the General Chair



Dr. (Mrs.) B. Mayurathan
General Chair/SL-SWCS'25
Head/Computer Science, Faculty of Science, University of Jaffna

It gives me immense pleasure to welcome all distinguished guests, presenters, reviewers, sponsors, and participants to the eighth Sri Lanka Student Workshop on Computer Science (SL-SWCS'25). SL-SWCS provides a dynamic and stimulating platform for young research students to engage in meaningful discussions with researchers from both local and international institutions. The workshop offers students the opportunity to present their research outcomes through oral/poster presentations in a friendly, supportive, and academically enriching environment. This friendly environment helps students share knowledge, grow their curiosity, and build useful collaborations with peers and experts.

The inaugural national workshop, SL-SWCS'11, was held on 8 December 2011. Since then, SL-SWCS has been conducted biennially and has steadily grown, attracting around 150 participants from diverse institutions across Sri Lanka. Over the years, it has earned a reputation as a respected platform for nurturing the next generation of researchers in Computer Science.

We are proud to conduct SL-SWCS'25 with the recognition of the IEEE Sri Lanka Section and in collaboration with the IEEE CIS Student Branch Chapter of the University of Jaffna. This year, we have also received generous support from our platinum (DAAD – through the DigSAL project), gold (University of Jaffna), silver (99x and WSO_2), and bronze sponsors (BOC and People's Bank), which has been instrumental in the successful organization of this workshop.

This year, we are proud to have 41 abstracts accepted for presentation, all of which were rigorously reviewed by national and international experts. I would like to extend my sincere gratitude to all the authors, reviewers, organising committee members, and

sponsors whose dedication and support have made SL-SWCS'25 possible. In addition, my special thanks go to our Vice Chancellor and the Dean of the Faculty of Science for their continued support in organizing this national event. I am confident that the workshop will provide a rewarding and inspiring experience for everyone, sparking new ideas, collaborations, and lifelong connections.

I wish all participants a successful and intellectually enriching workshop.



Message from the Editors





Prof. M. Siyamalan and Mr. S. Suthakar
Editors, SL-SWCS'25

SL-SWCS'25 gives final-year undergraduates and postgraduates the chance to showcase research, get constructive feedback from local and international experts, and build the research skills that will power Sri Lanka's future in computing.

For SL-SWCS'25, we received 49 extended-abstract submissions representing 22 institutions across Sri Lanka and overseas, including several works conducted in collaboration with international partners. After a rigorous double-blind peer-review process – with each submission evaluated by at least two domain experts – 41 abstracts were accepted for presentation.

Among the accepted submissions, 25 (60%) focus on Machine Learning, Deep Learning, and Computer Vision; 8 (20%) on Natural Language Processing; 4 (10%) on Cyber Security; 2 (5%) on Computer Networks and Mobile Computing; and 2 (5%) on Web Technologies. Overall, data-centric domains – Machine Learning, Deep Learning, Computer Vision, and NLP – constitute 33 submissions (80%), underscoring the strong national shift toward data-driven research and applications with high societal relevance.

The accepted research addresses a wide range of real-world challenges, including precision agriculture, medical diagnostics and local-language technologies, illustrating the growing relevance of computing research in everyday applications.

We extend our sincere thanks to all authors for their submissions, to the reviewers for their timely and constructive evaluations, and to the keynote speakers and organising team for their invaluable contributions. We hope these proceedings inspire further collaboration and innovation within Sri Lanka's computing community.

Table of Contents

Agenda **1**

Invited Plenary Speakers **3**

Visualization of Large Networks (Information Visualization)

– *Prof. Dr. Oliver Deussen, University of Konstanz, Germany* 3

Global Connectivity: Breaking Barriers with Technology

– *Prof. S. Vasanthapriyan, Sabaragamuwa University of Sri Lanka* 4

From Research to Reality: Building Tech Products in the Real World

– *Mr. S. Uthaiyashankar, WSO_2, Sri Lanka* 5

Abstracts and Posters **7**

Machine Learning, Deep Learning and Computer Vision **9**

Aedes Mosquito Larvae Classification Using Swin Transformers

Gangana Upasena, Somaskandan Suthakar, Thanikasalam Kokul, Sinnathamby Noble Surendran 10

Age and Gender Classification from Facial Images Using VGG16 and Canny Edge Detection

Kirushanth Uthayakumar, Kokul Thanikasalam 12

Alternating Binary SVM Strategies in a Hierarchical Framework for Multiclass Classification

Tharmaraj Thayatharsan, Amirthalingam Ramanan 14

An Efficient Hybrid Model for Early Detection of Mango Leaf Pests

Sankeetha Vamadevan, Agalya Vinoharan, Veerapathirapillai Vinoharan 16

An Evaluation Study of Transformer-Based Dehazing Models on Real World Datasets

Vishnuga Sivakumaran, Amirthalingam Ramanan, Thanikasalam Kokul 18

Attention-Enhanced Deep Learning for Diabetic Retinopathy Detection

Mahesan Natheesan, Rajendran Nirthika 20

Automated Fungi Classification with VGG16-based Feature Fusion and Contrast Enhancement
Vimalenthiran Thusani, Rajendran Nirthika 22

Automated Foundation Structure Identification from Floorplans using Deep Learning
Ravindu Deshitha, Somaskandan Suthakar, Siyamalan Manivannan 24

Comparative Analysis in Textual and Audio Modality on Cyberbullying Context
Arthika Anandapalan, Yuraj Perera, Jananie Jarachanthan 26

Deep Learning-Based Classification of Cerebellum Failure Using MRI Imaging and Ensemble Modeling
Mahendrarajah Vasavan, Luxshi Karunakaran, Chathurani Nadika Hettiarachchi 28

Detection and Classification Betel Leaf Diseases using Image Processing Techniques
H. P Isuru Lakmal, Barathy Mayurathan 30

Enhanced Prediction of Non-Communicable Diseases using Ensemble Machine Learning Approaches
Lenuga Mahendran, Banujan Kuhaneshwaran, Luxshi Karunakaran, Vadivel Abishethvarman, Senthan Prasanth, Banage. T. G. S. Kumara 32

Enhancing Paddy Leaf Disease Classification through Integrated Deep and Handcrafted Features
Jegatheesan Laavanja, Amirthalingam Ramanan 34

Enhancing Paddy Leaf Disease Classification: A Dual Stream CNN and Handcrafted Feature Approach
Tharmini Kiriharan, Amirthalingam Ramanan, Thanikasalam Kokul, Subha Fernando . 36

Evaluating Layout Detection Models For Historical Tamil Printed Texts
Suthakar Sivashanth, Kengatharaiyer Sarveswaran, Eugene Yugarajah Andrew Charles . 38

Explainable AI for Heart Disease Prediction Using CNNs
Preethy Srinevasan, Eugene Yugarajah Andrew Charles 40

Ordinal-Aware Contrastive Learning for Enhanced Diabetic Retinopathy Grading
Alagiah Suthakaran, Siyamalan Manivannan 42

Multimodal Complementary Fashion Recommendation using FashionCLIP and HBDSCAN
Nigee Hettige, Rajendran Nirthika 44

Multimodal Emotion Recognition with Emotion-Specific Cross-Modal Attention Blocks
Dharshika Gnanaseelan, Amirthalingam Ramanan 46

Potato Plant Disease Recognition using Lightweight Deep Neural Networks
Anjalika Ekanayaka, Dimuthu Rajamuni, Jananie Jarachanthan, Thuseethan Selvarajah, Nishankar Sathiyamohan 48

Recognition of Electrical Components in Hand-Drawn Schematics
S. M. Semasingha Mudiyanselage, Samantha Tharani Jeyakumar, Kesavi Nadesh 50

Segmentation of Fruit Surface Defects using Semi Supervised Learning
Saarukesan Premkumar, Siyamalan Manivannan 52

Self-Supervised Ultrasound Image Denoising Using Noise2Void and Noisier2Noise
Roy Nesarajah Viththagan, Siyamalan Manivannan, Thanikasalam Kokul 54

Transformer based Semi-Supervised Learning for Skin Lesion Classification
Jathursika Velummayilum, Siyamalan Manivannan 56

Transformer Tracking with Controlled Attention for Background Suppression
Janani Kugarajeevan, Thanikasalam Kokul, Amirthalingam Ramanan, Subha Fernando . 58

Computer Networks, and Mobile Computing **61**

Bridging the GPS Divide: A Dual-Mode Real-Time Tracking App for Highway Buses in Sri Lanka
Tharindu Rajasinghe, S. A. A. H. Samaraweera 62

Decoding Phrasal Verbs with Context-Aware BERT Representations
Sujair Ibrahim, Abishethvarman Vadivel, Kumara T. G. S. Banage 64

Cybersecurity **67**

Adaptive Adversarial Defense for Network Intrusion Detection Systems Using Dynamic GAN Ensemble
Vieronicka Kanesamoorthy, Shriparen Sriskandarajah 68

ARP Spoofing Prevention in SDN Using Host Agents and Cryptographically Verified DHCP IP-MAC Bindings
Sankalpa Fernando, Shriparen Sriskandarajah 70

Behavioural Analysis for Suspicious Transaction in Bitcoin Network
Raseswaran Kanistan, Jeyakumar Samantha Tharani 72

Zero Trust-based Network Traffic Anomaly Detection using Rule-based Approaches
S. W. Chamindu Lakshan Wijerathna, Shriparen Sriskandarajah 74

Natural Language Processing **77**

Sinhala Handwritten Optical Character Recognition
P. N. Nipun Sadeepa Weerasinghe, Kengatharaiyer Sarveswaran 78

Analysing Ambiguity in Anaphoric Reference:
A Rule-Based Framework and a Proposal for Tamil
Lukas Beiske 80

Building A Baseline Sinhala Text-To-Speech System
Vigneswaran Wanith Nirman, Ahrane Mahaganapathy, Kengatharaiyer Sarveswaran . . . 82

Dialect Divergence in Sri Lankan Tamil: A Corpus-Based
Computational Analysis
Gobalakrishnar Nilusana, Kengathaiyar Sarveswaran 84

Improving Label Consistency in Tamil Emotion Datasets through a Multi-Method
Pipeline
Sankavi Mohan, Andrew Charles Eugene Yugarajah 86

Modeling Classical Nahuatl Verbal Nuclear Clauses with
a Finite State Transducer
Francesca Hoffmeier, Natalia Villanueva 88

Perplexity-Based Evaluation of Open-Source LLM Adaptation for Tamil
Jeyarajalingam Varsha, Tharmaraj Thayatharsan, Kengatharaiyer Sarveswaran 90

Target Identification in Tamil Hate Speech
Sathasivam Nerujan, Kengatharaiyar Sarveswaran 92

Web Technologies **95**

A Performance Comparison Between Rest API and GraphQL in
Microservice Architecture
Thulashika Balachandran, Kengatharaiyar Sarveswaran, Samantha Tharani Jeyakumar . 96

CertiRide: A Tamper-Free Vehicle Service Record Ledger
Naveen Hettiwaththa, Samantha Jeyakumar, Haritha Senadheera, Hiroshan Hettiarachchi,
Sandaru Gunarathna, Tharushi Adhikari 98

List of Reviewers **100**

Organising Committee **102**

Achievers of Past SL-SWCS **103**

Sponsors **105**

Agenda

08:00 – 08:30	**Registration**
08:30 – 08:35	**Welcome Address** Dr. (Mrs.) B. Mayurathan General Chair
08:35 – 08:45	**Vice Chancellor's Address** Prof. S. Srisatkunarajah Vice-Chancellor, University of Jaffna
08:45 – 08:55	**Dean's Address** Prof. P. Ravirajan Dean, Faculty of Science, University of Jaffna
09:00 – 09:55	**Keynote Address** Prof. Dr. Oliver Deussen University of Konstanz, Germany
10:00 – 10:55	**Keynote Address** Mr. S. Uthaiyashankar Senior VP & GM – Integration, WSO_2, Sri Lanka
11:00 – 11:30	**Tea Break**
11:30 – 12:30	**Keynote Address** Prof. S. Vasanthapriyan Professor and Dean, Faculty of Computing, Sabaragamuwa University of Sri Lanka
12:30 – 12:45	**Oral Presentation** Perplexity-based Evaluation of Open-Source LLM Adaptation for Tamil Ms. Jeyarajalingam Varsha

12:45 – 01:00	**Oral Presentation** Alternating Binary SVM Strategies in a Hierarchical Framework for Multiclass Classification Mr. Tharmaraj Thayatharsan
01:00 – 02:00	**Lunch**
02:00 – 02:15	**Oral Presentation** Self-Supervised Ultrasound Image Denoising Using Noise2Void and Noisier2Noise Mr. Roy Nesarajah Viththagan
02:15 – 02:30	**Oral Presentation** Analysing Ambiguity in Anaphoric Reference: A Rule-Based Framework and a Proposal for Tamil Mr. Lukas Beiske
02:30 – 02:45	**Oral Presentation** Ordinal-Aware Contrastive Learning for Enhanced Diabetic Retinopathy Grading Mr. Alagiah Suthakaran
02:45 – 03:00	**Oral Presentation** Multimodal Emotion Recognition with Emotion-Specific Cross-Modal Attention Blocks Ms. Dharshika Gnanaseelan
03:00 – 03:15	**Spotlight Session**
03:15 – 04:45	**Poster Session with Tea Break**
04:45 – 05:00	**Awarding and Closing Ceremony**

Invited Plenary Speakers

Keynote 1

Visualization of Large Networks (Information Visualization)

Prof. Dr. Oliver Deussen
Professor for Visual Computing
University of Konstanz,
Germany



Abstract of the Talk: Inevitably, the projection of most graph structures on two-dimensional screens will create errors and therefore visually wrong impressions. In the past, two types of methods have been developed to minimize such projection errors and create layouts in a visually pleasing way. The first group of methods, force-directed layouts, interpret the links of a graph as physical springs, while stress-based methods minimize an energy function that aims to map graph distances faithfully. With Prof. Yunhai Wang I developed a unified description that allows to create optimal parameters for both modeling types and even to specify new, better methods for most graphs. Adding vector-based constraints allows systems to render graphs with different layouts and under varying perspectives. We show that deviating from all kinds of physical metaphors is a good idea to create even better projection methods.

Biographical Profile: Prof. Deussen graduated at Karlsruhe Institute of Technology and is a professor at University of Konstanz, one of the Excellence Universities in Germany. 2008-2020 he was a visiting professor at the Chinese Academy of Science in Shenzhen (SIAT). 2019-2020 he served as President of the Eurographics Association. His areas of interest are Information Visualization, non-photorealistic rendering techniques, sampling theory. He also contributed papers to geometry processing and image-based modeling methods. Currently he is one of the speakers of the Excellence Cluster "Centre for the Advanced Study of Collective Behaviour" at University of Konstanz.

Keynote 2

Global Connectivity: Breaking Barriers with Technology – Building Inclusive Networks that Unite Communities Worldwide

Prof. S. Vasanthapriyan
Professor and Dean,
Faculty of Computing,
Sabaragamuwa University of Sri Lanka
Chair, IEEE Sri Lanka Section



Abstract of the Talk: In an era defined by rapid technological evolution, the power of connectivity transcends geographical, cultural, and socio-economic boundaries. Today, technology serves as the cornerstone for fostering inclusivity, collaboration, and innovation – bridging divides and enabling communities to thrive together. As the world becomes increasingly interconnected, it is imperative to ensure that this digital transformation benefits all, leaving no one behind. The IEEE, as a global community of engineers, researchers, and innovators, continues to play a vital role in advancing these frontiers – encouraging cross-border collaboration, nurturing young minds, and supporting initiatives that strengthen global networks of knowledge and impact. This talk explores how emerging technologies – from next-generation communication networks to artificial intelligence and cloud-driven platforms – are redefining global collaboration and empowering diverse communities. It highlights real-world initiatives and innovations that demonstrate how inclusive technological ecosystems can drive equitable access to knowledge, opportunity, and sustainable growth.

Biographical Profile: Professor S. Vasanthapriyan, founding Dean of the Faculty of Computing at Sabaragamuwa University (2023), has made remarkable contributions to ICT and academia. An award-winning scholar from Jaffna, he holds degrees from Peradeniya, Eindhoven, and Wuhan. With over 19 years of experience, he served in key leadership roles including Head of CIS, Coordinator of the Faculty of Computing, and creator of the Software Engineering degree programme. A Senior Member of IEEE, he received major recognitions such as "ICT Educator of the Year 2022" and "ICT Leader of the Year 2023." In 2025, he became Chair of IEEE Sri Lanka and joined IEEE TPIC.

Keynote 3

From Research to Reality: Building Tech Products in the Real World

Mr. S. Uthaiyashankar
Senior Vice President & General Manager - Integration, WSO_2,
Sri Lanka.



Abstract of the Talk: What does it take to turn academic research into a real-world tech product? This session aims to answer that question through a dynamic conversation led by an industry expert with firsthand experience in bridging theory and practice. Designed specifically for student researchers across the computing spectrum, including artificial intelligence, cybersecurity, computer networks, robotics, data engineering, software, and hardware systems, this session will provide insight into how ideas nurtured in university labs can evolve into impactful, scalable technologies. Participants will gain a behind-the-scenes look at the development process: how companies align research with user needs, what trade-offs are made during product design, and how technical teams adapt when shifting from proof-of-concept to production-ready solutions. You'll also explore the challenges of working within tight real-world constraints like time, budget, and platform limitations. Most importantly, this session is built to inspire. It will encourage young researchers to think beyond academic papers, to see their work as a foundation for innovation that can improve lives, solve local and global problems, and drive career-defining opportunities. Whether you're working on a neural network, a mobile app, a network protocol, or a robotics system, you'll walk away with a clearer vision of how to move from research to relevance.

Biographical Profile: Mr. Selvaratnam Uthaiyashankar has over 17 years of experience in the technology space spanning multiple functions, including web services, cloud deployments, solutions architecture, and engineering. He started his career at Millennium IT in a project that involved telecommunication billing systems. He joined WSO_2 in 2007 and played many roles before taking on his current role of overseeing engineering.

His previous work at WSO_2 included contributing to web services implementation for C/C++/PHP/Perl, managing StratosLive deployment in the Cloud, and working directly with customers as part of the solutions architecture teams.

In his current role, Shankar oversees WSO_2's integration business unit and is responsible for driving product and go-to-market strategies. He holds a BSc Engineering in Computer Science and Engineering from the University of Moratuwa, Sri Lanka.

When he's not absorbed in product development in the technology world, you will find Shankar buried in books on various topics and genres. He claims he could spend hours or days doing nothing but reading.

Abstracts and Posters

Machine Learning, Deep Learning and Computer Vision

Aedes Mosquito Larvae Classification Using Swin Transformers

Gangana Upasena[1,*], Somaskandan Suthakar[1], Thanikasalam Kokul[1], Sinnathamby Noble Surendran[2]

[1]Department of Computer Science, University of Jaffna, Sri Lanka,
[2]Department of Zoology, University of Jaffna, Sri Lanka

[*]g.dilshanupasena@gmail.com

Dengue fever is a major public health concern in Sri Lanka, where early identification of *Aedes aegypti* and *Aedes albopictus* larvae is essential for effective vector control. Manual larval identification is time-consuming, expertise-dependent, and often unreliable under field conditions. This study presents an automated deep learning approach for classifying mosquito larvae using Swin Transformer architectures. A dataset of 5,939 images was curated by combining images captured by the researchers with publicly available collections from platforms such as Mendeley, Kaggle, and Roboflow, and annotated into three classes: Aedes aegypti, Aedes albopictus, and non-Aedes (comprising other common mosquito genera). Images were pre-processed with augmentation techniques, and three Swin Transformer variants Swin-T, Swin-S, and Swin-B were fine-tuned using a customized Linear $\rightarrow$ ReLU $\rightarrow$ Dropout $\rightarrow$ Linear head, AdamW optimization, StepLR scheduling, and Label Smoothing Cross-Entropy. Dataset refinement, guided by Gradient-weighted Class Activation Mapping (Grad-CAM), enabled the removal of ambiguous samples. Swin-B achieved the highest accuracy of 96.86%, outperforming conventional convolutional neural network (CNN) models reported in the literature. The hierarchical shifted-window attention mechanism of Swin Transformers effectively captured fine-grained morphological features, including siphon shape, head structure, and thorax bristles, supporting reliable species-level differentiation. These findings demonstrate the potential of transformer-based models for dengue vector surveillance.

Keywords: Aedes aegypti, Aedes albopictus, deep learning, Swin Transformer, image classification

Aedes Mosquito Larvae Classification Using Swin Transformers.

G. D. Upasena[1*], S. Suthakar[1], T. Kokul[1], S. N. Surendran[2]

1Department of Computer Science, University of Jaffna, Sri Lanka,
2Department of Zoology, University of Jaffna, Sri Lanka
Email: g.dilshanupasena@gmail.com

Introduction

❖ **Problem Statement:**

Dengue fever: A major global health threat, spread by Aedes aegypti and Aedes albopictus.

Vector control is key: Targeting larvae is most effective.

Current methods: Manual identification is slow, error-prone, and requires expertise.

❖ **Research Gap:**

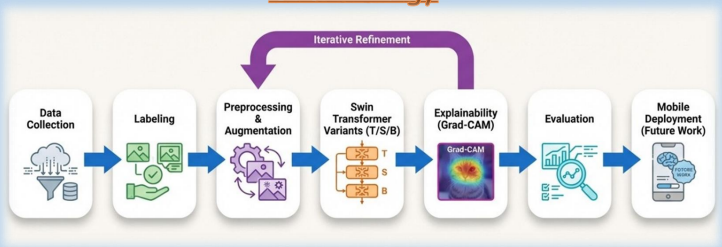
Limited studies on differentiating A. aegypti vs. A. albopictus.

Lack of research in Sri Lanka for this specific task.

Previous models rely on dead larvae, missing real-world variation

❖ **Our Goal:** To develop an automated, accurate, and accessible method for classifying dengue vector larvae using deep learning.

Methodology



Dataset:

Total Images: 5,939

Classes: Aedes aegypti (1,854), Aedes albopictus (1,932), Non-Aedes (2,153)

Image Preprocessing & Augmentation:Cropping, Flipping, Color Jittering (brightness, contrast, saturation, hue), Random Erasing.

Classification Model:

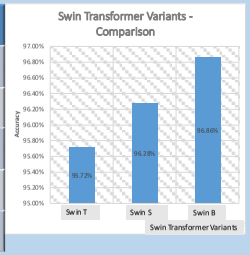
Customization: Replaced the default head with a Linear-ReLU-Dropout-Linear sequence.

Training: 20 epochs, AdamW optimizer, StepLR scheduler, Label Smoothing Cross-Entropy loss. Fully fine-tuned.

Results & Discussion

Class	Precision	Recall	F1 Score	Support
Aedes Aegypti	0.9539	0.9676	0.9607	278
Aedes Albopictus	0.9654	0.9621	0.9637	290
Non Aedes	1.0000	0.9907	0.9953	323
Macro Avg	0.9731	0.9735	0.9733	891
Weighted Avg	0.9744	0.9742	0.9742	891

Table 1 – Per Class Matrix



❖ **Key Findings:**

High accuracy achieved despite dataset variability (diverse sources, live/dead larvae, uncontrolled conditions).

Previous models rely on dead larvae, missing real-world variation.

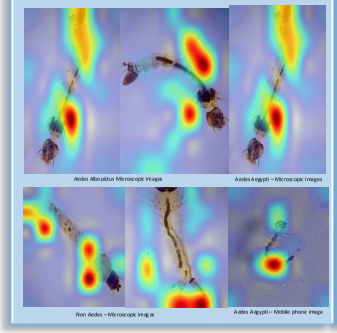
Model variant	Accuracy	Precision	F1 Score
Swin - T	95.42%	95.36%	95.33%
Swin - S	96.28%	96.21%	96.19%
Swin - B	96.86%	96.76%	96.75%

Table 2 – Confusion Matrix

❖ **Discussion:**

- Why the high accuracy? Swin Transformer's hierarchical attention effectively captures local and global features.
- Addressing initial challenges: Iterative data refinement was crucial.
- Real-world applicability: Robustness demonstrated by performance on varied "in-the-wild" images.

GRAD – CAM Images



Data (Images)



Conclusions

❖ Successfully developed a high-accuracy (96.86%) automated classifier for Aedes larvae using Swin Transformers.

❖ Demonstrated the importance of data refinement and explainability for robust models.

❖ Pioneered Swin Transformer use and live larvae classification for this task.

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)

Poster ID: P01

Age and Gender Classification from Facial Images Using VGG16 and Canny Edge Detection

Kirushanth Uthayakumar*, Kokul Thanikasalam

Department of Computer Science, Faculty of Science, University of Jaffna, Jaffna, Sri Lanka

*2020sp200@univ.jfn.ac.lk

Estimating age and gender simultaneously from a single facial image is important for many applications. Although previous deep learning approaches have achieved strong performance in gender classification, their performance in age estimation remains limited, as they often fail to capture the shape and structural features crucial for accurate prediction. To overcome this limitation, this study proposes a dual-branch CNN model for joint age and gender estimation. The proposed approach takes a raw facial image and generates its corresponding edge map using the Canny edge detection algorithm. Features from the raw image are extracted using a pre-trained VGG-16 network, while those from the edge map are obtained through a customised CNN. The inclusion of edge-based features captures fine structural details, whereas channel and spatial attention modules enhance discriminative learning across both branches. Finally, the features from both branches are concatenated and processed through a series of fully connected layers for classification. In this work, age estimation is treated as a multi-class classification problem rather than regression, enabling group-wise age prediction. The proposed approach was trained and evaluated on the publicly available UTKFace dataset, attaining classification accuracies of 87.88% for gender and 60.00% for age prediction. Compared to state-of-the-art methods, our approach achieves competitive accuracy while using less than half the training samples, underscoring its efficiency and novelty.

Keywords: Age and gender estimation, CNN, edge detection

Age and Gender Classification from Facial Images Using VGG16 and Canny Edge Detection

Kirushanrth, U., and Kokul, T.
Department of Computer Science, University of Jaffna, Jaffna, Sri Lanka
2020sp200@univ.jfn.ac.lk, kokul@univ.jfn.ac.lk

Introduction

Estimating age group and gender simultaneously from facial images is important for a wide range of applications, including healthcare, security, and marketing. However, in many existing approaches, age and gender are predicted individually, which requires more resources and time, and may potentially lead to inconsistencies. This research aims to develop a deep learning model to predict both age group and gender from a single image, offering greater efficiency and consistency. Although a few Convolutional Neural Network (CNN) based approaches [1–10] have been proposed to detect age and gender jointly, their performance in age estimation remains considerably limited, as age prediction relies more heavily on structural and shape-related cues such as wrinkles, and facial contours, which cannot be effectively extracted from raw images alone. To overcome these limitations, this study employs a dual-branch CNN for joint age and gender estimation. The proposed model processes both the raw facial image and its corresponding edge map, enabling the network to capture low-level textural information as well as high-frequency structural details.

Dataset

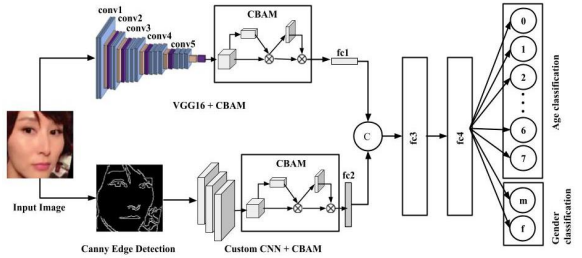


UTKFace dataset

Attribute	Details
Dataset Name	UTKFace (subset)
Total Images	12,756 (from ~23,000)
Resolution	128x128 pixels (RGB)
Age Range	8 groups: 0–2, 4–6, 8–12, 15–20, 25–32, 38–43, 48–53, 60–100
Gender Categories	Male (0), Female (1)
Partitioning	Train set: 9,600 images, Validation set: 1,552 images, Test set: 1,604 images

- **Image augmentation techniques**: rotation, zoom, flipping, brightness adjustment.

Methodology



Results and Discussion

Approach	Test Accuracy	
	Gender	Age
Akgül, İ., et al. [10]	90.50%	72.00%
Singhla, D., et al. [7]	93.82%	75.69%
Kalpana, R., et al. [5]	89.00%	78.00%
Proposed Approach	**87.88%**	**60.00%**

Table 1

Models	Test Accuracy	
	Gender	Age
ResNet50	55.31%	48.19%
EfficientNetV2S	57.19%	28.19%
EfficientNetV2M	62.25%	25.56%
VGG16	85.93%	60.43%
VGG16+CBAM	87.06%	56.87%
VGG16+CBAM+ Canny Edge Map Branch	**87.88%**	**60.00%**

Table 2

Conclusion

In conclusion, this study introduced a novel CNN architecture for joint age and gender classification. The proposed dual-branch design, which integrates edge map features with raw image features, together with channel and spatial attention modules, enhanced the model's capability. Experimental results on the UTKFace dataset demonstrated that the proposed approach achieves competitive performance, highlighting its effectiveness for age and gender classification tasks while attaining these results using less than half the training samples employed by other state-of-the-art methods.

References

[1] R. Thaneeshan, K. Thanikasalam, and A. Pinidiyaarachchi, "Gender and Age Estimation From Facial Images using Deep Learning," 2022 7th International Conference on Information Technology Research (ICITR), Moratuwa, Sri Lanka, 2022, pp. 1-6, doi: 10.1109/ICITR57877.2022.9993277.
[2] Y. Khurana, S. Gupta, A. Sood, D. Gupta, K. Sagar, and Shivani, "Leveraging Computer Vision for Gender, Age, and Ethnicity Prediction Using Deep CNN," in Data Science and Artificial Intelligence for Digital Healthcare, Springer, Cham, 2024, pp. [Page Numbers], doi: 10.1007/978-3-031-56818-3_2.
[3] M. M. Iqbal, A. Rukhsar, and S. K. Baliarsingh, "A CNN-based Prediction Model for Age, Gender, and Ethnicity Using Facial Images," 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT), Delhi, India, 2023, pp. 1-7, doi: 10.1109/ICCCNT56998.2023.10306826.
[4] A. Uniyal, Pooja, H. Narang, S. Vats, V. Sharma, and S. P. Yadav, "Employing DL-Based Algorithm for Gender and Age Identification," 2024 IEEE International Conference on Computing, Power and Communication Technologies (IC2PCT), Greater Noida, India, 2024, pp. 1459-1464, doi: 10.1109/IC2PCT60090.2024.10486667.
[5] R. Kalpana, D. Deepika, A. Kavya, P. Himabindhu, and S. Kethavi, "Facial Age and Gender Prediction using Deep Learning," 2024 10th International Conference on Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2024, pp. 1571-1576, doi: 10.1109/ICCSP60870.2024.10543983.
[6] T. P. Kancharlapalli and P. Dwivedi, "A Novel Approach for Age and Gender Detection using Deep Convolution Neural Network," 2023 10th International Conference on Computing for Sustainable Global Development (INDIACom), New Delhi, India, 2023, pp. 873-878.
[7] D. Singhla, N. Aggarwal, and N. Gupta, "Inferencing Age and Gender from Facial Images," 2024 International Conference on Computational Intelligence and Computing Applications (ICCICA), Samalkha, India, 2024, pp. 418-423, doi: 10.1109/ICCICA60014.2024.10585045.
[8] A. Kanwar and K. D. Singh, "Prediction of Age, Gender, and Ethnicity Using CNN and Facial Images in Real-Time," 2023 IEEE World Conference on Applied Intelligence and Computing (AIC), Sonbhadra, India, 2023, pp. 668-674, doi: 10.1109/AIC57670.2023.10263824.
[9] D. B. Kinthadi, A. Burugu, A. Rumandla, and S. S, "Prediction of Human Age and Gender Using Deep Learning For Smart Security Systems," 2023 International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE), Ballar, India, 2023, pp. 1-6, doi: 10.1109/ICDCECE57866.2023.10151119.
[10] İ. Akgül, "Deep convolutional neural networks for age and gender estimation using an imbalanced dataset of human face images," Neural Comput. & Applic., vol. 36, pp. 21839–21858, 2024, doi: 10.1007/s00521-024-10390-0.

Alternating Binary SVM Strategies in a Hierarchical Framework for Multiclass Classification

Tharmaraj Thayatharsan*, Amirthalingam Ramanan

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*thayatharsan70@gmail.com

Support Vector Machines (SVMs) are widely used in machine learning due to their strong theoretical foundations and kernel flexibility, but extending them to multiclass problems remains challenging. Traditional approaches like One-Versus-All (OVA) and One-Versus-One (OVO) suffer from overlapping boundaries and quadratic classifier growth, respectively, while hierarchical methods such as Directed Acyclic Graph (DAG) and Unbalanced Decision Tree (UDT) struggle with error propagation due to its final decision being dependent on the last classification at the leaf node, and imbalanced datasets at decision nodes, respectively. This work proposes A-OV-AO (Alternating OVA and OVO) SVMs, a novel hybrid decision tree architecture that alternates between OVA classifiers at higher levels for separating distinguishable classes and OVO classifiers at deeper levels for refined pairwise classification. The method incorporates dynamic reordering to prioritize confident decisions and reduce error propagation. Experiments on six UCI benchmark datasets (Iris, Wine, Glass, Isolet, DNA, and Vowel) employed RBF kernels with comprehensive hyperparameter tuning via grid search, where C ranges from 2^{12} to 2^{-2} and gamma ranges from 2^{4} to 2^{-10}. Results demonstrate that A-OV-AO achieves competitive or superior accuracy compared to baseline methods, matching perfect classification on Iris (100%), reaching 98.15% on Wine, 67.69% on Glass, 96.41% on Isolet, 95.51% on DNA, and 96.30% on Vowel. The proposed architecture balances OVO's accuracy with OVA's computational simplicity, requiring fewer classifiers than OVO while mitigating imbalances inherent in hierarchical approaches. This contribution demonstrates an efficient, scalable framework for multiclass SVM classification with improved adaptability for real-world applications.

Keywords: Support Vector Machines, Multiclass Classification, Decision Trees, Hybrid Methods, Machine Learning

Alternating Binary SVM Strategies in a Hierarchical Framework for Multiclass Classification

Thayatharsan, T. and Ramanan, A.

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka
thayatharsan70@gmail.com, a.ramanan@univ.jfn.ac.lk

Introduction

- Multiclass classification is challenging for Support Vector Machines (SVMs), despite their strong theoretical basis.
- Drawbacks of Existing Methods:
  - OVA: Suffers from overlapping decision boundaries.
  - OVO: Causes quadratic growth in the number of classifiers.
  - Hierarchical:
    - DAG: suffers from error propagation due to its final decision.
    - UDT: struggles with imbalanced datasets at decision nodes.
- Proposed Solution: We propose A-OV-AO (Alternating OVA and OVO) SVMs, a novel hybrid decision tree architecture.

Objectives

- **Analyse** the performance and complexity of existing SVM multiclass classification strategies.
- **Propose** A-OV-AO, a novel alternating OVA-OVO hierarchical architecture incorporating **dynamic reordering**.
- **Validate** the proposed framework's accuracy and efficiency using hyperparameter tuning on benchmark datasets.

Literature Review

Table 1: Existing SVM-based multiclass classification methods for N classes

Methods	Classification Strategy	#Classifiers	
		Training	Testing
OVA [1]	Winner-takes all	N	N
OVO [1]	Max-wins voting	N(N-1)/2	N(N-1)/2
DAG [2]	Leave-one out	N(N-1)/2	N-1
UDT [3]	Knock-out	N!/2	Best case: 1 Worst case: N-1
Ensemble SVM [4]	Averaging multiple tree outputs	N!/2	k*(N-1)
LCDG D-Tree [5]	Centroid distance and grouping	N-1	Best case: 1 Worst case: N-1

Methodology

The proposed A-OV-AO (Alternating OVA and OVO) SVMs architecture builds a decision tree alternating OVA and OVO classifiers.

- Top-Level Classifiers: OVA classifiers separate easily distinguishable classes based on separability metrics.
- Alternating Classifier Levels: Subsequent levels alternate to OVO for pairwise refinement among remaining classes.
- Dynamic Reordering: shifts low-confidence classifiers lower, reducing error propagation.

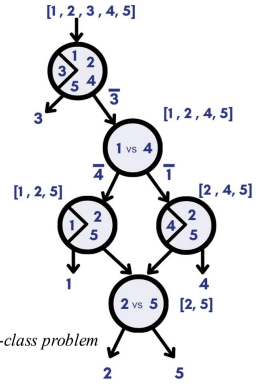


Figure 1: Proposed A-OV-AO hybrid decision tree structure for a 5-class problem

Experimental Setup and Test Results

- Kernel: Linear, RBF with hyper parameter tuning (C, gamma) via grid search, where *C* ranges from 2^{12} to 2^{-2} and gamma ranges from 2^{4} to 2^{-10}.
- The RBF kernel significantly outperformed the linear kernel, delivering much stronger results.

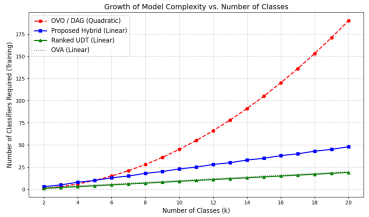


Figure 2: Comparison of classifier count (model complexity) vs. the number of classes for various SVM methods

Table 2: UCI benchmark datasets statistics used in the experiments

Dataset	Classes	Attributes	Instances	Train	Test
Iris	3	4	150	105	45
Wine	3	13	178	125	53
Glass	6	9	214	150	64
Vowel	11	10	528	370	158
DNA	3	180	3,186	2,230	956
Isolet	26	617	7,797	5,458	2,339

Table 3: Performance comparison of SVMs using RBF Kernel

Dataset	OVO		OVA		DAG		UDT		Proposed Method	
	C, γ	Rate(%)	C, γ	Rate(%)	C, γ	Rate(%)	C, γ	Rate(%)	C, γ	Rate(%)
Iris	$2^{12}, 2^{-3}$	**100.0**	$2^{12}, 2^{-4}$	**100.0**	$2^{7}, 2^{1}$	**100.0**	$2^{8}, 2^{1}$	**100.0**	$2^{8}, 2^{1}$	**100.0**
Glass	$2^{4}, 2^{0}$	66.15	$2^{4}, 2^{-3}$	69.23	$2^{9}, 2^{-1}$	67.69	$2^{7}, 2^{0}$	60.0	$2^{8}, 2^{-1}$	**70.77**
Wine	$2^{12}, 2^{0}$	96.30	$2^{12}, 2^{-8}$	**98.15**	$2^{11}, 2^{-9}$	**98.15**	$2^{10}, 2^{0}$	96.30	$2^{11}, 2^{-3}$	**98.15**
Isolet	$2^{12}, 2^{-8}$	96.15	$2^{12}, 2^{-7}$	96.67	$2^{8}, 2^{-8}$	**96.92**	$2^{3}, 2^{-9}$	95.38	$2^{4}, 2^{-8}$	96.41
DNA	$2^{11}, 2^{-2}$	94.88	$2^{11}, 2^{2}$	95.30	$2^{12}, 2^{-2}$	94.98	$2^{10}, 2^{2}$	95.40	$2^{12}, 2^{2}$	**95.51**
Vowel	$2^{12}, 2^{2}$	96.97	$2^{12}, 2^{2}$	**97.31**	$2^{12}, 2^{1}$	96.63	$2^{9}, 2^{0}$	96.30	$2^{7}, 2^{2}$	96.30

Conclusion

- The research introduces a novel architecture SVM decision tree that alternates between OVA and OVO classifiers in a hierarchical decision tree.
- Experiments showed that the A-OV-AO architecture achieves competitive or slightly better accuracy than DAG and UDT methods, while requiring fewer classifiers than OVO, thus proving to be an efficient and scalable approach for multiclass learning.

References

[1] N. Cristianini and J. Shawe-Taylor, "An introduction to Support Vector Machines and other kernel-based learning methods". Cambridge University Press, March 2000.
[2] J. Platt, N. Cristianini, and J. Shawe-Taylor, "Large margin DAGs for multiclass classification," in Advances in Neural Information Processing Systems, vol. 12, pp. 547–553, 2000.
[3] A. Ramanan, S. Suppharangsan, and M. Niranjan, "Unbalanced decision trees for multi-class classification," in Proceedings of the Second IEEE International Conference on Industrial and Information Systems, pp. 291–294, 2007.
[4] T. Phillips and W. Abdulla, "Developing a new ensemble approach with multi-class SVMs for Manuka honey quality classification," Applied Soft Computing, vol. 111, p. 107710, 2021.
[5] P. Kantavat, P. Songsiri, and B. Kijsirikul, "Efficient decision trees for multi-class support vector machines using large centroid distance grouping," Engineering Journal, vol. 26, no. 5, pp. 13–23, 2022.
[6] Dua, D., & Graff, C. (2019). UCI Machine Learning Repository. University of California, Irvine, School of Information and Computer Sciences. Retrieved from: http://archive.ics.uci.edu/ml

An Efficient Hybrid Model for Early Detection of Mango Leaf Pests

Sankeetha Vamadevan[1], Agalya Vinoharan[2], Veerapathirapillai Vinoharan[1,*]

[1]Department of ICT, Faculty of Technological Studies, University of Vavuniya, Sri Lanka,
[2]Department of Research and Development, National Institute of Education, Sri Lanka

[*]v.vinoharan@vau.ac.lk

Mango is a vital cash crop widely cultivated across tropical and subtropical regions, including Sri Lanka, where it contributes significantly to agricultural income and rural livelihoods. However, mango cultivation is often hindered by pest infestations, which cause substantial yield losses and reduce fruit quality, adversely affecting farmers' economic stability. Traditional pest detection relies on manual inspection, a labor-intensive, time consuming process prone to human error, especially on large-scale plantations. To overcome these limitations, this research proposes an efficient automated pest identification system specifically designed for mango cultivation. The framework employs a hybrid approach that integrates advanced deep learning models, VGG19, ResNet50, DenseNet121, MobileNetV2, InceptionV3, and UNet, for effective feature extraction, combined with machine learning classifiers such as Support Vector Machines (SVM), Random Forest, and Logistic Regression for accurate pest classification. The system is evaluated using 200 images per class across 10 classes in the*MangoPestClassification* dataset, which contains real-world images of healthy and pest-affected mango leaves captured under diverse environmental conditions, reflecting practical field scenarios. Experimental results show that Logistic Regression, trained on a combined feature set derived from ResNet50 features extracted after the fully connected layer, achieved the highest test accuracy of 97%. This performance surpasses all other feature-classifier combinations, demonstrating the robustness and strong generalization capability of Logistic Regression in distinguishing pest categories. Overall, the proposed framework offers a scalable and effective solution for early pest detection, enabling timely interventions, reducing dependence on manual inspection, and supporting technology-driven precision agriculture in mango cultivation.

Keywords: mango pest detection, machine learning, deep learning, feature extraction, precision agriculture

An Efficient Hybrid Model for Early Detection of Mango Leaf Pests

V. Sankeetha[1], V. Agalya[2] and V. Vinoharan[1]

[1]*Department of ICT, Faculty of Technological Studies, University of Vavuniya, Sri Lanka*
[2]*Department of Research and Development, National Institute of Education, Sri Lanka*

Introduction

The agricultural sector in Sri Lanka, like many other parts of the world, faces persistent challenges from pest infestations in key crops such as mangoes [1]. Early and accurate identification of these pests is paramount for effective disease management, minimizing crop damage, and ensuring food security [2]. Traditional pest monitoring techniques are often insufficient for the scale and speed required in modern agriculture [3]. The advent of artificial intelligence, particularly deep learning and machine learning, offers unprecedented opportunities to automate and enhance these processes [4]. Building upon the existing efforts in expert systems for pest identification, this study seeks to integrate state-of-the-art computational methods to create a highly accurate and efficient pest detection framework specifically tailored for mango crops in the Sri Lankan context.

Methodology

This study utilized a dataset of over 20,000 images collected from Sri Lankan mango farms, representing 10 distinct pest classes. For experimentation, a balanced subset of 200 randomly selected images per class was used, totalling 2,000 images, to train and validate the models.

- The methodology followed a two- stage approach: feature extraction and classification.
  - In the feature extraction stage, pre-trained deep learning models were used to generate high-dimensional representations of mango pest images.

Methodology...

- The architectures included: VGG19, ResNet50, MobileNetV2, DenseNet121, InceptionV3, and UNet, each offering unique strengths, such as deep hierarchical features, lightweight efficiency, dense connectivity, and multi-scale representation through inception modules and factorized
- In the classification stage, the extracted features were input into traditional machine learning algorithms, namely Support Vector Machine (SVM), Random Forest, and Logistic Regression.

Results

- Table 1 summarizes the classification accuracy of the different feature extractor–classifier combinations evaluated for mango pest detection.
- For all experiments, the dataset was divided into an 80:20 split, with 80% of the images used for training and the remaining 20% reserved for testing to ensure reliable evaluation of model performance.

Table 1: Classification Accuracy (%) of Different Feature Extractor–Classifier Combinations for Mango Pest Detection

Feature	SVM(Linear)	SVM(RBF)	Random forest	Logistic Regression
VGG19	92.50	93.75	87.00	92.50
RestNet50	96.00	95.25	91.00	**97.00**
DenseNet121	94.75	95.75	94.50	94.75
MobileNetV2	95.00	95.50	90.00	94.75
UNET	93.25	93.50	89.25	95.00
InceptionV3	96.00	96.25	90.00	95.50

- As shown in Table 1, InceptionV3 combined with SVM (RBF) achieved an accuracy of 96.25%, while ResNet50 features paired with Logistic Regression attained the highest accuracy of 97.00%, illustrating that robust deep feature representations can enable even simple classifiers to perform exceptionally well.

Conclusions

This study developed a hybrid framework for automated mango pest detection using deep learning feature extractors and machine learning classifiers. The framework achieved up to 97% accuracy with ResNet50 features and Logistic Regression, demonstrating that robust deep features enable even simple classifiers to perform effectively. The system can support mango farmers by enabling early and accurate pest identification, reducing crop losses, and promoting sustainable agriculture. Future work will expand the dataset and explore real-time field deployment.

References

1. Wijekoon, W.M.C.D., Ganehiarachchi, G.A.S.M. and Vidanage, S.P., The variation of abundance and fruit damage of Bactrocera dorsalis on two commercial varieties of Mangifera indica in main Bio- climatic Zones of Sri Lanka. Sri Lankan Journal of Agriculture and Ecosystems, 4(1). 2022.
2. Pansy, C., & Murali, B. An Accurate Mango Pest Identification Employing the Gaussian Mixture Model and Expectation-Maximization (EM) Algorithm. Journal of Computer Science, 21(1), 52–70. 2025.
3. Severtson, D., et al. (2020). Advantages and limitations of using conventional traps for insect pest monitoring in agriculture. Journal of Pest Science, 93, 693–706, 2020.
4. Ünal, E., & Türkoğlu, M. (2025). A comparative study of deep feature extraction models in mango leaf disease detection. European Journal of Science and Technology, 90309, 2025.

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)

An Evaluation Study of Transformer-Based Dehazing Models on Real World Datasets

Vishnuga Sivakumaran*, Amirthalingam Ramanan, Thanikasalam Kokul

Department of Computer Science, University of Jaffna, Sri Lanka

*vishnukumaran257@gmail.com

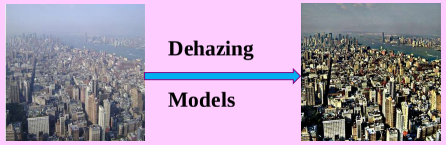
The presence of haze significantly reduces the visibility of outdoor scenes, which not only degrades the performance of many computer vision tasks but also compromises safety and reliable decision-making in practical systems. In recent years, transformer-based models have demonstrated remarkable performance in single-image dehazing by effectively capturing long-range dependencies and global context. However, their effectiveness under real-world haze conditions remains largely unexplored, as most models are trained and evaluated on synthetic or artificially generated datasets. To address this limitation, we evaluate two state-of-the-art transformer-based dehazing models, DehazeFormer and UMWTFormer, on both real-world datasets (LM-Haze, RW-Haze and RTTS) and synthetic datasets (Dense-Haze, RESIDE-OUT, O-HAZE and NH-HAZE), and analyse their performance using NIQE and BRISQUE scores. Experimental results reveal a substantial performance drop, with average dehazing effectiveness reduced by 40% on real-image datasets compared to synthetic ones. This performance gap highlights a critical limitation in current dehazing models, suggesting overfitting to synthetic data distributions and limited generalisation to real-world complexities. These findings underscore the need for new models or training strategies explicitly designed to cope with the variability and challenges inherent in real-world hazy images.

Keywords: Computer vision, Transformer-based architecture, Single-image dehazing

An Evaluation Study of Transformer-Based Dehazing Models on Real World Datasets

S. Vishnuga[1], A. Ramanan[2] and T. Kokul[2]
[1]Faculty of Graduate Studies, University of Jaffna, Sri Lanka
[2]Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

1. Introduction



- Haze is caused by atmospheric particles reducing image visibility, affecting various vision-based applications such as autonomous driving, remote sensing, aviation, agriculture, and surveillance.
- Recently, Transformer-based dehazing models achieve state-of-the-art results on synthetic datasets but struggle to generalise to complex real-world hazy scenes.
- This study evaluates Transformer-based dehazing models on three real-world datasets, revealing their limitations and suggesting future directions for robust single-image dehazing.

2. Related Works

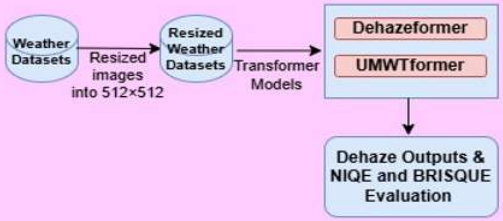
Study	Dataset	Results
Dehazeformer [1]	SOTS-indoor	37.84 PSNR, 0.99 SSIM
	SOTS-outdoor	34.95 PSNR, 0.98 SSIM
	RESIDE-6k	31.45 PSNR, 0.98 SSIM
	RS-Haze	39.87 PSNR, 0.97 SSIM
UMWTformer [2]	Outdoor-Rain	27.49 PSNR, 0.94 SSIM
	SOTS	31.41 PSNR, 0.98 SSIM
Mulsromer [3]	RESIDE-6k	31.45 PSNR, 0.98 SSIM
	SOTS	31.45 PSNR, 0.98 SSIM
	RTTS	3.09 NIQE

All models perform excellently on synthetic datasets, but their real-world generalisation remains limited, emphasizing the need for more diverse and realistic training data.

3. Datasets

Datasets	Number of images
LM-Haze	5040
RW-Haze	210
RTTS	4322
Dense-Haze	55
O-Haze	45
RESIDE-OUT	500
NH-Haze	55

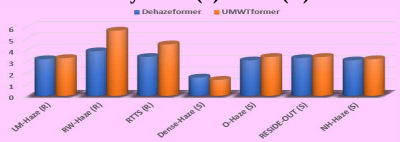
4. Methodology



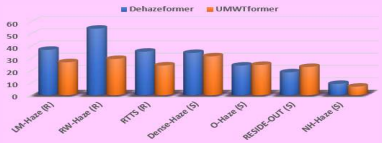
- This study evaluates two leading transformer-based dehazing models on multiple real and synthetic datasets without any retraining, ensuring a fair comparison and same conditions.
- All images were uniformly resized for consistent processing, and NIQE/BRISQUE were used to assess perceptual quality in the absence of ground truth.
- This evaluation highlights the true real-world effectiveness of transformer models beyond controlled synthetic benchmarks.

5. Results

NIQE Evaluation on Synthetic(S) & Real(R) Datasets



BRISQE Evaluation on Synthetic(S) & Real(R) Datasets



- The results reveal a clear performance gap: the models work very well on synthetic images but not as well on real haze images.
- The models perform strongly on synthetic datasets but drop in quality when used on real-world hazy images.

6. Conclusion

- This study highlights the need for more robust and generalised Transformer-based dehazing models to handle real-world haze.
- Future work should emphasize diverse, realistic training data and domain adaptation to close the synthetic-real performance gap.

7. References

1. Y. Song, Z. He, H. Qian, X. Du. "Vision transformers for single image dehazing", IEEE Transactions on Image Processing, IEEE, vol. 32, pp. 1927-1941, 2023.
2. A. Kulkarni, S. Phutke, S. Murala. "Unified transformer network for multi-weather image restoration", European Conference on Computer Vision, Springer, pp. 344-360, 2022.
3. Z. Wang. "Efficient multi-scale transformer with convolutional attention for high-definition image dehazing", IEEE Access, IEEE, 2025.

Attention-Enhanced Deep Learning for Diabetic Retinopathy Detection

Mahesan Natheesan*, Rajendran Nirthika

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*mnatheesan@gmail.com

The detection and classification of diabetic retinopathy (DR) using fundus images is essential for preventing vision loss in diabetic patients. While existing deep learning approaches have achieved reasonable performance, their accuracy can be significantly improved through advanced attention mechanisms that enable models to focus on critical retinal features. This study proposes an enhanced framework for DR detection by integrating the Convolutional Block Attention Module (CBAM) with EfficientNetB5 to achieve improved classification performance. The proposed approach utilizes EfficientNetB5 as the backbone network, enhanced with CBAM attention modules that combine channel and spatial attention mechanisms to improve feature representation. The CBAM module sequentially applies channel and spatial attention to emphasize important feature channels and identify critical spatial locations within fundus images. The framework processes fundus images through comprehensive preprocessing, including image enhancement, class balancing, and data augmentation, to ensure robust model training. We combined the APTOS and IDRiD datasets, yielding a total of 4,178 fundus images. The integrated dataset was preprocessed to remove irrelevant regions, noise, and artifacts. Our CBAM-enhanced EfficientNetB5 model achieved a 1.14% accuracy improvement over the existing state-of-the-art performance (96.06%). While the integration of CBAM improved accuracy, it also introduced a slight increase in model complexity. The attention-enhanced architecture enables the model to automatically focus on clinically significant features such as microaneurysms, hemorrhages, and exudates, which play a vital role in precise DR classification, highlighting the effectiveness of attention mechanisms in medical image classification for DR detection.

Keywords: Diabetic Retinopathy Detection, Deep Learning, Image Enhancement, Attention Mechanism, Data Fusion

Attention-Enhanced Deep Learning for Diabetic Retinopathy Detection

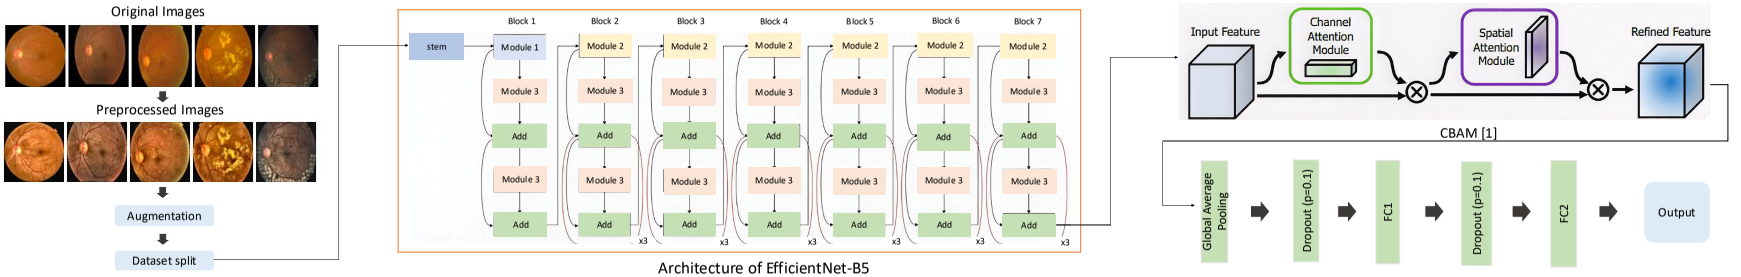
Natheesan, M., and Nirthika, R.
Department of Computer Science, University of Jaffna, Jaffna, Sri Lanka
mnatheesan@gmail.com, nirthika@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

Introduction and Motivation

Diabetic Retinopathy (DR) is one of the leading causes of blindness worldwide. Early detection is essential, yet it is often hindered by slow and subjective clinical diagnosis. Traditional automated methods depend heavily on large labeled datasets, which are costly and time-consuming to obtain. Although existing deep learning approaches have shown promising results, their accuracy can be further improved by incorporating advanced attention mechanisms that help models focus on the most critical retinal features. This study proposes an enhanced DR detection framework by integrating the Convolutional Block Attention Module (CBAM) [1] with EfficientNetB5 to achieve more accurate and robust classification performance. Our CBAM-enhanced EfficientNetB5 model achieved a 1.14% accuracy improvement over the existing state-of-the-art performance.

Methodology



Architecture of EfficientNet-B5

Loss Function

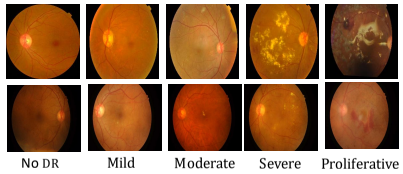
Categorical Cross-Entropy Loss

$$\text{Loss} = -\sum_{i=1}^{\text{output size}} y_i \cdot \log \hat{y}_i$$

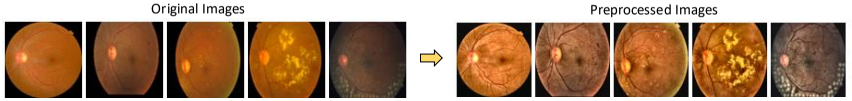
y_i: True label; $\hat{y}_i$: Predicted probability

Datasets

Combined APTOS and IDRiD for a total of 4,178 images in five classes.



Preprocessing



A four-step preprocessing pipeline was applied.

1. Fundus images were cropped to eliminate irrelevant background regions.
2. Then resized to a consistent dimension of 180 × 180 × 3.
3. CLAHE [2] was applied to improve image contrast and enhance the visibility of retinal features.
4. SMOTE [3] was employed to address class imbalance issues.

	No DR	Mild	Moderate	Severe	Proliferative	Total
Original Images	1973	395	1167	286	357	4178
Balanced Images	1973	1880	1650	1750	1650	**8903**

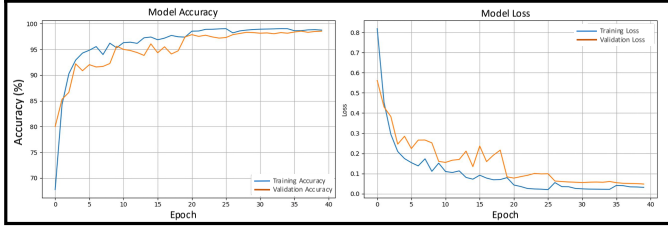
Number of image distribution per class after SMOTE

Data augmentations - Vertical & Horizontal flips, Brightness (10%), Contrast (10%)

Experimental Setup

- **Backbone:** EfficientNet-B5 (pre-trained on ImageNet)
- **Attention Module:** CBAM (Channel & Spatial Attention)
- **Data Split:** Training - 80%, Validation - 10%, Testing - 10%
- **Learning rate:** 0.001
- **Batch size:** 32
- **Epochs:** 40
- **Optimizer:** Adam
- **Loss Function:** Categorical Cross-Entropy

Results and Discussion



Model	Accuracy (%)
EfficientNet-B5 [2]	96.06
EfficientNet-B5 + CBAM (between 6th-7th blocks)	**96.09 ± 0.06**
EfficientNet-B5 + CBAM (after 7th block)	**97.20 ± 0.05**

Mean ± standard deviation over 3 experiments

Conclusion

This study introduced an attention-enhanced EfficientNet-B5 model for automated diabetic retinopathy classification using a combined APTOS and IDRiD dataset. The model effectively improved classification performance by leveraging the CBAM module to focus on key retinal regions. Experimental results showed that integrating CBAM enhanced feature representation and significantly boosted overall classification accuracy.

References

[1] Woo, S., Park, J., Lee, J. Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 3-19.
[2] Akhtar, S., & Aftab, S. (2024). A framework for diabetic retinopathy detection using transfer learning and data fusion, International Journal of Information Technology and Computer Science, vol. 16, no. 6, pp. 61-73.
[3] Aftab, S., & Akhtar, S. (2025). Diabetic retinopathy severity classification using data fusion and ensemble transfer learning, Journal of Software Engineering and Applications, vol. 18, no. 1, pp. 1–23.
[4] https://www.kaggle.com/datasets/mariaherrerot/aptos2019
[5] https://ieee-dataport.org/open-access/indian-diabetic-retinopathy-image-dataset-idrid

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P06
21

Automated Fungi Classification with VGG16-based Feature Fusion and Contrast Enhancement

Vimalenthiran Thusani*, Rajendran Nirthika

Department of Computer Science, University of Jaffna, Sri Lanka

*vimalanthushani1122@gmail.com

This study introduces a deep learning framework for microscopic fungi image classification, addressing the critical need for timely and accurate identification of fungal infections in medical diagnostics. Traditional microscopic examination often struggles to differentiate fungal species due to high visual similarity, requiring biochemical and culture-based tests that delay treatment and increase patient risks. To overcome these limitations, this research leverages the DeFungi dataset, a curated repository of microscopic images of yeasts, moulds, and dermatophytes. The proposed approach integrates Contrast Limited Adaptive Histogram Equalization (CLAHE) to enhance local contrast and highlight subtle morphological variations that are difficult to detect under standard imaging. A pretrained VGG16 was used for hierarchical feature extraction, with feature fusion applied after block 4 to enrich representations and improve discriminative power. The baseline model achieved 85.04validation accuracy, while the CLAHE-enhanced VGG16 with feature fusion reached a mean class accuracy of 95.01%, surpassing baseline models and demonstrating the importance of both contrast enhancement and multilevel feature integration. This advancement supports automated mycological diagnostics, offering progress toward precision medicine, faster clinical decision-making, and improved outcomes. Future work will focus on lightweight models for real-time deployment and extending the framework to larger, more diverse datasets to further validate robustness and practical utility in medical mycology.

Keywords: CLAHE, Feature Fusion, Fungal Classification, Image Enhancement, VGG16

Automated Fungi Classification with VGG-16-Based Feature Fusion and Contrast Enhancement

Thusani, V., and Nirthika, R.
Department of Computer Science, University of Jaffna, Jaffna, Sri Lanka
vimalanthushani1122@gmail.com, nirthika@univ.jfn.ac.lk

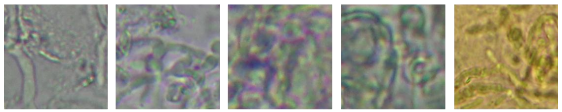
SLSWCS Student Workshop on Computer Science

Introduction

Fungi are essential to ecosystems, agriculture, and biotechnology, but some species can cause serious infections in humans, animals, and plants. Accurate identification of fungi is crucial for medical diagnosis, disease management, and ecological studies. Traditional microscopic identification is slow, subjective, and highly dependent on expert knowledge [1]. Deep learning has advanced biological image analysis, yet fungal classification remains challenging due to high inter-class similarity, class imbalance, and noisy microscopic images. Previous works did not explore multistream enhancement or feature fusion, limiting complementary feature capture. We propose an efficient image enhancement-based feature fusion framework using VGG16 to improve visual quality and discriminative features in fungi image classification.

Dataset

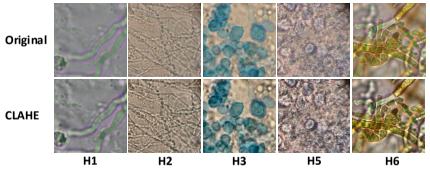
- The publicly available DeFungi dataset[1] contains 9,114 microscopic fungal images, each of size 500 × 500 pixels.
- It includes five fungal classes H1, H2, H3, H5, and H6 corresponding to TSH (4,404 images), BASH (2,334), GMA (819), SHC (818), and BBH (739).
- Sample images from each class demonstrate the visual diversity and subtle morphological differences within the dataset.



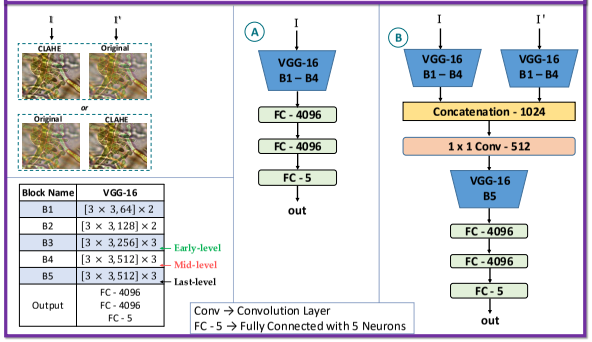
[1]https://www.kaggle.com/camilovu/lemm-pre-processed

CLAHE Enhancement

- Contrast Limited Adaptive Histogram Equalisation (CLAHE) is a contrast enhancement method that improves visibility in images with uneven illumination by applying histogram equalisation in small local regions.
- It reveals fine fungal details such as subtle textures, hyphal structures, and spore boundaries while preventing noise amplification.
- Sample CLAHE-enhanced images are shown below:



VGG16 Feature Fusion Architectures



Block Name	VGG-16
B1	[3 × 3, 64] × 2
B2	[3 × 3, 128] × 2
B3	[3 × 3, 256] × 3
B4	[3 × 3, 512] × 3
B5	[3 × 3, 512] × 3
Output	FC - 4096 FC - 4096 FC - 5

Experimental Setup

- Data augmentation:
  - Random rotations (± 20°)
  - Horizontal flips
  - Random affine transformations with shear and scale
- Epochs: 100, Batch size: 32
- Dataset: 85% - training, 15% - validation

Weighted Cross-Entropy Loss:

$$w_i = \frac{N}{C \times n_i}$$

$$L = -\sum_{i=1}^{c} w_i y_i \log(\hat{y}_i)$$

w_i - Weight for each class
N - Number of samples
C - Number of classes
n_i - Samples in class i
y_i - One hot encoded ground truth
$\hat{y}_i$ - Predicted probability

Experiments and Results

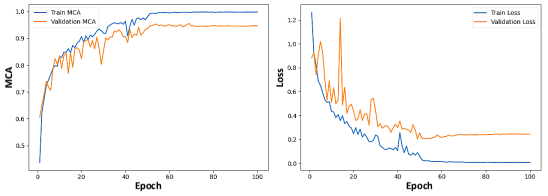
Optimizer	Method	Architecture	MCA ± SD (%)
Adam	Original	A	92.12 ± 0.38
SGD	Original	A	94.62 ± 0.39
	CLAHE	A	94.29 ± 0.08
	After B3	B	91.93 ± 0.39
	After B4	B	95.01 ± 0.37
	After B5	B	92.45 ± 0.49

Validation results

Adam → learning rate (lr) = 1e-5
SGD → lr = 1e-3
lr scheduler reduces by 0.1,
Milestone = [50, 70, 90]

Model	Accuracy (%)	MCA (%)
InceptionV3 [1]	82.80 ± 3.50	-
ResNet-50 [1]	83.04 ± 1.86	-
VGG-16 [1]	85.04 ± 1.97	-
ViT + VGG-16 [2]	85.43	-
ViT + InceptionV3 [2]	86.53	-
ViT + ResNet-50 [2]	90.12	-
MobileNetV3 [3]	92.89	-
ResNet-101 [4]	93.00	-
GNN [5]	93.10	-
After B4	94.82 ± 0.45	95.01 ± 0.37

Comparison with the state-of-the-art approaches on the DeFungi dataset



Training and validation MCA and loss curves showing the best-performing model

The proposed mid-level fusion strategy yielded the highest performance among all configurations.

Conclusion

- This work presents a contrast-enhanced feature fusion framework that enhances microscopic fungal classification.
- Mid-level fusion of original and CLAHE-enhanced features in VGG-16 captures richer textures.
- The method achieves 95.01% MCA, surpassing all baselines and previous state-of-the-art results on the DeFungi dataset.
- Overall, enhancement-driven multistream fusion provides a simple, reliable, and scalable approach for automated fungal identification.

References

[1] C. J. P. Sopo, F. Hajati, and S. Gheisari, "DeFungi: Direct my cological examination of microscopic fungi images," arXiv preprint arXiv:2109.07322, 2021.
[2] S. I. Ahmed and A. O. Haque, "Microscopic fungi classification using vision transformer guided by transfer learning approach," in 26th International Conference on Computer and Information Technology, pp. 1–6, IEEE, 2023.
[3] J. Bhimavarapu, A. Chinta, S. V. Movva, and J. P. Jampani, "Fungi classification: Enhancing diagnosis using deep learning," in 2nd World Conference on Communication & Computing, pp. 1–6, IEEE, 2024.
[4] M. S. Rashid, M. S. Morshed, M. U. Islam, S. Rashid, A. Mahmud, and A. Islam, "Mycological Examination of Microscopic Fungi Images with Deep Learning and Gradient Weighted Class Activation Mapping Visualization," in Advances in Science and Engineering Technology International Conferences, pp. 01–08, IEEE, 2024.
[5] P. Sajjadi, *et. al.*, "Classification of direct microscopic fungi images using optimized graph networks," Biomedical Signal Processing and Control, vol. 109, p. 108035, 2025.

Automated Foundation Structure Identification from Floorplans using Deep Learning

Ravindu Deshitha*, Somaskandan Suthakar, Siyamalan Manivannan

Department of Computer Science,University of Jaffna, Sri Lanka

*2020csc067@univ.jfn.ac.lk

Floorplan understanding is critical for architectural design, reconstruction, and urban planning. Accurate detection of foundation structures help to determine the stability and the functionality of buildings. We present a deep learning framework for automated foundationstructure identification based on a multi-head U-Net with a ResNet encoder, that jointly segmentations foundation-lines and junction heatmaps. The proposed approach was evaluated on a real-image dataset of 1,000 floorplans annotated with foundation structure lines and junction points. The framework achieved a precision of 0.9438 and a recall of 0.9484 for junction detection, as well as an intersection-over-union (IoU) of 0.8806 for foundation-line segmentation. These results outperform other segmentation models, including DeepLab-v3+ (precision 0.9273, recall 0.9329, IoU 0.8569) and the vanilla U-Net (precision 0.3731, recall 0.6804, IoU 0.5784), demonstrating robust learning across diverse layout configurations. Qualitative analysis shows accurate localization of key junctions and coherent foundation structure predictions across diverse floorplans. This work contributes a practical pipeline for real-image dataset preparation, a multi-task architecture tailored to foundation cues, and empirical evidence supporting its effectiveness for structural recognition. Future work will involve curating and annotating more diverse real-world floorplans to fine-tune the model and extending the system to explicitly map connections between junctions and foundation structures for enhanced structural reasoning and downstream engineering applications.

Keywords: foundation detection, floor plan segmentation, junction analysis, floor plan analysis, heatmap segmentation

Automated Foundation Structure Identification from Floorplans using Deep Learning

Ravindu Deshitha, Somaskandan Suthakar, Manivannan Siyamalan

Department of Computer Science, University of Jaffna, Sri Lanka

Introduction and Background

Floorplans contain walls and rooms, but foundation structures and junctions are often overlooked, limiting accurate reconstruction and structural analysis.

We present a multi-head U-Net with ResNet encoder that detects foundation lines and junctions using heatmaps, enabling precise and coherent identification across diverse real-image floorplans.

Dataset Details

- Real Floorplan Images: We collected 1,000 real floorplan images (512×512 pixels), manually annotated with foundation lines and junction points.
- Training Set: 700 images were augmented to a total of 2,100 for robust training.



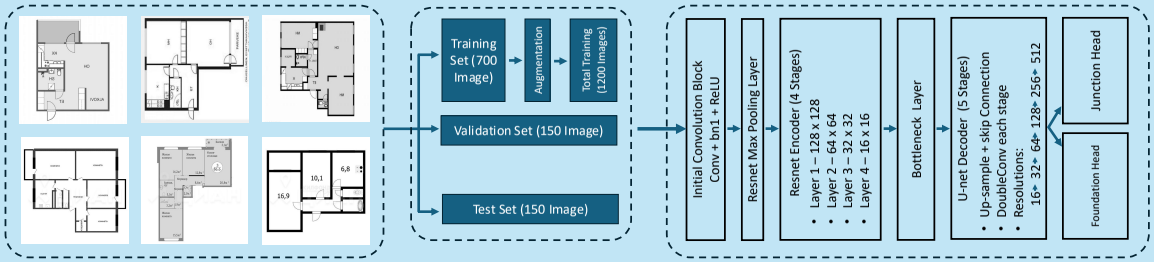
Original Floorplan Image



Annotations drawn on the floorplan

Methodology and Diagram

- Heatmap Generation: Foundation lines and junctions are represented as heatmaps to guide learning.
- ResNet-U-Net: Multi-head U-Net with a ResNet encoder predicts both foundation structure and junction heatmaps simultaneously for accurate structure detection.



Results and Discussion

- Best Model: The ResNet-encoder U-Net achieved the highest performance with strong junction F1 (0.9461) and the best foundation-line IoU (0.8806).
- Quality of Predictions: The model produced clean, continuous foundation lines and accurately localized junctions, outperforming all baseline architectures.

Model	Junction Point			Foundation structure		Final Score (Mean IoU +F1)/2
	Precision	Recall	F1-Score	Mean IoU	Median IoU	
DeepLabV3+	0.9273	0.9329	0.9301	0.8569	0.8719	0.8935
Vanilla U-Net	0.3731	0.6804	0.4819	0.5784	0.5920	0.5302
U-Net with Batch Normalize	0.9166	0.9683	0.9418	0.8679	0.8822	0.9048
ResNet-encoder U-Net	0.9438	0.9484	0.9461	0.8806	0.8935	0.9133

Table: Evaluation results comparing all trained model variations

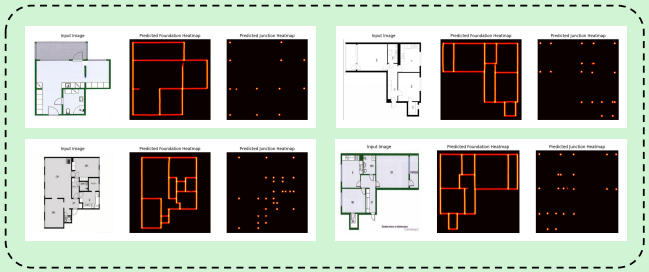


Image: Sample outputs - Predicted foundation structure and junction heatmaps (test set)

Conclusion

- The ResNet-U-Net model achieved the strongest performance for foundation-line and junction detection, demonstrating reliable structural extraction from floorplans.
- Future work will expand the dataset and incorporate connectivity reasoning to reconstruct full foundation geometry for real-world architectural use.

Comparative Analysis in Textual and Audio Modality on Cyberbullying Context

Arthika Anandapalan*, Yuraj Perera , Jananie Jarachanthan

Faculty of Engineering, University of Jaffna, Sri Lanka

*arthikaanandapalan@gmail.com

Cyberbullying, defined as repeated aggressive or harmful behavior conducted through digital platforms, has emerged as a critical issue. The rise of digital platforms has created new avenues for such behavior, leading to psychological distress, disrupted learning, and strained interpersonal relationships. While many studies focus on textual data or image cues, the role of audio, particularly emotional signals in speech, remains comparatively underexplored. This study presents a comparative analysis of textual and audio modalities for cyberbullying detection. For text, transformer-based models such as BERT, Mistral-7B, Gemma-7B, and LLaMA-2 are fine-tuned on publicly available cyberbullying datasets to classify offensive, racist, sexual, and neutral content. For audio, emotional speech corpora including CREMA-D, RAVDESS, SAVEE, and TESS are employed to benchmark classical machine learning models (Logistic Regression, Random Forest, Support Vector Machines, Multilayer Perceptron, and Naive Bayes) against CNN-based deep learning models with channel, spatial, and multidimensional attention to highlight important audio features. Results show that transformer-based models achieve F1-scores above 0.90 in text detection, while attention-enhanced CNNs improve recall in aggression detection but remain below 0.50 of F1-score overall. These findings highlight the complementary strengths and limitations of each modality and establish a foundation for future multimodal fusion research in cyberbullying detection.

Keywords: Cyberbullying Detection, Textual Modality, Audio Modality, Deep Learning

Comparative analysis in Textual and Audio Modality on Cyberbullying context

Arthika A.[1], Perera A.Y.I.D.[1], Jananie J.[1]
Department of Computer Engineering Faculty of Engineering University of Jaffna[1]

1. INTRODUCTION

Cyberbullying has become a pervasive problem across online platforms, causing emotional harm, anxiety, and long-term psychological effects among users. Although extensive research has been conducted using text-based datasets, audio cues—especially emotional signals in speech—remain underexplored in the context of cyberbullying detection. Most existing studies rely heavily on text or images, which overlook important vocal characteristics such as tone, anger, stress, and sarcasm.

This research addresses that gap by examining cyberbullying detection through a dual-modality perspective. It compares the performance of transformer-based text classification models and attention-enhanced audio models to determine how each modality contributes to detecting abusive, offensive, or aggressive behavior online. The comparison provides insight into their complementary strengths and serves as a foundation for future multimodal cyberbullying detection systems.

2. METHODOLOGY

The study employs two separate pipelines for text and audio.

2.1. TEXT MODALITY

The textual component of this study focuses on classifying cyberbullying-related content using transformer-based language models. Kaggle cyberbullying dataset were cleaned, tokenized, encoded, and stratified into an 80% training, 10% validation, and 10% testing split to preserve class balance.

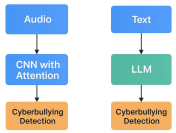
Four large language models—*BERT-base, Mistral-7B, Gemma-7B*, and *LLaMA-7B* were fine-tuned to detect cyberbullying text categories. During training, attention was placed on optimizing contextual understanding, reducing false positives, and improving generalization across ambiguous or borderline phrases. Each model learned semantic patterns associated with harassment through supervised learning, and performance was evaluated using accuracy, precision, recall, and F1-score. Confusion matrix analysis was used to assess misclassification behavior and identify category-specific weaknesses. This modality establishes a clear benchmark for how effectively modern transformer architectures detect textual cyberbullying.

2.1. AUDIO MODALITY

The audio component investigates aggression and hostility expressed through speech by analyzing emotional and acoustic cues. Publically available emotional speech datasets *(CREMA-D, RAVDESS, SAVEE TESS)* were preprocessed by converting the raw audio into Mel-spectrograms and extracting acoustic features including MFCCs, zero-crossing rate, spectral energy, chroma, RMSE, and tonal characteristics.

CNN and CNNs with channel, special and multidimensional (spatial + channel) were trained by partitioning the datasets into 64% training, 16% validation, and 20% testing splits focusing on identifying meaningful audio cues such as MFCCs, chroma, and mel-spectrogram features while reducing the influence of irrelevant patterns. Channel Attention recalibrated feature-channel importance, Spatial Attention highlighted key time-frequency regions, and CBAM combined both mechanisms for multidimensional weighting.

Performance was assessed using accuracy, cross-validation mean, and class-specific metrics, with the multidimensional attention model achieving the best overall results due to its superior ability to emphasize aggression-related acoustic features. This modality establishes a benchmark for how effectively attention-enhanced CNN architectures detect cyberbullying aggression from audio signals.



Architectur
e

3. RESULTS

Below is the performance comparison of the LLMs where the text was trained.

Text Modality

Model	Accuracy	Precision	Recall	F1-Score
BERT-base-uncased	0.8780	0.8621	0.9000	0.8806
MistralAI-7B-V.01	0.8830	0.8839	0.8830	0.8830
Gemma - 7B	0.9020	0.9026	0.9020	0.9020
Llama 2 - 7B	0.8790	0.8825	0.8790	0.8786

Audio Modality

Model	Accuracy	Precision	Recall	F1-Score
CNN	0.820	0.735	0.788	0.745
ACNN - Channel	0.800	0.676	0.741	0.694
ACNN – Spatial	0.836	0.781	0.794	0.775
ACNN – Multi- Dimensional	0.872	0.829	0.842	0.819

4. DISCUSSIONS

Text models detect explicit abusive language reliably due to clear and abundant written data, while audio captures emotional cues like tone and intensity that reveal hidden aggression. However, audio performance is affected by limited real cyberbullying datasets, noise, accents, and imperfect ASR transcripts. Text datasets also suffer from class imbalance. Overall, each modality offers distinct strengths but also practical limitations, highlighting the need for multimodal systems for more complete cyberbullying detection.

5. CONCLUSION

This study shows that text models excel at detecting explicit abusive content, while audio models capture emotional aggression that text alone may miss. Both modalities contribute unique strengths, and combining them offers a more complete and reliable approach to cyberbullying detection.

6. REFERENCES

[1] W. A. K. M. Wickramaarachchi, S. S. Subasinghe, K. K. R. T. Wijerathna, A. S. U. Athukorala, L. Abeywardhana, and A. Karunasena, "Identifying False Content and Hate Speech in Sinhala YouTube Videos by Analyzing the Audio," 2023 5th International Conference on Advancements in Computing (ICAC). IEEE, Dec. 07, 2023. doi: 0.1109/icac60630.2023.10417565.

[2] R. A. Silakhal, S. Zekri, O. Salem, and A. Mehaoua, "Cyberbullying Detection Through Acoustic and Linguistic Analysis," 2023 International Symposium on Networks, Computers and Communications (ISNCC). IEEE, pp. 1–5, Oct. 23, 2023. doi: 10.1109/isncc58260.2023.10323647.

Deep Learning-Based Classification of Cerebellum Failure Using MRI Imaging and Ensemble Modeling

Mahendrarajah Vasavan, Luxshi Karunakaran*, Chathurani Nadika Hettiarachchi

Department of Physical Sciences and Technology, Faculty of Applied Science,
Sabaragamuwa University of Sri Lanka

*klluxshi99@gmail.com

Cerebellum failure significantly impairs motor coordination and cognitive function, making its early detection essential for improving patient outcomes. This study focuses on the timely and precise diagnosis of cerebellum failure using deep learning-based medical image classification techniques. The research utilizes a Kaggle-sourced 2D MRI brain dataset comprising four classes: ataxic, degenerative, vascular, and no failure, collected from publicly available medical imaging repositories. The dataset was preprocessed through augmentation and noise reduction to improve quality and diversity, then split into 80% training and 20% testing before applying the deep learning models. Three architectures were explored: a custom Convolutional Neural Network (CNN), MobileNetV2 (a pre-trained CNN variant optimized for lightweight deployment), and a Multi-Layer Perceptron (MLP), combined through a stacking ensemble where the MLP served as the meta-learner. The model performance was evaluated using 5-fold cross-validation to ensure robustness and generalizability. The ensemble achieved the highest accuracy of 0.97, outperforming individual modelsâĂŤcustom CNN (0.93) and MobileNetV2 (0.94) as well as baseline models such as VGG16 (0.90) and ResNet50 (0.92). The ensemble also demonstrated strong reliability, with precision of 0.96, recall of 0.95, and macro-F1 of 0.97, indicating balanced classification across all categories. The final system was implemented as a Streamlit-based web application for real-time clinical applicability, supporting quick and accurate cerebellum failure diagnosis. Despite standard model architectures being used, the integration of stacking ensemble learning with clinical deployment represents a practical advancement for automated MRI-based neurological assessment.

Keywords: Medical image classification, cerebellum failure detection, MRI imaging, ensemble deep learning, CNN, MobileNetV2

DEEP LEARNING-BASED CLASSIFICATION OF CEREBELLUM FAILURE USING MRI IMAGING AND ENSEMBLE MODELING

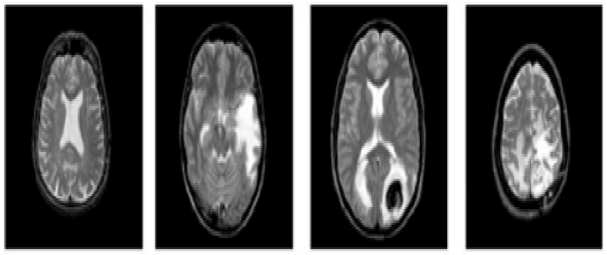
Vasavan Mahendrarajah[1], Luxshi Karunakaran[1*] and CN Hettiarachchi [1]

[1]*Faculty of Applied Sciences, Sabaragamuwa University of Sri Lanka*

*mahenvas@gmail.com

Introduction

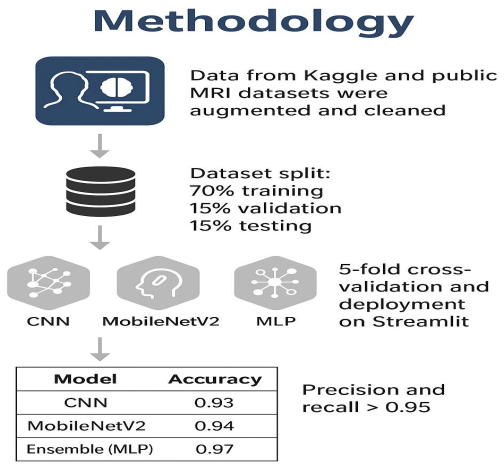
Cerebellum failure is a serious neurological disorder affecting motor coordination, balance, and cognitive functions, greatly reducing quality of life. Delayed diagnosis may cause irreversible damage and increased mortality. Manual MRI analysis is time-consuming and prone to error, making automated diagnosis essential. Recent advances in deep learning have shown strong potential for medical image classification, but single models often suffer from overfitting and poor generalization. Therefore, this study proposes an **ensemble-based deep learning framework** to improve reliability and diagnostic accuracy.



a) Normal b) Glioma c) Meningioma d) Pituitory

Methodology

Deep Learning (DL) provides a promising solution by enabling automated and accurate detection of cerebellum failure from MRI images using data-driven learning models.



Step	Details
Dataset Source	MRI images collected from Kaggle and publicly available medical imaging datasets
Data Collected	Cerebellum MRI scans (Ataxic, Degenerative, Vascular, Normal)
Preprocessing	Noise removal, normalization, resizing, data augmentation
Models Tested	CNN, MobileNetV2, Multi-Layer Perceptron (MLP)
Training setup	Adam optimizer, categorical cross-entropy loss
Evaluation	Accuracy, Precision, Recall, F1-score

Results

Algorithms	Accuracy	Precision	F1 score
CNN	0.93	0.95	0.94
Mobile net v2	0.94	0.94	0.94
Ensemble model	0.97	0.97	0.97

Table 1: Algorithm's accuracy, precision and F1 score

Conclusion

This study classified cerebellum failure (ataxic, degenerative, vascular, no failure) using CNN, MobileNetV2, and an ensemble model, with the latter achieving 0.97 accuracy and F1-score. The ensemble outperformed individual models and baselines (VGG16, ResNet50), enhancing diagnostic confidence. Future work will explore larger datasets and segmentation techniques to further improve detection.

Detection and Classification Betel Leaf Diseases using Image Processing Techniques

H. P. Isuru Lakmal*, Barathy Mayurathan

Department of Computer Science, University of Jaffna, Sri Lanka

*99isurulakmal@gmail.com

Betel leaf (Piper betle) is a widely cultivated Sri Lankan crop that is highly susceptible to fungal and bacterial infections, yet manual diagnosis remains slow, subjective, and dependent on expert availability. This research aims to develop an automated and reliable detection pipeline by integrating image processing, traditional machine learning, and deep learning models. A dataset of 2,589 betel leaf images was collected from the Mendeley repository and split into 70% training and 30% testing. Preprocessing included resizing, normalization, noise reduction, and data augmentation, while background removal was applied using U-Net, rembg, and the Snake algorithm to address the difficulty of isolating diseased regions from complex backgrounds. Texture-based features such as Local Binary Patterns and Gray Level Co-occurrence Matrices were used with Support Vector Machines and Random Forest classifiers, whereas a custom Convolutional Neural Network with three convolutional blocks was implemented to learn spatial features automatically. Transfer learning using DenseNet121 demonstrated the highest classification accuracy of 99.44%. Results indicate that background removal improves performance by eliminating noise and highlighting disease-relevant regions, with the Snake algorithm offering the best trade-off between boundary preservation and generalization. Compared with state-of-the-art plant disease models, the proposed method performs competitively while addressing the underexplored domain of betel leaf disease classification. The findings establish that integrating preprocessing, segmentation, and deep learning significantly enhances accuracy, and future work will focus on hybrid models and larger datasets to improve robustness.

Keywords: betel leaf, disease classification, image processing, CNN, transfer learning

DETECTION AND CLASSIFICATION BETEL LEAF DISEASES USING IMAGE PROCESSING TECHNIQUES

H.P. Isuru Lakmal, Barathy Mayurathan
Department of Computer Science, Faculty of Science, University of Jaffna
2020csc015, barathym@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

Abstract

Betel leaf (Piper betle) is a widely cultivated Sri Lankan crop that is highly susceptible to fungal and bacterial infections, yet manual diagnosis remains slow, subjective, and dependent on expert availability. This research aims to develop an automated and reliable detection pipeline by integrating image processing, traditional machine learning, and deep learning models. A dataset of 2,593 betel leaf images was collected from the Mendeley repository and split into 70% training and 30% testing. Preprocessing included resizing, normalization, noise reduction, and data augmentation to improve model generalization.

Texture-based features such as Local Binary Patterns and Gray Level Co-occurrence Matrices were extracted and used with Support Vector Machines and Random Forest classifiers to evaluate the effectiveness of handcrafted features. Additionally, a custom Convolutional Neural Network with three convolutional blocks was developed to automatically learn spatial features, while transfer learning using DenseNet121 achieved the highest classification accuracy of 99.44%.

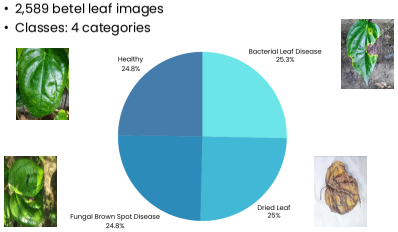
Results show that deep learning models significantly outperform traditional machine learning approaches, highlighting the importance of automated feature extraction for betel leaf disease classification. Compared with state-of-the-art plant disease models, the proposed method performs competitively while addressing the underexplored domain of betel leaf disease detection. Future work will focus on hybrid feature learning methods and larger datasets to further enhance robustness and real-world applicability.

Introduction

- Betel leaf (Piper betle) is an important cash crop in Sri Lanka, but its quality and market value are often reduced.
- Farmers traditionally identify these conditions through manual visual inspection, which is slow, subjective, and varies based on the experience of the observer.
- These results are then compared with deep learning techniques, especially Convolutional Neural Networks (CNNs), which automatically learn and extract features directly from images.
- By providing a fast and accurate detection method, this system helps farmers reduce crop losses, improve grading accuracy, and increase overall income through better disease management and higher-quality produce.

Dataset

- Mendeley Data Repository (Bangladesh)
- 2,589 betel leaf images
- Classes: 4 categories



- Split into 70% for Training and 30% Testing
- Image Resolution: 600 × 800 pixels
- RGB leaf images captured under natural lighting

Methodology

Data Collection and Preprocessing

- Image Resizing : Uniform size 224x224 pixels
- Data Split : 70% training , 30% testing
- Image Augmentation : To prevent overfitting and increase dataset size, applied Random flip (horizontal & vertical),Random rotation up to 15%, Random zoom up to 12%, Random contrast adjustments
- Normalization : Scaled pixel values to the range [0 – 1]

Model Selection

1. Deep Learning - DenseNet121

Used CNN models such as VGG16, ResNet50 and DenseNet121
DenseNet121 model show better performance with our Dataset.



Training Configuration

- Optimizer : Adam
- Learning rate : 0.0001
- Loss function : Sparse categorical entropy
- Fine-tune : Early-Stopping

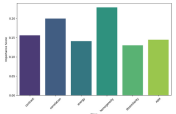
Experimental Setup

- Hyperparameters
  - Batch Size : 32
  - Epochs : 20
  - Seed : 123
  - Autotune : None

2.Machine Learning - GLCM + SVM

Feature Extraction - GLCM + LBP

- contrast
- correlation
- energy
- homogeneity
- dissimilarity
- ASM



Feature Selection - RF

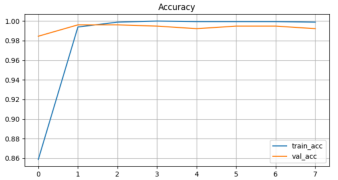
Feature Classification - RF + KNN + SVM

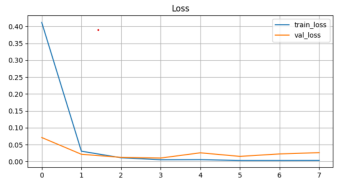
Results

- Comparison of Previous models

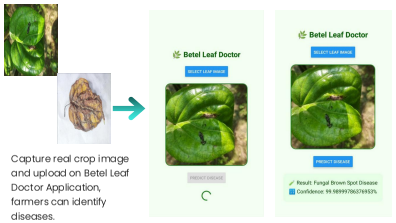
Research Paper	Model	Accura
[2]	DenseNet121 - Pumpkin	97.82%
Our Method	Machine Learning - SVM	78.21%
	DenseNet121	99.44%

- Loss & Accuracy versus epochs for deep learning :





- Betel Leaf Doctor : Mobile Application for Farmers



Capture real crop image and upload on Betel Leaf Doctor Application, farmers can identify diseases.

Discussion & Conclusion

- The comparison table highlights the performance differences between a previous study and our experiments. The reference study [2] achieved 97.82% accuracy using DenseNet121 on pumpkin leaf disease images. Since pumpkin leaves have different texture patterns compared to betel leaves, this result cannot be directly applied to our domain.
- In our work, the Machine Learning (SVM) model using texture features reached 78.21% accuracy, showing that handcrafted features provide moderate performance but are limited in capturing complex disease patterns. However, when we fine-tuned DenseNet121 on our betel leaf dataset, the accuracy increased significantly to 99.44%, demonstrating that deep learning is more effective for identifying subtle variations in betel leaf diseases.

The results demonstrate that data augmentation significantly improves model performance by increasing the variability of training samples and reducing overfitting. Background removal further enhances accuracy by allowing the model to focus only on the leaf region, eliminating unwanted noise. The fine-tuning of a transfer learning model such as DenseNet121 proved highly effective, as deep learning architectures can automatically extract complex spatial and texture patterns that traditional methods often miss. While texture-based machine learning methods provide useful insights, their accuracy is limited compared to deep learning. Overall, the experiment confirms that automatic feature extraction through deep learning, supported by augmentation and background preprocessing, is the most reliable approach for betel leaf disease classification.

References

[1] J. Biswas, A. R. U. Islam Rafim, T. Paul, M. M. Hossain, and others, "Automatic Detection of Pumpkin Leaf Diseases Using Transfer Learning and a Custom Dataset from Bangladesh," in *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)*, September 2024. doi:10.1109/COMPAS60761.2024.10796774.

[2] A. Author(s), "Betel Leaf Disease Classification Using Data Augmentation and Convolutional Neural Network," *[Journal/Conference]*, 2025.

[3] R. Polly and E. Anna Devi, "Semantic Segmentation for Plant Leaf Disease Classification and Damage Detection: A Deep Learning Approach," *[Journal/Conference]*, Sathyabama Institute of Science and Technology, Chennai, India.

[4] F. David and M. A. Mukunthan, "Betel Leaf Diseases Classification using Machine Learning Algorithm: A Feasible Approach," Dept. of Computer Science and Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Inst. of Science and Technology, Chennai, Tamil Nadu, India, 2022.

[5] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," Computational intelligence and Neuroscience, vol. 2016, pp. 1–11, 2016

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P11

Enhanced Prediction of Non-Communicable Diseases using Ensemble Machine Learning Approaches

Lenuga Mahendran[1,2], Banujan Kuhaneshwaran[1,3], Luxshi Karunakaran[4], Vadivel Abishethvarman[1,*], Senthan Prasanth[4,5], Banage. T. G. S. Kumara[1,6]

[1]Department of Computing and Information Systems, Sabaragamuwa University of Sri Lanka,
[2]School of Computer Science, Northumbria University, United Kingdom,
[3]Faculty of Science and Engineering, Southern Cross University, Australia,
[4]Dept. of Physical Sciences and Technology, Faculty of Applied Sciences, Sabaragamuwa University of Sri Lanka,
[5]Faculty of Engineering and Applied Science, Memorial University of Newfoundland, Canada,
[6]Department of Software Engineering, Sabaragamuwa University of Sri Lanka.

[*]abishethvarman@gmail.com

Non-communicable diseases (NCDs) are a major global health issue, and early detection is key to prevention and management. Conventional predictive models cannot always reflect the complex, non-linear associations between anthropometric indicators, limiting their clinical application. This study aims to develop a predictive model for identifying NCDs using anthropometric measurements and to determine the most accurate supervised learning technique and the most influential anthropometric parameters. The dataset comprised 300 adult subjects (164 males and 136 females) from the Jaffna Teaching Hospital and Sabaragamuwa University, of Sri Lanka. Data were collected using body composition analyzers and divided into training (80%) and testing (20%) sets. The key characteristics included age, body mass index (BMI), waist-hip ratio (WHR), visceral fat area (VFA), body fat percentage, total body water, and muscle body fat. Ensemble machine learning (ML) methods, such as Random Forest, AdaBoost, CatBoost, and XGBoost, were compared with traditional classifiers such as Support Vector Machine, Artificial Neural Network, Logistic Regression, and Decision Tree. The ensemble models consistently outperformed the conventional classifiers, with Random Forest achieving the highest accuracy of 98.9%, followed by XGBoost and ANN at 97.8%. The most significant predictors were VFA, BMI, and WHR. The results highlight the strong potential of ensemble ML methods to deliver powerful, precise, and scalable tools for early NCD risk assessment, supporting their integration into clinical decision-support systems.

Keywords: Non-communicable diseases, Ensemble machine learning, Random Forest, XGBoost, Anthropometric indicators, Predictive modeling

ENHANCED PREDICTION OF NON-COMMUNICABLE DISEASES USING ENSEMBLE MACHINE LEARNING APPROACHES

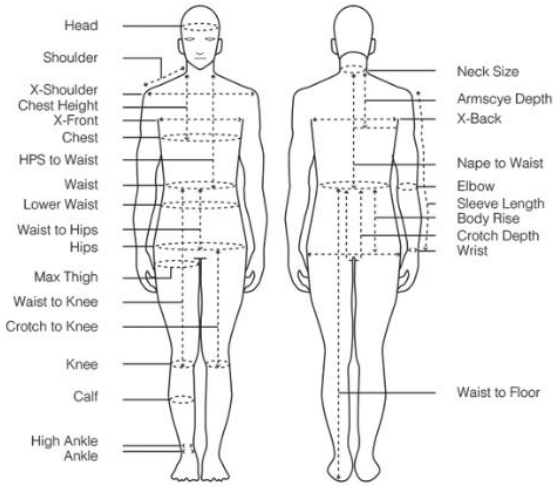
Lenuga Mahendiran[1,2], Banujan Kuhaneshwaran[1,3], Luxshi Karunakaran[4], Vadivel Abishethvarman[1*], Senthan Prasanth .[4,5], and Banage. T.G.S. Kumara[1,6]

[1]Faculty of Computing, Sabaragamuwa University of Sri Lanka [2]School of Computer Science, Northumbria University, United Kingdom,
[3] Faculty of Science and Engineering, Southern Cross University, Australia [4]Faculty of Applied Sciences, Sabaragamuwa University of Sri Lanka
[5] Faculty of Engineering and Applied Science, Memorial University of Newfoundland, Canada
*Abishethvarman@gmail.com

Introduction

Non-communicable diseases (NCDs) account for over **70% of global deaths**, posing a major public health challenge.

In Sri Lanka, rising cases of cardiovascular and metabolic disorders highlight the need for **early detection tools** that are both affordable and accessible.

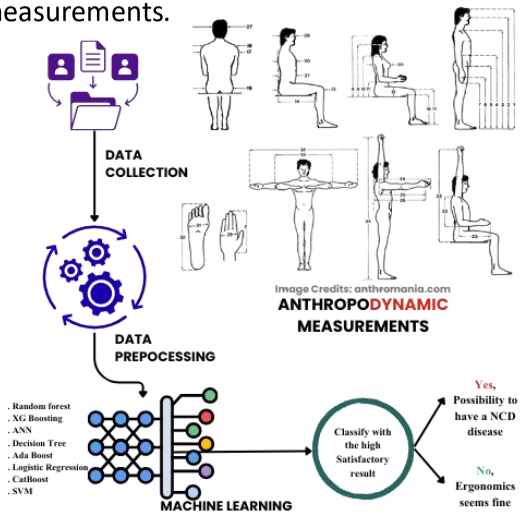


ANTHROPOMETRIC DATA

MEASUREMENTS OF THE SIZE, SHAPE, AND PROPORTIONS OF THE HUMAN BODY.

Methodology

Machine Learning (ML) provides a promising path by enabling **data-driven prediction of disease risk** using simple anthropometric measurements.



Step	Details
Participants	300 individuals from Jaffna Teaching Hospital & Sabaragamuwa University
Data Collected	Anthropometric measurements: height, weight, waist, hip, BMI, etc.
Preprocessing	Noise removal, normalization, outlier correction
Models Tested	Random Forest, XGBoost, AdaBoost, CatBoost, ANN, Decision Tree, LR, SVM
Optimization	Grid search with cross-validation
Evaluation	Accuracy & performance comparison across models

Results

Algorithms	Accuracy (%)	Mean Squared Error [MSE] (%)	Absolute Squared Error [ASE] (%)
Random Forest	98.90	1.09	1.09
XG Boosting	97.80	2.19	2.19
ANN	97.80	2.19	2.19
Decision Tree	96.05	3.94	3.94
Ada Boost	93.40	6.59	6.59
Logistic Regression	88.52	11.47	11.47
Cat Boost	87.91	12.08	12.08
SVM	85.24	14.75	14.75

Table 1: Algorithm's accuracy, mean squared error and absolute squared error

Conclusion

Random Forest achieved 98.9% accuracy. Ensemble ML models can reliably predict NCD risk using simple body measurements, enabling **affordable and practical early screening** in resource-limited healthcare settings.

Enhancing Paddy Leaf Disease Classification through Integrated Deep and Handcrafted Features

Jegatheesan Laavanja*, Amirthalingam Ramanan

Department of Computer Science, University of Jaffna, Sri Lanka

*laavanja1999@gmail.com

This research presents a paddy leaf disease classification framework that integrates deep learning features with handcrafted features to improve accuracy and generalisation. The proposed model combines the features extracted from the VGG-19 convolutional block 5 (conv5) with texture and shape information extracted from the Gray-Level Co-Occurrence Matrix (GLCM) and Hu Moments. This fusion makes the system to extract high-level and low-level features from the diseased regions, addressing the limitations of approaches that rely solely on deep or handcrafted features. To enhance the performance, two feature selection methods – Particle Swarm Optimization (PSO), a wrapper-based method, and Mutual Information (MI), a filter-based method were applied to remove redundant or irrelevant features. Then the selected features were used and classified for the diseases using fine-tuned XGBoost classifier. Experiments were conducted on two publicly available Kaggle datasets containing six and five paddy disease categories, including bacterial leaf blight, brown spot, blast, sheath blight, and others. The proposed framework achieved classification accuracies of 99.24% for the six-class dataset and 99.60% for the five-class dataset, outperforming existing methods reported in recent literature. The strong performance of the Conv5 + GLCM combination, along with the superior feature selection capability of MI, highlights the effectiveness of the hybrid approach. From this, we can make sure that integrating deep and handcrafted features, followed by feature selection methods, offers an optimized and efficient solution for the paddy disease classification.

Keywords: Paddy leaf disease classification, GLCM, PSO, XGBoost, Feature Fusion

Enhancing Paddy Leaf Disease Classification through Integrated Deep and Handcrafted Features

J. Laavanja and A. Ramanan
Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka
laavanja1999@gmail.com, a.ramanan@univ.jfn.ac.lk

Introduction

Paddy cultivation plays a vital role globally, serving as a major source of food and income for a large portion of the population. However, its productivity is frequently threatened by diseases such as bacterial leaf blight, brown spot, and other diseases. Early and accurate detection is therefore essential to support timely disease management and reduce crop losses. However, many recent studies rely solely on either deep learning features or handcrafted features, which prevents them from fully capturing both semantic and textural characteristics in leaf images. To address this limitation, this study integrates deep learning features with handcrafted features, supported by feature selection and machine learning techniques, to develop a more accurate, efficient, and reliable paddy leaf disease classification system.

Methodology

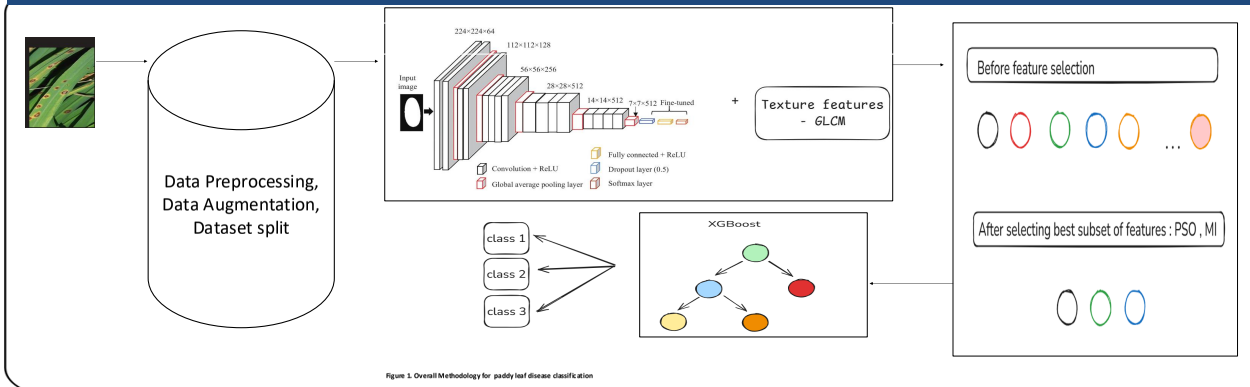


Figure 1. Overall Methodology for paddy leaf disease classification

Results & Discussions

Table 1: Dataset Statistics

Dataset	#Diseases	# Training Imgs/class	# Testing Imgs/class
[1]	6	350	44
[2]	5	510	50

Table 2 : Features Tested using XGBoost Classifier

Features	Dataset 1 [1]	Dataset 2 [2]
Conv5 (VGG-19)	95.45 %	98.00 %
GLCM	86.36 %	82.00 %
Hu-Moments	75.76 %	70.40 %
Conv5 + GLCM	**98.86 %**	**99.20 %**

Table 3: Test Results After Feature Selection for the Proposed Methodology

Conv5+GLCM	Dataset 1 [1]	Dataset 2 [2]
#Features selected out of total features	**MI - 400/532**	**MI - 420/532**
	PSO - 327/532	**PSO - 332/532**
MI	**99.24 %**	**99.60%**
PSO	97.72 %	96.8 %

Table 4: Performance Comparison of Accuracies

Authors	Dataset	Accuracy
Kumar & Sakthivel (2025) [3]	[1]	96.70 %
Alsakar *et al.* (2024) [4]	[2]	99.40 %
Upadhyay & Dwivedi (2025) [5]	[2]	99.30 %
Ours	**[1]**	**99.24 %**
	[2]	**99.60 %**

Conclusion

- The proposed model achieved superior performance, reaching 99.24% and 99.60% accuracy on two benchmark Kaggle datasets, outperforming existing state-of-the-art methods.
- Demonstrates strong potential for real-world agricultural applications.

References

1. Kaggle Rice Leafs Disease Dataset https://www.kaggle.com/datasets/dedeikhsandwisaputra/rice-leafs-disease-dataset
2. Kaggle rice_diseases_using_CNN_and_SVM Dataset : https://www.kaggle.com/datasets/rajeshbhattacharjee/rice-diseases-using-cnn-and-svm
3. B. N. Kumar and S. Sakthivel, "Rice leaf disease classification using a fusion vision approach," Scientific Reports, vol. 15, no. 1, Mar. 2025, https://doi.org/10.1038/s41598-025-87800-3.
4. Y. M. Alsakar, N. A. Sakr, and M. Elmogy, "An enhanced classification system of various rice plant diseases based on multi-level handcrafted feature extraction technique," Scientific Reports, vol. 14, no. 1, Dec. 2024, doi: https://doi.org/10.1038/s41598-024-81143-1.
5. S. K. Upadhyay and A. K. Dwivedi, "Rice Plant Disease Diagnosis Using SqueezeNet and Deep Transfer Learning," Journal of Phytopathology, vol. 173, no. 3, May 2025, doi: https://doi.org/10.1111/jph.70092.

Enhancing Paddy Leaf Disease Classification: A Dual Stream CNN and Handcrafted Feature Approach

Tharmini Kiriharan[1,2,*], Amirthalingam Ramanan[3], Thanikasalam Kokul[3], Subha Fernando[2]

[1]Department of Computer Science, Eastern University, Sri Lanka,
[2]Department of Computational Mathematics, University of Moratuwa, , Sri Lanka,
[3]Department of Computer Science, University of Jaffna,, Sri Lanka

[*]tharminik@esn.ac.lk

Rice is a vital global food crop, yet its production faces significant threats from leaf diseases that reduce both yield and quality. While automated disease classification models exist, current approaches face fundamental limitations: Deep learning approaches overlook critical texture patterns, and handcrafted features based approaches, which lack high level semantic understanding. To address these limitations, this approach proposes a hybrid dual stream architecture integrating Haar wavelets with EfficientNetV2M, achieving state-of-the-art performance across multiple benchmark datasets. Haar wavelets effectively capture both spatial and frequency characteristics, enabling explicit encoding of texture, edge, and structural information, while EfficientNetV2M extracts hierarchical semantic features. This method processes wavelet transformed images via convolutional layers and fuses them with EfficientNetV2M's features. This fusion enables the model to simultaneously leverage high level semantic context and detailed local patterns critical for distinguishing visually similar diseases. This approach was evaluated on three benchmark datasets: Mendeley, UCI, and RiceLeafV3 encompassing seven distinct (Bacterial Blight, Blast, Brown Spot, Leaf Scald, Tungro, Leaf Smut and Narrow Brown Spot) rice diseases. The model achieved superior testing accuracies of 99.94%, 96.44%, and 96.68% respectively, outperforming all existing state-of-the-art methods on each dataset. This work demonstrates that the fusion of wavelet based texture analysis and deep learning creates a more interpretable and accurate model for precision agriculture, offering farmers reliable disease classification while advancing computer vision methodologies for agricultural applications.

Keywords: EfficientNetV2M, Wavelet features, Dual stream, Paddy leaf diseases

Enhancing Paddy Leaf Disease Classification A Dual-Stream CNN and Handcrafted Feature Approach

Tharmini Kiriharan[1,2,*], Amirthalingam Ramanan[3], Thanikasalam Kokul[3], Subha Fernando[2]

[1]Department of Computer Science, Eastern University, Sri Lanka, [2]Department of Computational Mathematics, University of Moratuwa, Sri Lanka
[3]Department of Computer Science, University of Jaffna, Sri Lanka

SLSWCS Student Workshop on Computer Science

Introduction

Rice is a vital global food crop threatened by leaf diseases that reduce yield and quality. Current automated classification faces limitations: deep learning overlooks critical texture patterns [1], while handcrafted features lack semantic understanding.

To overcome this, we present a novel hybrid architecture integrating Haar wavelet decomposition with EfficientNetV2M. This dual-stream approach uniquely preserves high-frequency texture information fused with deep features, enabling superior distinction between visually similar diseases for robust agricultural diagnostics.



Objective

To develop a hybrid dual-stream model integrating Haar wavelet texture with EfficientNetV2M deep features for classifying visually similar rice leaf diseases.

Methodology

1. **Input Preparation**
   paddy leaf image of size 224×224×3 fed simultaneously to both processing streams.

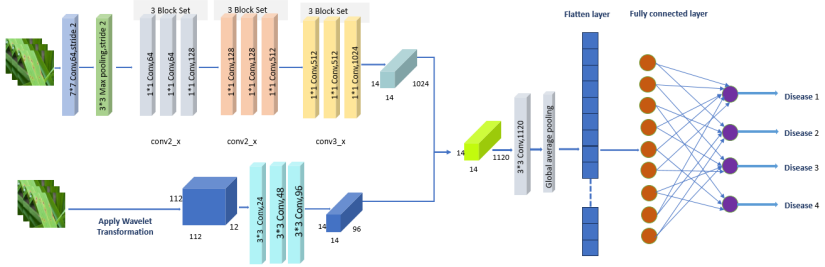
2. **Deep Learning Stream (EfficientNetV2M)**
   Features are extracted from block6b_add layer, combining multiscale semantic patterns with preserved spatial details via skip connections.

Methodology

3. **Handcrafted Feature Stream (Haar Wavelet)**
   Applies Haar wavelet transformation, decomposing into 4 sub-bands per RGB channel (12 total maps); processes through convolutional layers to learn compact texture descriptors from frequency components.

4. **Feature Fusion**
   Wavelet features reshaped to match deep feature dimensions; direct concatenation of both streams; refined using 3×3 convolution with 1120 filters.

5. **Classification**
   Global average pooling compresses features; fully connected layer produces final disease classification.



Results

Performance Comparison with SOTA Approaches

Dataset	Method	Author (Year)	Accuracy
UCI [2]	SVM + HOG	Pothen et al. (2021)	94.60%
	ResNetV2	Sowmyalakshmi et al. (2021)	94.20%
	Pretrained DenseNet	Chen et al. (2020)	92.22%
	Random Forest	Mekha and Teeyasuksaet (2021)	69.44%
	Dual Stream Approach	Ours	**96.44%**
Mendeley [3]	KLD-CNN	Navaneetha et al. (2025)	99.84%
	Custom VGG16	Ritharson et al. (2024)	99.70%
	Vision Transformer-RDTNet	Kumar et al. (2024)	99.54%
	Darknet53 + SVM	Padhi et al. (2025)	99.43%
	Dual Stream Approach	Ours	**99.94%**
RiceLeafV3 [4]	ResNet152V2	Waghmare et al. (2024)	95.04%
	VGG19	Latif et al. (2022)	96.08%
	Random Forest	Benita et al. (2023)	89.83%
	Dual Stream Approach	Ours	**96.68%**

Experimental Setup

1. **Dataset Specifications:**
   UCI (120 images) :Bacterial Blight (40), Brown Spot (40), and Leaf Smut (40)
   Mendeley (5932 images) :Bacterial Blight (1584), Blast (1440), Brown Spot (1600), and Tungro (1308)
   RiceLeafV3 (2710 images) :Bacterial Blight (438), Blast (456), Brown Spot (466), Leaf Scald (449), and Narrow Brown Spot (440)

2. **Training Parameters:**
   Epochs: 50. Batch Size: 16, Optimizer: Adam (learning rate = 0.0001), Loss Function: Categorical Cross-Entropy

3. **Implementation Details:**
   Input Resolution: 224×224×3, Early Stopping: 15 epochs, Weight Initialization: ImageNet

Conclusion

This study presents a dual-stream framework for rice disease classification by integrating Haar wavelet textures with EfficientNetV2M. The hybrid model achieves state-of-the-art performance across multiple datasets, excelling at distinguishing visually similar diseases. By preserving both high-frequency details and semantic features, it proves highly suitable for real-world agricultural applications.

References

[1] Petchiammal, A., Murugan, D., 2025. Automated paddy leaf disease identification using visual leaf images based on nine pre-trained models approach. Procedia Computer Science 252, 118–126. doi:10.1016/j.procs.2024.12.013.

[2] Shah, J., Prajapati, H., Dabhi, V., 2017. Rice Leaf Diseases. UCI Machine Learning Repository. URL: https://doi.org/10.24432/C5R013

[3] Sethy, P.K., 2024. Rice leaf disease image samples. URL: https://data.mendeley.com/datasets/fwcj7stb8r/2, doi:10.17632/fwcj7stb8r.2

[4] Tjikoa, A.F., 2022. Rice leaf 5 diseases. doi:10.34740/KAGGLE/DSV/3726583

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)

Poster ID: P14

Evaluating Layout Detection Models For Historical Tamil Printed Texts

Suthakar Sivashanth*, Kengatharaiyer Sarveswaran, Eugene Yugarajah Andrew Charles

Department of Computer Science, University of Jaffna, Sri Lanka

*ssivashanth@univ.jfn.ac.lk

Tamil, one of the oldest living classical languages, boasts a rich literary heritage preserved over centuries in printed forms. Digitizing these materials is essential for their preservation and transmission to future generations. A critical step in this process is document layout detection, which enables accurate Optical Character Recognition (OCR) and information retrieval from documents with complex layouts. This study evaluates multiple layout detection models on a benchmark dataset of diachronic Tamil documents that we curated, with the aim of identifying the most suitable model for integration into our ongoing OCR system development. We evaluated models pre-trained on the DocLayNet dataset, including YOLO-based architectures, Docling variants, and custom-trained Faster R-CNN models. Evaluation metrics such as precision, recall, and F1-score were used to analyze model performance. Notably, the docling-layout-heron model achieved high recall (0.714) and Faster R-CNN (R50-FPN) (3x schedule) model achieved high precision (0.661). However, certain limitations were observed, including over-segmentation of tables, predictions on unannotated elements such as stamps, stickers, and handwritten markings, difficulty in distinguishing between plain text and table regions due to variations in table formats, and occasional incorrect bounding box predictions caused by damaged pages. Despite these challenges, the findings provide valuable insights for refining layout detection pipelines for complex historical documents.

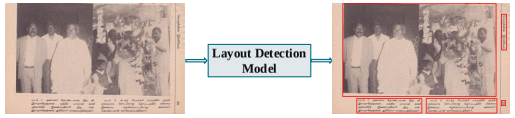
Keywords: Layout Detection, Tamil, Benchmark Dataset, DocLayNet

Evaluating Layout Detection Models for Historical Tamil Printed Texts

Suthakar **Sivashanth**, Kengatharaiyer **Sarveswaran**, and Eugene Yugarajah Andrew **Charles**
Department of Computer Science, University of Jaffna
ssivashanth@univ.jfn.ac.lk

Introduction

- Tamil, a classical language with a long literary history, preserves valuable knowledge in diverse printed documents.
- Printed documents require digitization for long term preservation and accessibility.
- Historical Tamil documents present significant challenges for digitization due to complex layouts, printing technologies, document degradation, and noise, as well as the difficulty of understanding their historical and cultural context, which is essential for post-processing.
- Layout detection is essential for building a reliable Tamil OCR system for digitizing these documents, enabling accurate and structured text extraction.



Input image — Output image with bounding-box segmentation

Motivation

- Although several layout detection models are available, there is currently no benchmark dataset for evaluating these models on complex document layouts and the diachronic nature of the documents.
- This research aims to create a benchmark dataset and systematically evaluate existing layout detection models to identify the most suitable model for layout detection in historical Tamil documents.

Data Collection

- 164 images have been collected from the Jaffna University Library and the Noolaham Foundation, covering Tamil printed documents from 1850 to the present at approximately 10-year intervals.
- The dataset includes a varied range of complex document layouts, covering cover pages, imprint pages, tables of contents, prefaces, plain text pages, tables, text with tables, text with images, and advertisements.



Figure1: Examples of Collected Images

Methodology

Research workflow

1. **Scanning the Documents.**
   Sources: books, magazines, pamphlets, and newspapers.
   Resolution: 300 DPI (dots per inch).
2. **Creating the Ground Truth Bounding Boxes.**
   Annotation tool: LabelImg.
3. **Benchmarking the Layout Detection Models.**
   Models evaluated: YOLOv11l, YOLOv12l, Docling variants (egret-medium, egret-large, egret-xlarge, heron, heron-101; Faster R-CNN (R50-FPN, R101-FPN; 3xschedule).
   Pre-training dataset: DocLayNet.

Benchmarking Steps

Intersection over Union (IoU)-based Matching
- Compare predicted and ground-truth boxes
- Compute overlap and union areas
- Match if IoU ≥ 0.5

↓

Classification of Detected Bounding Boxes
- True Positive (TP): correct matches
- False Positive (FP): extra or incorrect detections
- False Negative (FN): missed ground-truth boxes

↓

Evaluation Metrics
- Precision, Recall, F1-score

Results

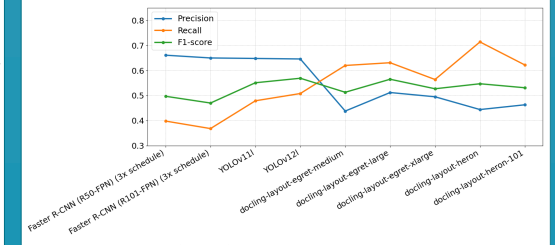


Figure2: Performance of Layout Detection Models on Printed Tamil Documents

Discussion

- Complex layouts in Tamil documents significantly impacts model precision and recall of the models.
- Faster R-CNN variants miss many layout elements, especially in mixed or degraded pages.
- YOLO-based models offer improved precision and recall compared to the Faster R-CNN model, making them suitable for general-purpose layout detection.
- Docling models reliably identify most components in complex documents.
- Non-ground-truth elements (e.g., stamps, stickers, handwriting) and noise trigger additional false positives.
- Table variability and page degradation lead to frequent confusion between tables and text, causing incorrectly identified bounding boxes.



Figure3: Examples of Challenging Images

Conclusion

- The pre-trained Docling models demonstrated strong baseline performance for layout detection in historical Tamil printed documents.
- Among them, the docling-layout-heron model achieved the highest recall, making it ideal for OCR pipelines that require detailed structural extraction in these documents.

References

1. Tzutalin, "LabelImg," 2015, accessed: Oct. 5, 2024. [Online]. Available: https://github.com/tzutalin/labelImg
2. ppaanngggg, "YOLO-DocLayNet: Repository for YOLO models pre-trained on DocLayNet," 2022, accessed: Jun. 19, 2025. [Online]. Available: https://github.com/ppaanngggg/yolo-doclaynet
3. I. D. Search, "Ds4sds models on hugging face," accessed: Jul. 12, 2025. [Online]. Available: https://huggingface.co/ds4sd/models
4. Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," 2019. [Online]. Available: https://github.com/facebookresearch/detectron2
5. B. Pfitzmann, C. Auer, M. Dolfi, A. S. Nassar, and P. Staar, "Doclaynet: A large human-annotated dataset for document-layout segmentation," *in Proc. 28th ACM SIGKDD Conf. Knowledge Discovery and Data Mining* (KDD '22). New York, NY, USA: Association for Computing Machinery, 2022, p. 3743-3751. [Online]. Available: https://doi.org/10.1145/3534678.3539043

Acknowledgement

- This research study was carried out under the DigSAL project, supported by the German Academic Exchange Service (DAAD) and funded by the Federal Ministry for Economic Cooperation and Development (BMZ), Germany through SDG Partnerships (2024-2027).
- We thank the Library of the University of Jaffna and the Noolaham Foundation for providing access to the historical Tamil documents utilized in this study.

Explainable AI for Heart Disease Prediction Using CNNs

Preethy Srinevasan* and Eugene Yugarajah Andrew Charles

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*2020sp183@univ.jfn.ac.lk

Heart disease remains a leading cause of mortality worldwide, causing millions of deaths each year. Accurate and timely diagnosis is essential for effective treatment and prevention, as early detection can reduce complications and improve patient outcomes. Electrocardiography (ECG) provides critical diagnostic information, yet interpreting complex 12-lead signals manually is challenging and prone to human error. Recent advances in deep learning, particularly Convolutional Neural Networks (CNNs), enable automatic learning of discriminative patterns from ECG data for improved cardiac assessment. However, most CNN and CNN-LSTM models emphasize predictive accuracy while offering limited interpretability. This study presents an InceptionTime-based one-dimensional CNN for five-class heart-disease prediction using the PTB-XL dataset. To address the black-box limitation, post-hoc Explainable Artificial Intelligence (XAI) methodsâĂŤGrad-CAM and LIMEâĂŤwere applied to highlight model reasoning. Grad-CAM visualizations revealed clinically significant regions, including the QRS complex and ST-segment, while LIME identified lead-specific waveform features most influential for predictions, providing interpretable insights to support clinical decision-making. Unlike intrinsic attention- or case-based approaches, the proposed model integrates complementary post-hoc explanations linking learned representations with physiological meaning. The model achieved 88.12% accuracy and 86.59% AUC, demonstrating competitive performance compared with existing CNN, LSTM, and Transformer-based methods on the PTB-XL dataset. By combining robust predictive capability with transparent interpretability, this framework enhances clinical trust and supports AI-assisted cardiovascular diagnostics.

Keywords: Heart disease prediction, Explainable AI, InceptionTime, Grad-CAM, LIME

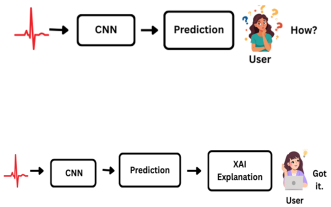
Explainable AI for Heart Disease Prediction Using CNNs

S. Preethy and E. Y. A. Charles
Department of Computer Science, Faculty of Science, University of Jaffna
2020sp183@univ.jfn.ac.lk, charles.ey@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

1. Introduction

- Heart disease is one of the leading global causes of death.
- ECG signals contain important diagnostic information but are difficult to interpret manually due to complex waveform patterns.
- Convolutional Neural Networks (CNNs) can automatically extract meaningful ECG features.
- However, CNNs act as "black-box" models, reducing clinical trust.
- This study uses a customized InceptionTime-based 1D CNN combined with Explainable AI (Grad-CAM & LIME) to provide transparent and interpretable heart disease predictions.



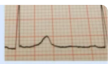
2. Objectives

- Develop a CNN-based model to predict heart disease from ECG signals.
- Evaluate the model using accuracy, AUC, Loss and per-class precision, recall, F1-score.
- Integrate and analyze Explainable AI (XAI) techniques for model interpretability.

3. Dataset

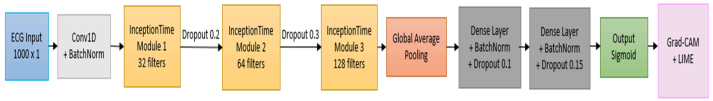
PTB-XL ECG dataset
Wagner et al Sci Data 7, 151 (2021)

21799 12-lead ECGs
from 18869 patients



- 10-second signals, 100 Hz (downsampled)
- Classes: NORM, MI, STTC, CD, HYP
- Annotations: SCP-ECG codes (cardiologist-verified)
- Format: WFDB + metadata, recommended train/test split

4. Methodology



- ECG Signals: Preprocessed by splitting into training, validation, and test sets, handling missing values, label encoding, addressing class imbalance, reshaping signals, and converting data types.
- Model: 1D InceptionTime CNN (~706K parameters) for 5-class classification (CD, HYP, MI, NORM, STTC) using 100 Hz signals.
- Training: Adam optimizer, 100 epochs, batch size 32, Binary Cross-Entropy loss, with EarlyStopping and ReduceLROnPlateau for regularization. Class imbalance handled using class weights.
- Explainable AI (XAI): Grad-CAM to visualize important ECG regions, LIME for feature attribution.
- Evaluation Metrics: Accuracy, AUC, loss, precision, recall, F1-score (per class).

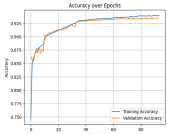
5. Results & Discussion

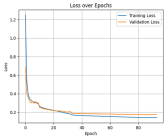
Study (Authors, Year)	Model / Method	Accuracy (%)	Use of Interpretability / XAI
[1] Elyamani et al. (2024)	Deep Residual 2D CNN	89.87	No
[2] Strodthoff et al. (2020)	Ensemble(CNN, ResNet, GRU, InceptionTime)	84.7	No
[3] Muñoz et al. (2024)	CNN, CNN-LSTM, CNN-LSTM-Attention	72.52-78.12	No
[4] Geng et al. (2023)	Multi-Task Learning + CoT Attention Transformer	88.7	No
[5] Kang et al. (2025)	xLSTM-ECG (Extended LSTM)	87.59	No
[6] Sethi et al. (2025)	ProtoECGNet	Not specified	Case-based interpretability
[7] Rodriguez & Nafea (2025)	Hierarchical Attention Network	75.59	Attention mechanism visualization
[8] Cheng et al. (2023)	MSW-Transformer (Multi-Scale Shifted Windows)	89.8	Attention mechanism visualization
[9] Ayano et al. (2024)	Hybrid Multichannel Deep Learning	89.84	Grad-CAM++ and SHAP
Proposed model (2025)	**InceptionTime 1D CNN**	**88.12**	**Grad-CAM + LIME**

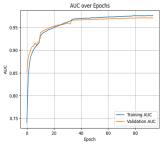
Overview of 5-Class ECG Prediction Models on PTB-XL Dataset

Test Results

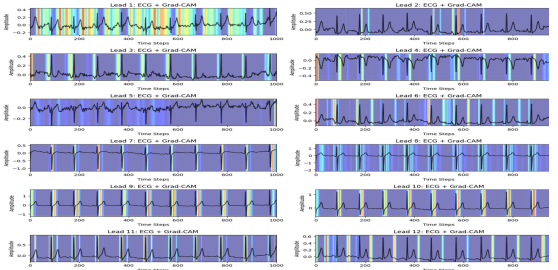
Accuracy	88.12%
AUC	0.8659
Loss	0.3541



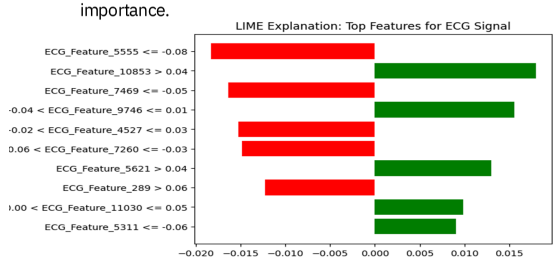




Grad-CAM (Gradient-weighted Class Activation Mapping)
Highlights which parts of the ECG signal the model focused on for its prediction.



LIME (Local Interpretable Model-Agnostic Explanations)
Explains individual predictions by showing feature importance.



6. Conclusion

The InceptionTime CNN model accurately classifies 12-lead ECG signals into five heart disease classes. The integration of Explainable AI techniques, Grad-CAM and LIME, provides clear visual insights, improving interpretability and clinical trust. This approach demonstrates potential for early detection and diagnosis of heart conditions, emphasizing the value of combining high-performance models with explainability in real-world healthcare applications.

7. References

[1] H. A. Elyamani, M. A. Salem, F. Melgani, and N. Yhiea, "Deep residual 2D convolutional neural network for cardiovascular disease classification," Scientific Reports, vol. 14, no. 1, p. 22040, 2024.
[2] N. Strodthoff, P. Wagner, T. Schaeffter, and W. Samek, "Deep learning for ECG analysis: Benchmarks and insights from PTB-XL," IEEE Journal of Biomedical and Health Informatics, vol. 25, no. 5, pp. 1519–1528, 2020.
[3] A. Muñoz, J. Torres-Santamaría, M. Iregui, E. Romero, and A. Cruz-Roa, "Comparative analysis of deep neural network architectures for heart disease classification in electrocardiography signals," in Proc. 20th Int. Symp. Med. Inf. Process. Anal. (SIPAIM), Nov. 2024, pp. 1–4.
[4] Q. Geng, H. Liu, T. Gao, R. Liu, C. Chen, Q. Zhu, and M. Shu, "An ECG classification method based on multi task learning and CoT attention mechanism," Biomedical Signal Processing and Control, vol. 87, p. 105523, 2023.
[5] L. Kang, X. Fu, J. Vazquez-Corral, E. Valveny, and D. Karatzas, "xLSTM-ECG: Multi-label ECG classification via feature fusion with xLSTM," Computer Vision Center, Universitat Autònoma de Barcelona, Spain, 2025.
[6] S. Sethi, D. Chen, T. Statchen, M. C. Burkhart, N. Bhandari, B. Ramadan, and B. Beaulieu-Jones, "ProtoECGNet: Case-based interpretable deep learning for multi-label ECG classification with contrastive learning," arXiv preprint arXiv:2504.08713, Apr. 2025.
[7] M. P. Rodriguez and M. Nafea, "Hierarchical attention network for interpretable ECG-based heart disease classification," arXiv preprint arXiv:2504.03703, Mar. 2025.
[8] R. Cheng, Z. Zhuang, S. Zhuang, L. Xie, and J. Guo, "MSW-Transformer: Multi-scale shifted window transformer for ECG-based arrhythmia classification," Computers in Biology and Medicine, vol. 167, p. 107641, 2023.
[9] Y. M. Ayano, F. Schwenker, B. D. Dufera, T. G. Debelee, and Y. G. Ejegu, "Interpretable hybrid multichannel deep learning model for heart disease classification using 12-leads ECG signal," IEEE Access, 2024.

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P16

Ordinal-Aware Contrastive Learning for Enhanced Diabetic Retinopathy Grading

Alagiah Suthakaran*, Siyamalan Manivannan

Department of Computer Science, University of Jaffna, Sri Lanka

*suthakarana@esn.ac.lk

Diabetic retinopathy (DR) is a leading cause of preventable blindness worldwide. Automated DR grading is inherently challenging due to the ordinal nature of disease stages, where misclassifications across distant stages (e.g., No DR vs. Severe DR) are more critical than between adjacent ones. Existing contrastive and classification-based methods often fail to model these ordinal relationships between severity stages, leading to clinically inconsistent grading outcomes. To address this, we propose an Order-aware Contrastive Loss (OCL) that incorporates inter-class distance weighting to enforce ordinal consistency during representation learning. Furthermore, we introduce a multi-branch architecture built upon a DenseNet-121 backbone, which integrates OCL with two complementary objectives in a unified multi-head framework. The entire network is jointly optimized to enhance both predictive accuracy and ordinal alignment. To the best of our knowledge, this is the first contrastive learning framework that explicitly integrates ordinal distance weighting for diabetic retinopathy grading. The proposed approach is evaluated on two benchmark datasets: EyePACS, a large-scale dataset containing 35,126 training and 42,670 test images across five stages, and APTOS2019, which includes 3,662 images. On EyePACS, our method achieves 86.93% accuracy, 0.8644 quadratic weighted kappa (QWK), and 0.8694 F1-score, surpassing the previous state-of-the-art (84.55% accuracy, 0.8631 QWK, 0.8440 F1-score). Similar improvements are observed on APTOS2019. These results highlight the effectiveness of incorporating order-aware representation learning with multi-objective optimization on a strong DenseNet based architecture, setting a new benchmark for ordinal DR classification.

Keywords: Ordinal Classification, Contrastive learning, Diabetic Retinopathy

Ordinal-Aware Contrastive Learning for Enhanced Diabetic Retinopathy Grading

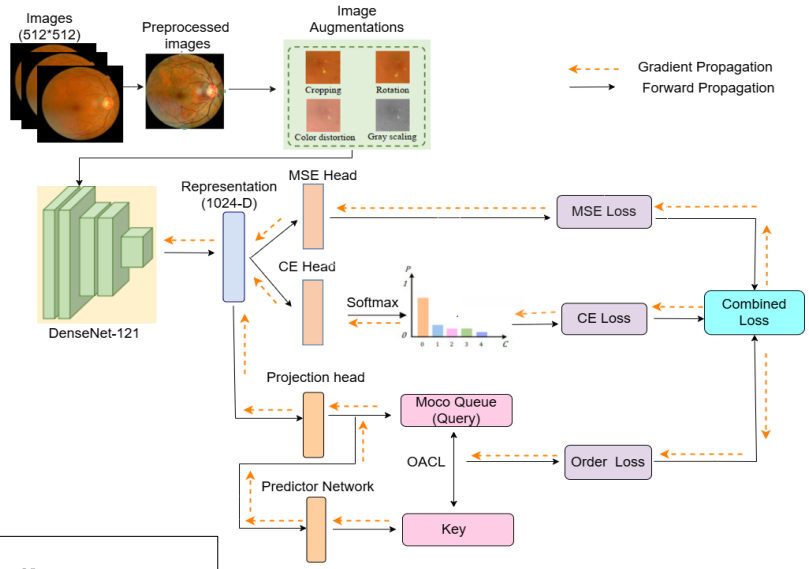
Alagiah Suthakaran[1], Manivannan Siyamalan[2]

[1]Faculty of Graduates Studies, University of Jaffna, [2]Department of Computer Science, University of Jaffna

Introduction

Diabetic retinopathy (DR), a leading cause of preventable blindness, requires automated grading methods that maintain both accuracy and clinical consistency. Traditional approaches often overlook the ordinal structure of DR severity, resulting in clinically significant misclassifications. To address this, we propose an Order-aware Contrastive Loss (OACL) that encodes inter-class distances to learn representations aligned with disease progression. Combined with a multi-branch DenseNet-121 that jointly optimizes contrastive, cross-entropy, and ordinal objectives, our framework improves both predictive performance and ordinal coherence, demonstrating strong effectiveness in DR grading.

Proposed Methodology



MSE Loss:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2$$

CE Loss:

$$\mathcal{L}_{\text{CE}} = -\sum_{n=1}^{N}\sum_{k=1}^{C} y_{n,k}\log(\hat{y}_{n,k})$$

Objective Function:

$$\mathcal{L}_{\text{Total}} = \lambda_1 * \mathcal{L}_{\text{MSE}} + \lambda_2 * \mathcal{L}_{\text{CE}} + \lambda_3 * \mathcal{L}_{\text{OCL}}$$

Ordinal Contrastive Loss (L_{OCL}):

$$\mathcal{L}_{\text{OCL}} = -\frac{1}{B}\sum_{i=1}^{B}\frac{1}{|P(i)|}\sum_{p\in P(i)}\log\frac{\exp(x_i^\top x_p/\tau)}{\sum_{k=1}^{K} w_{ik}\sum_{j\in N_k(i)}\exp(x_i^\top x_j/\tau)}$$

Batch size; # of positive samples; Positive pair similarity term; Weighted negative pair similarity term

Class distance weights for each pair (x_i, x_j): $w_{ij} = \frac{|y_i - y_j|^\gamma}{|B_{y_j}|}$

Datasets & Experimental Setup

	EyePACS		APTOS-2019	
Class	Training	Testing	Training	Testing
No DR	25,810	31,403	1,444	361
Mild	2,443	3,042	296	74
Moderate	5,292	6,282	799	200
Severe	873	977	154	39
Proliferative DR	708	966	236	59
Total	**35,126**	**42,670**	**2929**	**733**

- EyePACS: The entire training set was used for training, and the entire testing set was used for evaluation.
- APTOS-2019: A 5-fold cross-validation protocol was employed.

Results and Analysis

Performance comparison with recent approaches on EyePACS dataset

Method	**Single Eye**				**Both Eyes**			
	Acc.	BAcc	QWK	F1	Acc.	BAcc	QWK	F1
Single Models								
Siamese Network [1]	85.68	-	0.8320	-	86.42	-	0.8530	0.8360
MediDRNet [2]	-	-	-	-	63.27	-	0.5409	0.6361
Zoom-in-Net [3]	-	-	-	-	-	-	0.8490	-
Key Components [4]	-	-	0.8310	-	-	-	**0.8581**	-
Proposed Method (Class-weighted)	83.00	**65.80**	0.8446	0.8314	83.61	**67.86**	0.8552	**0.8362**
Proposed Method	**86.34**	59.60	**0.8456**	0.8614	**86.58**	**60.96**	0.8571	**0.8658**

Ensemble (Both Eyes)				
	Acc.	BAcc	QWK	F1
Siamese Network [1]	84.55	-	0.8600	0.8440
CABNet	-	-	0.8456	-
Zoom-in-Net [3]	-	-	0.8540	-
Key Components[4]	-	-	0.8631	-
Proposed Method	**86.93**	**62.72**	**0.8644**	**0.8694**

Performance comparison with recent approaches on APTOS-2019 dataset

Method	Acc.	BAcc	QWK	F1
KMoCoL [5]	85.20	70.20	0.9129	0.8479
Proposed Method	**85.72**	69.49	**0.9242**	**0.8571**

Conclusion

- We introduced an ordinal-aware contrastive framework that embedded disease-stage distances to learn clinically consistent features.
- Our unified CE–MSE–contrastive formulation improved prediction accuracy while enforcing ordinal structure in the representations.
- The approach established a general strategy for integrating ordinal information into contrastive models, with clear potential for semi-supervised and cross-domain extensions.

References

[1] R. Nirthika, et al."Siamese network based fine grained classification for Diabetic Retinopathy grading," Biomed Signal Process Control, vol. 78, Sep. 2022,
[2] S. Teng, et al.,"MediDRNet: Tackling category imbalance in diabetic retinopathy classification with dual-branch learning and prototypical contrastive learning," Comput Methods Programs Biomed, vol. 253
[3] Z. Wang, et al. ,"Zoom-in-Net: Deep Mining Lesions for Diabetic Retinopathy Detection," Jun. 2017
[4] Y. Huang, et al.,"Identifying the Key Components in ResNet-50 for Diabetic Retinopathy Grading from Fundus Images: A Systematic Investigation," Diagnostics, vol. 13, no. 10, May 2023
[5] L. Li, et al. "KMoCoL: k-Positive Momentum Contrastive Learning for Imbalanced Diabetic Retinopathy Grading," IEEE Trans Instrum Meas, vol. 74, 2025

Multimodal Complementary Fashion Recommendation using FashionCLIP and HBDSCAN

Nigee Hettige*, Rajendran Nirthika

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*nigeehettige20@gmail.com

Fashion recommendation systems enhance user experience by suggesting complementary items to create cohesive outfits, yet managing style diversity in e-commerce remains challenging. This study presents a scalable, context-aware framework for compatibility-driven recommendations, integrating multimodal embeddings with density-based clustering. Visual and textual features are extracted using FashionCLIP to produce 512-dimensional embeddings per modality, reduced to 50 dimensions via Principal Component Analysis (PCA) for computational efficiency. To capture latent compatibility patterns, Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) clusters approximately 33,000 fashion items across 52 article types, effectively handling noise and identifying variable-shaped clusters missed by traditional methods. A two-stage retrieval pipeline leverages text-based cluster retrieval followed by image-based re-ranking, with Facebook AI Similarity Search (FAISS) indexing enabling efficient scalability for real-world deployment. Experimental evaluation with uncertainty quantification demonstrates HDBSCAN significantly outperforms baseline methods, achieving a Silhouette score of 0.3777 compared to KMeans (0.2362 ± 0.0088), Gaussian Mixture Model-diagonal (0.2650 ± 0.0097), Agglomerative Clustering (0.2452), MiniBatch KMeans (0.1960), and Birch (0.2300), with the lowest Davies-Bouldin score (0.9994) indicating superior cluster separation. Statistical significance testing across multiple random initializations and cross-validation folds confirms these performance improvements are robust. This framework provides a scalable, practical solution for real-world e-commerce applications, with user-centered validation identified as a critical next step.

Keywords: FashionCLIP, Fashion Recommendation, Clustering, Embeddings

Multimodal Complementary Fashion Recommendation Using FashionCLIP and HDBSCAN

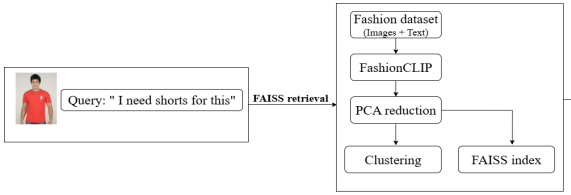
Nigee, H., and Nirthika, R.
Department of Computer Science, University of Jaffna, Jaffna, Sri Lanka
nigeehettige20@gmail.com, nirthika@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

Introduction and Motivation

Fashion recommendation has evolved from simple similarity-based retrieval [1] to compatibility-driven systems that suggest coherent outfit combinations rather than individual items. Multimodal models such as FashionCLIP [2] leverage both visual and textual information to produce rich embedding, improving retrieval performance even with noisy or limited fashion data. Effectively modeling style coherence and cross-category compatibility remains challenging [3], as items must complement each other across different categories while accommodating diverse fashion preferences. This motivates scalable approaches that can uncover latent compatibility patterns in large e-commerce catalogs, enabling personalized, context-aware recommendations that go beyond visual similarity to capture the nuanced relationships between fashion items.

Methodology



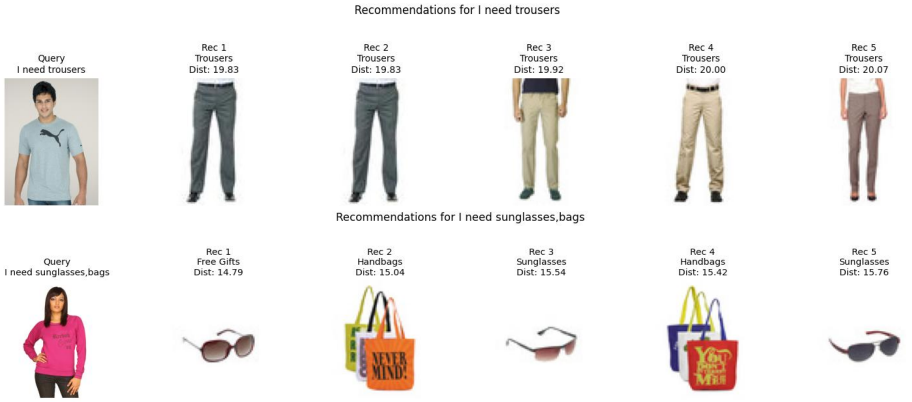


- Dataset of around 33,000 fashion images across 52 categories, with cleaned textual metadata and 70/15/15 for train-validation-test split.
- Multimodal embeddings from FashionCLIP [4] (image + text), reduced via PCA for efficiency.
- Evaluation incorporates multiple random initializations and 5-fold cross-validation to quantify uncertainty and ensure robustness.
- Two-stage retrieval: FAISS-based cluster candidate search followed by image re-ranking, evaluation with clustering metrics.

Results and Discussion

Algorithm	Silhouette	Calinski-Harabasz	Davies-Bouldin
HDBSCAN	**0.3777 ± 0.0000**	**0114.77 ± 00.00**	**0.9994 ± 0.0000**
K Means	0.2362 ± 0.0088	1616.21 ± 45.32	1.6582 ± 0.0421
MiniBatch	0.1960 ± 0.0059	1214.01 ± 35.23	1.8511 ± 0.0612
Agglomerative	0.2452 ± 0.0000	1438.94 ± 00.00	1.6050 ± 0.0000
Birch	0.2300 ± 0.0000	1686.23 ± 00.00	1.6308 ± 0.0000
GMM (diagonal)	0.2650 ± 0.0097	1625.26 ± 52.18	1.5534 ± 0.0856

- HDBSCAN achieves the highest Silhouette score (0.3777), outperforming K-Means++ (0.2362), GMM (0.2650), and Agglomerative Clustering (0.2452), indicating strong cluster cohesion and separation.
- Its Calinski–Harabasz score (114.77) is lower than Birch (1686.23) and GMM (1625.26), reflecting its ability to handle noise and non-spherical clusters.
- Davies–Bouldin index (0.9994) is the lowest among all baselines, confirming clearer cluster distinction.
- HDBSCAN outperforms other clustering methods (MiniBatch, Birch, Agglomerative, GMM) consistently across metrics.
- Recommendations are contextually relevant and complementary, capturing cross-category relationships.



Conclusion

This study presents a multimodal fashion recommendation framework using FashionCLIP embeddings and HDBSCAN clustering, effectively capturing cross-category compatibility and handling noisy, real-world data. Experimental results show that HDBSCAN outperforms baseline clustering methods, achieving superior Silhouette and Davies–Bouldin scores. The framework produces coherent clusters and contextually relevant, complementary recommendations, demonstrating its effectiveness for practical e-commerce applications.

References

[1] Addagarla, S. K., & Amalanathan, A. (2020). Probabilistic unsupervised machine learning approach for a similar image recommender system for E-commerce. Symmetry, vol.12(11), p.1783.
[2] Chia, P. J., Attanasio, G., Bianchi, F., Terragni, S., Magalhaes, A. R., Goncalves, D., & Tagliabue, J. (2022). Contrastive language and vision learning of general fashion concepts. Scientific Reports, vol.12(1), p.18958.
[3] He, R., Packer, C., & McAuley, J. (2016). Learning compatibility across categories for heterogeneous item recommendation. In 2016 IEEE 16th International Conference on Data Mining (ICDM) pp. 937-942.
[4] Balasubramanian, S., Basu, S., & Feizi, S. (2024). Decomposing and interpreting image representations via text in ViTs beyond CLIP. Advances in Neural Information Processing Systems, vol.37, pp.81046-81076.

Poster ID: P18

Multimodal Emotion Recognition with Emotion-Specific Cross-Modal Attention Blocks

Dharshika Gnanaseelan*, Amirthalingam Ramanan

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*dharshikagnanaseelan4@gmail.com

Multimodal emotion recognition is a critical research field that enables the development of human–centered applications, such as healthcare, education, and human–computer interaction. Although existing models achieve reasonable performance, many treat all emotions uniformly, limiting their ability to capture emotion-specific characteristics. To address this challenge, the study proposes a novel multimodal architecture that incorporates Emotion-Specific Cross-Modal Attention Blocks to enhance the model's ability to capture and integrate affective cues across modalities. Each block is dedicated to a single emotion and receives shared keys and values from Wav2vec speech features and prosody features, combined via a learnable gated fusion mechanism that balances the contributions of each modality at every time step. Each block is assigned a separate emotion-specific query, projected from BERT-based text embeddings corresponding to its target emotion, enabling the model to attend to cues most relevant for that emotion. This allows each block to independently learn specialised and discriminative representations. By explicitly modeling emotion-wise interactions through per-emotion queries and attention, the model captures subtle differences between emotions that shared-attention methods may overlook, resulting in more accurate and interpretable predictions. Experiments demonstrate that the proposed model significantly outperforms baseline methods, achieving an unweighted accuracy of 82.12% on the benchmark IEMOCAP corpus across four primary emotions – Angry, Sad, Neutral, and Happy – highlighting its robustness and potential for real-world applications.

Keywords: Multimodal Emotion Recognition, Speech and Text, Deep Learning, Gated Fusion, Emotion-Specific Attention

Multimodal Emotion Recognition with Emotion-Specific Cross-Modal Attention Blocks

G. Dharshika and A. Ramanan
Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka
dharshikagnanaseelan4@gmail.com, a.ramanan@univ.jfn.ac.lk

Introduction

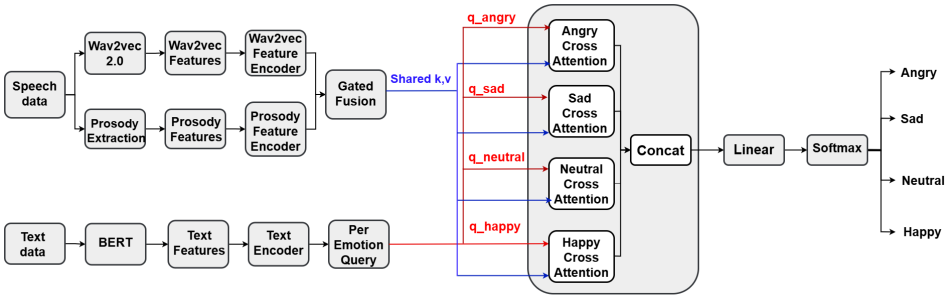
- Understanding human emotions is essential for natural Human–Computer Interaction.
- Multimodal Emotion Recognition combines speech and text to detect emotions more accurately, overcoming the limitations of single-modality approaches.
- Problem: Existing multimodal models treat all emotions equally, failing to capture emotion-specific patterns.
- This study proposes Emotion-Specific Cross-Modal Attention Blocks:
  - Each block focuses on one emotion.
  - Learns emotion-specific patterns from speech and text.
  - Improves discriminative and interpretable emotion recognition.

Dataset

IEMOCAP
- Emotions used: Angry, Sad, Neutral, Happy
- Modalities used: Speech and Text
- Duration: ~12 hours of dyadic conversational interactions



Methodology



Training Strategy: Each Emotion-Specific Attention Block is supervised using Binary Cross-Entropy loss for its corresponding emotion during training.

Test Results

Table 1: Performance comparison of the proposed model with baselines on the IEMOCAP corpus

Methods	Year	Modality	UW Acc (%)
Liu *et al.* [1]	2022	Audio + Text	69.74
Chen *et al.* [2]	2022	Audio + Text	74.30
Wang *et al.* [3]	2023	Audio + Text	75.08
Naderi *et al.* [4]	2023	Audio	75.63
Khan *et al.* [5]	2025	Audio + Text	81.33
Proposed Method	**2025**	**Audio + Text**	**82.12**

Table 2: Ablation study of the proposed model

Model Variant	UW Acc (%)
Text + Wav2vec (without Prosody)	79.23
Simple concatenation instead of gated fusion	80.90
Shared query instead of per-emotion queries	80.88
Single shared attention block instead of emotion-specific blocks	80.23
Proposed Method	**82.12**

Conclusion

- The proposed model captures subtle emotion-specific cues, improving accuracy.
- Achieved state-of-the-art performance, obtaining an unweighted accuracy of 82.12% on the benchmark IEMOCAP dataset.
- Offers strong potential for real-world applications in emotion-aware AI systems, including virtual assistants, intelligent tutoring, and affective healthcare.

References

[1] Y. Liu, H. Sun, W. Guan, Y. Xia, and Z. Zhao, "Multi-modal speech emotion recognition using self-attention mechanism and multi-scale fusion framework," *Speech Commun.*, vol. 139, pp. 1–9, 2022.
[2] W. Chen, X. Xing, X. Xu, J. Yang, and J. Pang, "Key-sparse transformer for multimodal speech emotion recognition," *Proc. ICASSP*, pp. 6897–6901, 2022.
[3] Y. Wang *et al.*, "Multimodal transformer augmented fusion for speech emotion recognition," *Front. Neurorobot.*, vol. 17, art. no. 1181598, 2023.
[4] N. Naderi and B. Nasersharif, "Cross-corpus speech emotion recognition using transfer learning and attention-based fusion of wav2vec2 and prosody features," *Knowl.-Based Syst.*, vol. 277, art. no. 110814, 2023.
[5] M. Khan, P. N. Tran, N. T. Pham, A. El Saddik, and A. Othmani, "MemoCMT: Multimodal emotion recognition using cross-modal transformer-based feature fusion," *Sci. Rep.*, vol. 15, no. 1, art. no. 5473, 2025.

Potato Plant Disease Recognition using Lightweight Deep Neural Networks

Anjalika Ekanayaka[1,*], Dimuthu Rajamuni[1], Jananie Jarachanthan[1], Thuseethan Selvarajah[2], Nishankar Sathiyamohan[3]

[1]Faculty of Engineering, University of Jaffna, Sri Lanka,
[2]Deakin University, Australia,
[3] Faculty of Computing, Sabaragamuwa University of Sri Lanka

[*]2021e013@eng.jfn.ac.lk

Potato leaf diseases cause significant yield losses worldwide, making early and accurate detection essential for protecting crop productivity. Deep learning can achieve high accuracy in recognizing diseases within controlled datasets. However, in real farm conditions, uncontrolled factors like lighting, angles, and complex backgrounds pose challenges for deployment. Many deep learning models are also too large for mobile and IoT devices, which limits their real-time application by farmers. This research seeks to balance accuracy with a smaller model size for potato leaf disease identification by using knowledge distillation to transfer information from large teacher models (DenseNet201 and ViT-Base, both achieving over 92% test accuracy) to lightweight student networks. The tested student models included ShuffleNetV2, MobileNetV2, EfficientNetB0, and SqueezeNetL1. EfficientNetB0 achieved the highest student accuracy of 91.26%, offering a strong balance between performance and compactness. The model size was cut by 50% through quantization, changing from 32-bit to 16-bit floating-point, while maintaining performance suitable for edge devices. The study used the Potato Leaf Disease Dataset in an uncontrolled environment, which included 3,076 images across seven classes: bacteria, fungi, healthy leaves, nematodes, pests, phytophthora, and viruses. Grad-CAM visualizations showed that the student models picked up the important features needed for disease classification.

Keywords: Knowledge Distillation, Lightweight CNN, Quantization, Potato Leaf Disease, Edge Devices

Potato Plant Disease Recognition using Lightweight Deep Neural Networks

A. J. K. Ekanayaka[1], R. D. D. S. Rajamuni[1], J. Jarachanthan[1], S. Thuseethan[2], S. Nishankar[3]
Faculty of Engineering University of Jaffna[1], Deakin University, Australia[2], Faculty of Computing University of Sabaragamuwa[3]

INTRODUCTION

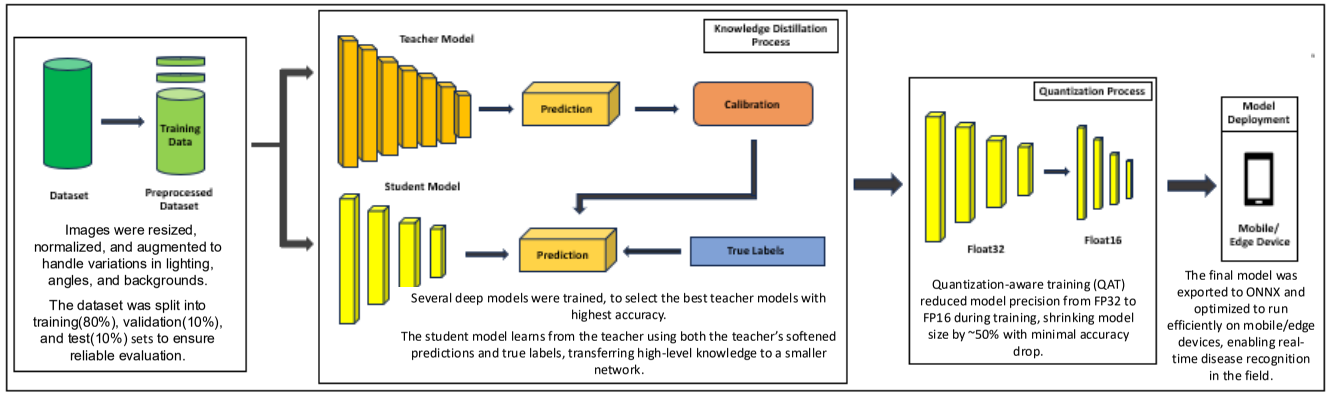
Potato crops are highly susceptible to leaf diseases that significantly impact global yield. Early disease detection is essential, but most deep learning models achieving high accuracy are computationally heavy and unsuitable for real-time field deployment.

This research focuses on developing **lightweight, high-accuracy models** suitable for mobile and edge devices by applying **knowledge distillation** from high-capacity teacher models to efficient student models. The study uses the **Potato Leaf Disease Dataset in Uncontrolled Environments** (3,076 images across 7 classes), reflecting real farm variability such as lighting changes, angles, and background clutter.

OBJECTIVES

Enhance the CNN model with improved accuracy and reduced model size to help cultivators with early disease detection

METHODOLOGY



RESULTS & DISCUSSION

- **Teacher Models:** Multiple models were evaluated, with **ViT-Base achieving 92.23% and DenseNet201 achieving 92.56% test accuracy**, making them the strongest candidates for knowledge distillation.
- **Student Model :** Multiple models were trained under both teacher models and EfficientNetB0 achieved **91.26% accuracy when distilled from DenseNet201 and 90.29% under ViT-Base**, providing the best performance among lightweight networks while significantly reducing computational requirements.
- **Quantization:** Converting EfficientNetB0 from FP32 to FP16 reduced model size from **16.34 MB to 7.87 MB (~50% reduction)** with minimal accuracy drop (**90.29% → 89.64%**), improving deployability on mobile and edge devices.
- **Training for Real-World Robustness:** To handle uncontrolled farm conditions (lighting, angle and background changes), the models were trained with strong data augmentation, **MixUp**, and **CBAM attention added to the CNN teacher model**, improving feature focus and generalization.
- **Feature Validation:** Grad-CAM visualizations confirmed that the model correctly highlighted key infection regions such as spots, lesions, and discoloration, indicating effective knowledge transfer from teacher to student models.

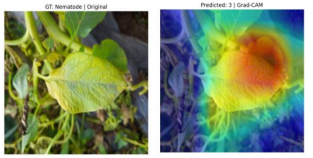


Figure1: Grad-CAM results of student model

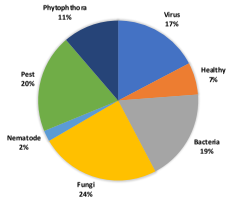


Figure2: Dataset Class Distribution

CONCLUSION

Knowledge distillation improved the accuracy of lightweight models by transferring rich representations from high-capacity teachers. Strong augmentation and MixUp helped counter dataset variability in uncontrolled farm conditions. CBAM enhanced feature focus in the teacher network, improving student learning. FP16 quantization halved model size with minimal accuracy loss, enabling edge deployment. Grad-CAM confirmed that both teacher and student models attended to correct diseased regions, demonstrating reliable and interpretable predictions.

REFERENCE

1. C.-Y. Chang and C.-C. Lai, "Potato Leaf Disease Detection Based on a Lightweight Deep Learning Model," Machine Learning and Knowledge Extraction, vol. 6, no. 4, pp. 114, Oct. 2024, doi: 10.3390/make6040114
2. X. Zhang, Q. Feng, D. Zhu, X. Liang, and J. Zhang, "Compressing recognition network of cotton disease with spot-adaptive knowledge distillation," Frontiers in Plant Science, vol. 15, p. 1433543, Sep. 2024. doi: 10.3389/fpls.2024.1433543.

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P20

Recognition of Electrical Components in Hand-Drawn Schematics

S. M. Semasingha Mudiyanselage*, Samantha Tharani Jeyakumar, Kesavi Nadesh

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*smschathu1234@gmail.com

Hand-drawn circuit schematics are widely used in engineering and education, remain difficult to recognize automatically due to inconsistent drawing styles, symbol overlap, incomplete strokes, or class imbalance in real-world sketches. Although the existing CNN and ANN-based classifiers and graph-based recognizers might reach good results, they mostly suffer from a lack of scalability, high computational intensity, and failure in real-time multi-component detection. This paper seeks to overcome these limitations by evaluating YOLOv8 and YOLO11 detectors for fast and accurate recognition of hand-drawn electrical components. There are 13 essential classes of components selected from the Handwritten Schematics – CGHD dataset and annotated in YOLO format, with further division into an 80% training set, a 10% validation set, and a 10% test set. Trained with Google Colab equipped with CUDA acceleration and evaluated using precision, recall, f1-score, and mean average precision at IoU thresholds 0.5 and 0.5 – 0.95, the YOLOv8s model achieved the highest overall accuracy with the highest f1-score of 0.987 and recall of 0.987 at 50 epochs, while YOLO11s reached the best balance of speed and performance, achieving the following results: F1-score of 0.977 and mAP@50 of 0.992. Comparing YOLOv8s with Faster R-CNN with a ResNet-50 backbone, the former has a reduction in trainable parameters of approximately 55%, allowing real time detection. Our experiments show the suitability of lightweight YOLO architectures for hand-drawn schematic analysis, opening up new ways toward scalable and efficient solutions. The next step would involve extension into broader sets of electrical and logical symbols.

Keywords: Hand-drawn schematics, YOLOv8, YOLO11, Object detection, Electrical components

Recognition of Electrical Components in Hand-Drawn Schematics

S.M.S. Chathuranga, J. Samantha Tharani, N. Kesavi
Department of Computer Science, University of Jaffna
smschathu1234@gmail.com

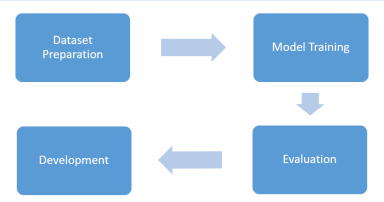
Introduction

- Hand-drawn schematics are common in education and engineering.
- Automatic recognition is difficult due to messy strokes, overlaps, and handwriting variations.
- Existing CNN/ANN/graph-based methods work only for clean, isolated symbols and are not real-time.
- Previous YOLO-based research is limited to small symbol sets or synthetic data.
- No prior work evaluates YOLOv8/YOLO11 on real hand-drawn electrical components.
- **Gap**: need for scalable, accurate, real-time detection models.

Objectives

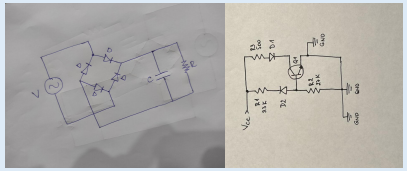
- To evaluate YOLOv8 and YOLO11 for multi-component detection in hand-drawn schematics.
- To compare performance with Faster R-CNN using standardized metrics.
- To analyze model accuracy, speed, and trainable parameter efficiency.
- build web-based interface for component detection.

Methodology

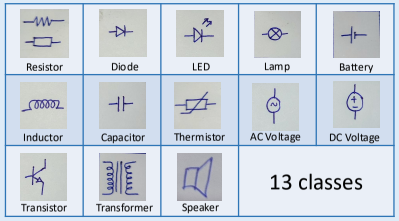


1. Dataset preparation → cleaning & annotating images
2. Model training → YOLOv8/YOLO11 + Faster R-CNN
3. Evaluation → metrics + comparison
4. Deployment → web system

Dataset

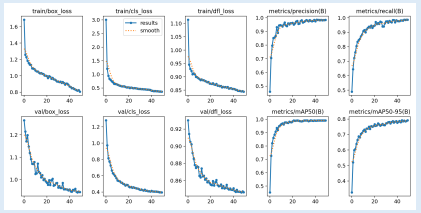


- **Source**: CGHD (Handwritten Schematics) Dataset
- Focused on 13 electrical component classes.

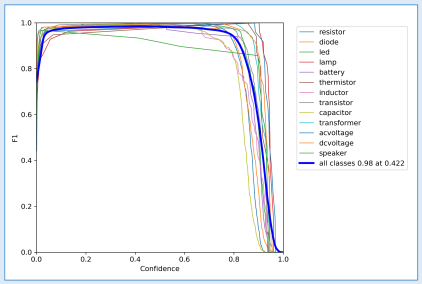


Results

Model	Epochs	Precision	Recall	F1-score	mAP@50	mAP@50-95	Time(h)
YOLO11s	20	0.968	0.946	0.957	0.980	0.772	0.577
	30	**0.988**	0.967	0.977	**0.992**	0.783	0.875
	50	0.986	0.977	0.981	0.990	0.793	1.468
YOLOv8s	20	0.957	0.955	0.958	0.983	0.765	0.665
	30	0.983	0.960	0.971	0.989	0.784	0.884
	50	0.984	0.987	**0.987**	0.991	0.791	1.455
YOLO11m	20	0.948	0.978	0.963	0.988	0.784	0.800
	30	0.984	0.780	0.871	0.988	0.793	0.918
	50	0.979	0.967	0.974	0.990	**0.811**	1.336
YOLOv8m	20	0.979	0.970	0.975	0.990	0.791	0.559
	30	0.987	0.982	0.986	0.991	0.805	0.928
	50	0.912	**0.990**	0.950	0.991	0.791	-



Training loss, Validation loss, Metrics Curves



F1-Confidence Curve

Conclusion

- Evaluated YOLOv8 and YOLO11 for hand-drawn electrical component detection. Both models achieved high accuracy and stable training convergence.
- YOLOv8 achieved the highest F1-score, balancing precision and recall.
- YOLO models handle messy strokes, shape variations, and overlapping symbols better than Faster R-CNN.
- Real-time detection achieved, enabling practical applications.
- Developed a web system for instant recognition of hand-drawn components.
- **Future work**: Expand to all 58 component classes, include logical symbols, and integrate circuit-level understanding.

References

[1] Thoma, F., Bayer, J., Li, Y., & Dengel, A. (2021, September). A public ground-truth dataset for handwritten circuit diagram images. In International Conference on Document Analysis and Recognition (pp. 20-27). Cham: Springer International Publishing.

[2] Günay, M., Köseoğlu, M., & Yıldırım, Ö. (2020). Classification of hand-drawn basic circuit components using convolutional neural networks. In 2020 International Congress on Human Computer Interaction, Optimization and Robotic Applications (HORA) (pp. 1-5). IEEE.

[3] Rachala, R. R., & Panicker, M. R. (2022). Hand-drawn electrical circuit recognition using object detection and node recognition. SN Computer Science, 3(3), 244.

[4] Rabbani, M., Khoshkangini, R., Nagendraswamy, H. S., & Conti, M. (2016). Hand drawn optical circuit recognition. Procedia Computer Science, 84, 41-48.

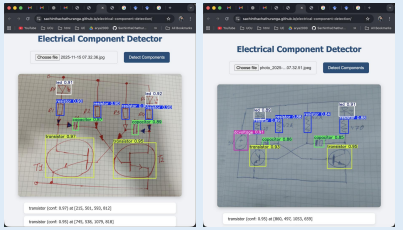
[5] Aksakallı, M., Gümüştekin, Ş. (2018). Hand-drawn Circuit Recognition. Circuit and System Journal

[6] Liu, S., Lee, L., Jiu, F., & Feng, G. (2020). Graph-Based Locality-Sensitive Circuit Sketch Recognizer. IEEE Access, 8, 204183-204193.

[7] Amraee, S., Chinipardaz, M., Charoosaei, M., & Mirzaei, M. A. (2022). Handwritten logic circuits analysis using the Yolo network and a new boundary tracking algorithm. IEEe Access, 10, 76095-76104.

Application

Integrated the best-performing YOLO model into a practical web-based application.



Segmentation of Fruit Surface Defects using Semi Supervised Learning

Saarukesan Premkumar*, Siyamalan Manivannan

Department of Computer Science, University of Jaffna, Sri Lanka

*saaru27kesan@gmail.com

Automated detection and segmentation of fruit-surface defects are essential for improving product quality, reducing dependence on manual inspection, and enabling scalable agricultural automation. Fruit defects vary widely in shape, texture, size, and severity, limiting the generalization of deep learning models. This study evaluates semantic segmentation of orange surface defects using a U-Net encoder–decoder with a ResNet34 backbone, trained on the Orange Navel dataset comprising 1,448 images annotated for three defect categories: sunburn, ulcer, and wind scarring. The model achieves strong quantitative performance: mean Pixel Accuracy (mPA) of 90.57%, mean Intersection over Union (mIoU) of 84.68%, and Dice coefficient of 91.37%. Class-wise IoU highlights near-perfect background segmentation (99.13%) and robust performance on sunburn (85.67%), while ulcer (83.61%) and wind scarring (70.29%) show reduced fidelity-likely due to subtler texture cues, lower contrast, and more irregular lesion boundaries. These results position U-Net with a ResNet34 encoder as a reliable baseline for fruit-defect segmentation. To strengthen generalization under limited annotations, a semi-supervised learning strategy based on the FixMatch paradigm is employed. In this approach, pseudo labels are generated for unlabeled samples, and only high-confidence predictions are used as supervisory signals. Training is guided by enforcing consistency between the model's predictions and these pseudo labels, allowing effective use of both labeled and unlabeled data and improving segmentation performance across all defect categories.

Keywords: Fruit Defect Segmentation, Image Segmentation, U-Net, ResNet34, Semi Supervised Learning

Segmentation of Fruit Surface Defects using Semi Supervised Learning

Saarukesan Premkumar, Siyamalan Manivannan
Department of Computer Science, University of Jaffna, Sri Lanka
saaru27kesan@gmail.com, siyam@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

Introduction

- Fruit quality inspection is important because surface defects reduce market value and consumer acceptance.
- Manual inspection is slow, inconsistent, and costly, creating a need for automated vision-based systems.
- Segmentation is challenging due to high variability in defect appearance, small or unclear defect regions, and visually similar patterns.
- This practical baseline is established by training a U-Net with a ResNet34[1] encoder within a semi-supervised FixMatch[4] framework, effectively leveraging unlabeled data and providing detailed per-class performance metrics.

Experimental Setup

Table 1 : Experimental Setup and Implementation Details

Model Architecture	U-Net encoder–decoder with ResNet34 backbone (initialized with ImageNet-pretrained weights).
Dataset	Orange Navel 1.5k (1,448 images)[2] 3 defect classes: sunburn, ulcer, and wind scarring.
Data Split	60:20:20 Ratio • Training: 868 images • Validation: 290 images • Test: 290 images
Preprocessing	• Optimizer: Adam • Learning Rate: 1e-4 • Batch Size: 32

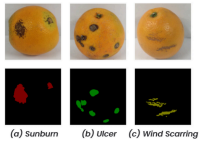


Figure 1 : Typical samples of Orange Navel 1.5k dataset

Methodology

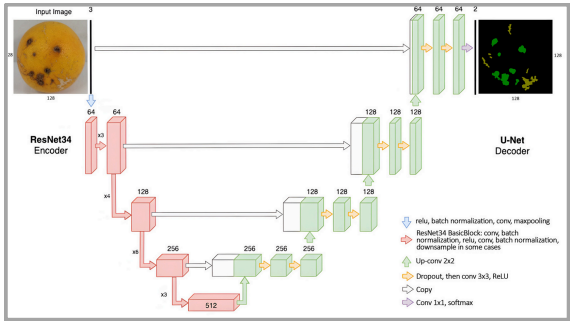


Figure 2 : UNet architecture with ResNet34 encoder

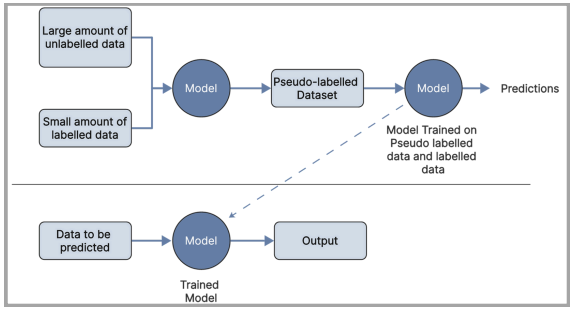


Figure 3 : Semi Supervised Learning

Results and Discussion

Table 2 : Comparison of ResNet34 with other representative models in Navel 1.5k dataset

Model	mPA(%)	mIoU(%)	Dice(%)	IoU(%)		
				Sunburn	Ulcer	Windscarring
ResUNet++[5]	85.0	79.3	74.9	81.3	80.5	56.3
DeepLabv3+[6]	86.7	81.4	86.4	84.1	82.6	59.6
SwinUNet[7]	76.8	69.4	55.0	67.3	70.8	40.4
TransUNet[8]	87.5	81.0	73.5	84.2	81.3	59.2
ResNet34	90.57	84.68	91.37	85.67	83.61	70.29
PriorIFormer[3]	92.4	87.4	86.4	90.0	87.8	72.3

Table 3 : Comparison of Supervised vs. Semi-Supervised performance across different labeling ratios

Label/Unlabeled Split	Type	mIoU(%)	Dice(%)
100% Labeled	Fully Supervised	84.68	91.37
50% Labeled + 50% Unlabeled	Semi Supervised	83.26 (+0.04 improvement)	90.45
50% Labeled	Fully Supervised	83.22	90.45
30% Labeled + 70% Unlabeled	Semi Supervised	80.83 (+0.17 improvement)	88.90 (+0.13 improvement)
30% Labeled	Fully Supervised	80.66	88.77
10% Labeled + 90% Unlabeled	Semi Supervised	75.93 (+1.20 improvement)	85.53% (+0.79 improvement)
10% Labeled	Fully Supervised	74.73	84.74

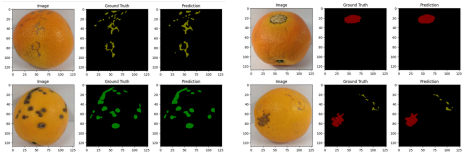


Figure 4 : segmentation results of the proposed approach

Conclusion

- This study introduced an automated citrus defect segmentation framework using a U-Net with a ResNet34 encoder, providing a solid supervised baseline for handling defect variability and limited labeled data.
- Additionally, the use of semi-supervised learning improved generalization by leveraging unlabeled samples and reducing annotation effort, supporting scalable and reliable real-time defect detection for agricultural quality control.

References

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2016, pp. 770–778.
[2] X. Cai, Y. Zhu, S. Liu, Z. Yu, and Y. Xu, "FastSegFormer: A knowledge distillation based method for real-time semantic segmentation of surface defects in navel oranges," Comput. Electron. Agric., vol. 217, p. 108604, 2024.
[3] X. Xu, T. Xu, Z. Wei, Z. Li, Y. Wang, and X. Rao, "Enhancing citrus surface defects detection: A priori feature guided semantic segmentation model," Artif. Intell. Agric., vol. 15, no. 1, pp. 67–78, 2025.
[4] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. A. Raffel, et al., "FixMatch: Simplifying semi-supervised learning with consistency and confidence," in Adv. Neural Inf. Process. Syst., vol. 33, 2020, pp. 596–608.
[5] D. Jha, P. H. Smedsrud, D. Johansen, T. de Lange, H. D. Johansen, P. Halvorsen, and M. A. Riegler, "A comprehensive study on colorectal polyp segmentation with ResUNet++, conditional random field and test-time augmentation," IEEE J. Biomed. Health Inform., vol. 25, pp. 2029–2040, 2021. doi: 10.1109/JBHI.2021.3049304.
[6] L. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in Proc. Eur. Conf. Comput. Vis. (ECCV), 2018, pp. 833–851. doi: 10.1007/978-3-030-01234-2_49.
[7] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, and M. Wang, "Swin-Unet: Unet-like pure transformer for medical image segmentation," in Proc. Eur. Conf. Comput. Vis. (ECCV), 2023, pp. 205–218. doi: 10.1007/978-3-031-25066-8_9.
[8] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, and Y. Zhou, "TransUNet: Transformers make strong encoders for medical image segmentation," arXiv preprint arXiv:2102.04306, 2021.

Self-Supervised Ultrasound Image Denoising Using Noise2Void and Noisier2Noise

Roy Nesarajah Viththagan*, Siyamalan Manivannan, Thanikasalam Kokul

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*viththagan.rn@gmail.com

Speckle noise in ultrasound imaging degrades image quality, leading to reduced diagnostic accuracy and challenges to automated image analysis. Conventional deep-learning methods are hindered by the scarcity of paired clean-noisy ultrasound data. To address this, we propose a self-supervised deep learning framework that integrates Noise2Void (N2V) pretraining and Noisier2Noise (N2N) finetuning for ultrasound image denoising. Unlike conventional supervised methods that require clean ground truth, N2V enables learning directly from single noisy images by masking random pixels and predicting their intensity based on the surrounding context, effectively exploiting spatial redundancy. Building on this, N2N finetuning leverages pairs of noisy images, where one image is synthetically degraded with additive Gaussian noise, allowing the network to learn robust noise mappings without ever accessing clean data. This two-stage strategy addresses the critical limitation of scarce paired datasets in medical imaging. The framework is evaluated against state-of-the-art methods on the publicly available Breast Ultrasound Images Dataset (BUSI) dataset containing 798 images against state-of-the-art methods, and experimental results demonstrate that MedSegUNet achieves an MSE of 0.0003 and a PSNR of 35.38 dB, enabling substantial speckle suppression while preserving fine anatomical details. Although effective on BUSI, performance may vary with device-specific noise characteristics. The model remains lightweight, making it suitable for deployment on small datasets and real-time clinical applications. This work highlights the potential of combining context-based self-supervision and noise-pair learning for practical ultrasound image enhancement.

Keywords: Ultrasound image denoising, Self-supervised learning, Noise2Void, Noisier2Noise, Speckle noise

Self-Supervised Ultrasound Image Denoising Using Noise2Void and Noisier2Noise

Roy Nesarajah Viththagan, Siyamalan Manivannan and Thanikasalam Kokul
Department of Computer Science, Faculty of Science, University of Jaffna
viththagan.rn@gmail.com, siyam@univ.jfn.ac.lk, kokul@univ.jfn.ac.lk

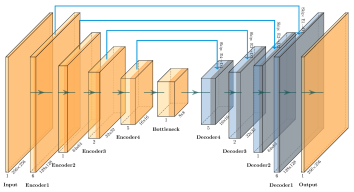
Introduction

Ultrasound imaging is widely used in clinical diagnostics for its real-time, cost-effective, and non-invasive nature. However, speckle noise—caused by coherent wave interference—degrades image quality by reducing contrast and obscuring fine anatomical structures [1, 2], negatively impacting both manual interpretation and automated tasks like segmentation. Traditional denoising methods (e.g., median filtering, wavelets, anisotropic diffusion) often oversmooth diagnostically relevant structures [3]. Deep learning approaches, particularly CNNs and GANs, achieve better noise suppression. GAN-based models like CGAN/WGAN [4], DEUS-Net [5], and Capsule GAN [1] enhance perceptual quality via adversarial learning but require large paired datasets—often unavailable in medical imaging. Self-supervised techniques such as Noise2Void (N2V) [6] and Noisier2Noise (N2N) [7] overcome this by training directly on noisy data. While effective in other modalities, they are underexplored in ultrasound. This study introduces a two-stage self-supervised framework—N2V pretraining followed by N2N finetuning—for ultrasound image denoising. The method is evaluated against state-of-the-art approaches using PSNR and SSIM. demonstrating effective speckle reduction without clean reference data.

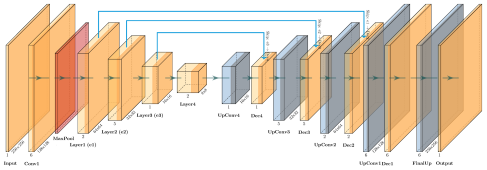
Dataset and Preprocessing

We use the BUSI dataset [8] (798 images: benign, malignant, normal), Images were resized to 256 × 256, and split 70%/30% for training/testing. Augmentations included rotation, flipping, contrast adjustments, and Gaussian noise.

Model Architectures



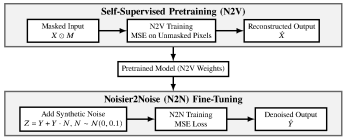
MedSegUNet: A lightweight MONAI UNet (32 base channels, 1.2 M parameters) trained from scratch to avoid overfitting on the modest BUSI dataset.



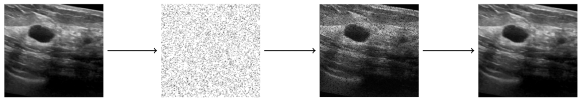
ModifiedResNet: ResNet50 encoder pretrained on ImageNet with UNet-style decoder, adapted for grayscale.

Methodology

N2V Pretraining: Randomly masked 10% of pixels; MSE loss on unmasked pixels
N2N Finetuning: Inputs with additive Gaussian noise (σ = 0.1, matching typical speckle variance observed in BUSI); MSE loss. This pipeline allows effective training without paired clean data.



Pretraining: Masked Image Regeneration (*Self-Supervised*)

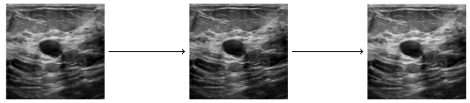


Loss Function:

$$\mathcal{L}_{MSE} = \frac{1}{|M|} \sum_{i \in M} \left(\hat{Y}_i - Y_i\right)^2 = MSE(\hat{Y} \times M,\ Y \times M)$$

Loss is computed only over masked pixels defined by binary mask M, encouraging reconstruction of occluded regions.

Fine-Tuning: Noisier2Noise (*Denoising*)



$$Z = Y + Y \cdot N_{spectral} \quad with \quad N_{spectral} \sim \mathcal{N}(0, \sigma^2_{spectral})$$

Loss Function: The loss function used for fine-tuning is the Mean Squared Error (MSE) between the predicted denoised image $\hat{Y}$ and the pseudo-clean image Y:

$$\mathcal{L}_{MSE} = \frac{1}{|M|} \sum_{i \in M} \left(\hat{Y}_i - Y_i\right)^2 = MSE(\hat{Y},\ Y)$$

Conclusion

We present a self-supervised N2V+N2N framework for ultrasound image denoising that eliminates the need for clean reference data. Experimental results show that MedSegUNet effectively reduces speckle noise while preserving fine anatomical structures, outperforming existing GAN-based and self-supervised methods on the BUSI dataset.

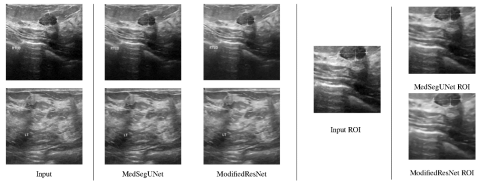
Results and Discussion

Table 1 summarizes denoising performance. MedSegUNet generally outperformed ModifiedResNet, achieving lower MSE and higher PSNR/SSIM, while ModifiedResNet achieved slightly better perceptual quality (BRISQUE). Compared to literature methods, MedSegUNet exceeded WGAN, DEUS-Net, and Capsule GAN. The results highlight that a lightweight UNet-based architecture is suitable for small datasets and effective speckle suppression, while pretrained ResNet features may require extended finetuning.

Table 1: Denoising performance on BUSI dataset (noise std=0.1). Input BRISQUE=39.66. For metrics: lower is better (↓); higher is better (↑).

Approach	MSE↓	PSNR↑	SSIM↑	SNR↑	CNR↑	BRISQUE↓
MedSegUNet (Ours)	**0.0003**	**35.38**	**0.9576**	**22.40**	**22.0685**	30.07
ModifiedResNet (Ours)	0.0005	32.69	0.8910	19.72	16.2057	**28.13**
WGAN [4]	-	33.00	0.92	-	-	-
DEUS-Net [5]	-	21.65	0.631	20.64	-	-
Capsule GAN [1]	-	34.16	0.90	-	-	-

Comparison of Denoising Results



Comparison of denoising results for breast ultrasound images. Left: full images (input, MedSegUNet, ModifiedResNet). Right: zoomed regions of interest (ROI) for detail assessment.

References

[1] A. Sivaanpu, K. Punithakumar, K. Thanikasalam, M. Noga, R. Zheng, D. Ta, E. H. Lou, and L. H. Le, "A lightweight ultrasound image denoiser using parallel attention modules and capsule generative adversarial network," Informatics in Medicine Unlocked, vol. 50, p. 101569, 2024.
[2] S. V. M. Sagheer and S. N. George, "A review on medical image denoising algorithms," Biomedical signal processing and control, vol. 61, p. 102036, 2020.
[3] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," IEEE transactions on image processing, vol. 26, no. 7, pp. 3142–3155, 2017.
[4] Y. Jimenez-Gaona, M. J. Rodr ´ iguez-Alvarez, L. Escudero, C. Sandoval, and V. Lakshminarayanan, "Ultrasound breast images denoising using generative adversarial networks (gans)," Intelligent Data Analysis, vol. 28, no. 6, pp. 1661–1678, 2024.
[5] C. Janorschke, J. Meyer, and D. Wulff, "Self-supervised denoising ultrasound network (deus-net)," IFAC-PapersOnLine, vol. 58, no. 24, pp. 614-619, 2024.
[6] A. Krull, T.-O. Buchholz, and F. Jug, "Noise2void-learning denoising from single noisy images," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 2129–2137, 2019.
[7] N. Moran, D. Schmidt, Y. Zhong, and P. Coady, "Noisier2noise: Learning to denoise from unpaired noisy data," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 12064–12072, 2020.
[8] W. Al-Dhabyani, M. Gomaa, H. Khaled, and A. Fahmy, "Dataset of breast ultrasound images," Data in brief, vol. 28, p. 104863, 2020

Transformer based Semi-Supervised Learning for Skin Lesion Classification

Jathursika Velummayilum*, Siyamalan Manivannan

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*jathukirubajathu@gmail.com

Skin cancer is one of the most prevalent cancers worldwide, and early detection is crucial for successful treatment. Deep learning has shown promise in automating skin lesion classification; however, most methods require large volumes of labeled data, which are scarce and costly to obtain in the medical domain. To address this limitation, we propose a Transformer-based Semi-Supervised Learning (SSL) framework for skin lesion classification. Our approach employs the Vision Transformer (ViT) as the backbone model, leveraging its strength in capturing long-range dependencies and global contextual information from dermoscopic images. By integrating SSL, the framework effectively utilizes both labeled and unlabeled data, thereby reducing reliance on extensive manual annotations. We evaluate the proposed method on the publicly available ISIC-2018 dataset. Experimental results show that our SSL framework achieves substantial performance improvements over a fully supervised baseline. Specifically, it achieves an accuracy of 74.3%, balanced accuracy of 62.9%, and F1 score of 62.7representing relative improvements of 9.3%, 11.7%, and 19.9%, respectively, compared to the baseline (65.0% accuracy, 56.3% balanced accuracy, and 52.3% F1 score). These findings demonstrate that SSL can significantly enhance skin lesion classification when labeled data are limited. The proposed Transformer-based SSL framework offers a promising direction for developing cost-effective, reliable, and scalable tools to support early skin cancer detection and improve clinical decision-making.

Keywords: Semi-supervised learning, Vision Transformers, Skin Lesion Classification

Transformer based Semi-Supervised Learning for Skin Lesion Classification

Jathursika V. , Siyamalan M.
Department of Computer Science, Faculty of Science , University of Jaffna
jathukirubajathu@gmail.com, siyam@univ.jfn.ac.lk

SLSWCS Student Workshop on Computer Science

Abstract

Skin cancer is one of the most prevalent cancers worldwide, and early detection is crucial for successful treatment. Deep learning has shown promise in automating skin lesion classification; however, most methods require large volumes of labeled data, which are scarce and costly to obtain in the medical domain. To address this limitation, we propose a Transformer-based Semi-Supervised Learning (SSL) framework for skin lesion classification. Our approach employs the Vision Transformer (ViT) as the backbone model, leveraging its strength in capturing long-range dependencies and global contextual information from dermoscopic images. By integrating SSL, the framework effectively utilizes both labeled and unlabeled data, thereby reducing reliance on extensive manual annotations. We evaluate the proposed method on the publicly available ISIC-2018 dataset.

Experiment

- Load dataset + augmentations.
- Create pretrained ViT-B/16.
- Split training into 20% labeled, 80% unlabeled.
- Stage 1: Train only classifier head.
- Stage 2: Unfreeze ViT + supervised fine-tuning.
- Evaluate supervised performance.
- Stage 3: Run FixMatch (weak/strong augmentation + pseudo labels).
- Final test evaluation with multiple metrics.

Loss Function

$$L_{\text{total}} = L_s + \lambda L_u$$

weight for unlabeled loss

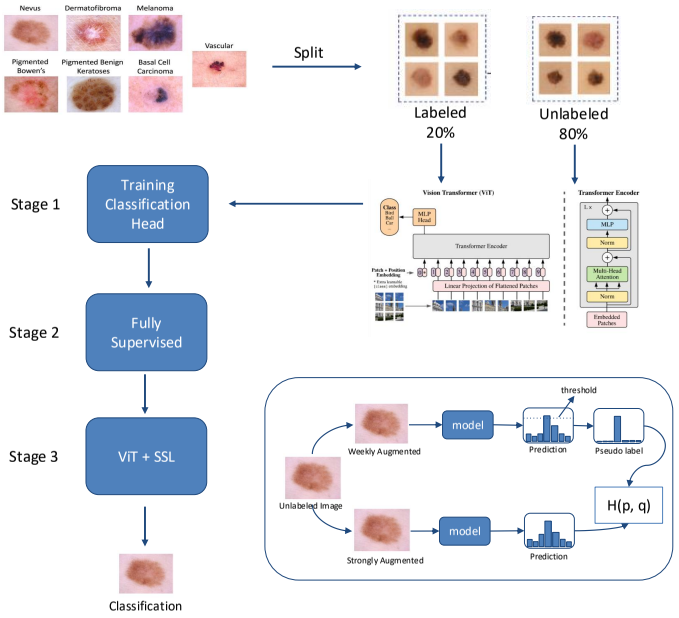
Supervised loss

Unsupervised loss

$$L_s = \frac{1}{B} \sum y_i \, p_\theta(\alpha(x_i)))$$

$$L_u = \frac{1}{\mu B} \sum_{i=1}^{\mu B} 1(\max(p_\theta(A_w(u_i)) \; H(\hat{y}_i, p_\theta(A_s(u_i))$$

Methodology



Dataset

The ISIC 2018 dataset was published by the International Skin Imaging Collaboration (ISIC) as a large-scale dataset of dermoscopy images. It includes 10,015 images.

Class	Count
Melanoma (MEL)	1,117
Nevus (NV)	6,705
Basal Cell Carcinoma (BCC)	514
Actinic Keratosis (AKIEC)	327
Benign Keratosis (BKL)	1099
Dermatofibroma (DF)	115
Vascular Lesion (VASC)	142

To simulate limited annotation scenarios, only 20% of the training set is designated as labeled data, while the remaining 80% is treated as unlabeled for semi-supervised learning.

Conclusion

This study presented a Transformer-based SSL framework for skin lesion classification, addressing the challenge of limited labeled data in medical imaging. By employing the ViT backbone and integrating a FixMatch-inspired strategy, the proposed approach effectively leveraged both labeled and unlabeled samples to enhance learning and improve classification performance. Experiments on the ISIC-2018 dataset demonstrated that our framework significantly outperforms a fully supervised baseline trained with the same proportion of labeled data, yielding notable improvements in accuracy, balanced accuracy, and F1 score. Although the performance remains below that of models trained with fully annotated datasets, the results highlight the potential of SSL to reduce reliance on costly manual labeling.

Results and Discussion

We compare our SSL framework against fully supervised ViT baselines trained with 100% and 20% labeled data as shown in the following table.

Model	Labeled	Unlabeled	Accuracy	Balanced Accuracy	F1 score
Fully Supervised	100%	-	83.5%	74.3%	75.3%
Fully Supervised (ViT)	20%	-	68.92%	57.55%	60.64%
Proposed ViT-SSL	**20%**	**80%**	**74.3%**	**62.9%**	**62.7%**
Ensemble of CNN [1]	**20%**	**80%**	**73.89%**	**72.05%**	**95.45%**

The results highlight that our SSL approach recovers much of the performance lost when only limited labeled data are available. Compared to the 20% supervised baseline, our method yields gains of 9.3% in accuracy, 6.6% in balanced accuracy, and 10.4% in F1 score.

References

1. Manivannan S. Semi-Supervised Learning with Online Knowledge Distillation for Skin Lesion Classification. arXiv preprint arXiv:2508.11511. 2025.
2. Van Engelen, Jesper E., and Holger H. Hoos., A survey on semi-supervised learning, Machine learning, vol. 109, no. 2, pp. 373-440, 2020.
3. Sohn, Kihyuk, et al., Fixmatch: Simplifying semi-supervised learning with consistency and confidence, Advances in neural information processing systems vol. 33: pp. 596-608, 2020.
4. Dosovitskiy, Alexey, et al., Mostafa Dehghani et al., An image is worth 16x16 words: Transformers for image recognition at scale., arXiv preprint arXiv:2010.11929, 2020.
5. Codella, Noel, et al., Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration, arXiv preprint arXiv:1902.03368, 2019.

Transformer Tracking with Controlled Attention for Background Suppression

Janani Kugarajeevan[1,2,*], Thanikasalam Kokul[3], Amirthalingam Ramanan[3], Subha Fernando[2]

[1]Department of Computational Mathematics, Faculty of Technology, University of Moratuwa, Sri Lanka
[2]Department of Interdisciplinary Studies, Faculty of Technology, University of Jaffna, Sri Lanka
[3]Department of Computer Science, University of Jaffna, Sri Lanka

[*]jananitha@univ.jfn.ac.lk

Visual object tracking aims to determine the state of a target object across video frames, starting from its initial position in the first frame. Recently, one-stream Transformer-based trackers have demonstrated impressive performance by jointly processing template and search region tokens through attention mechanisms, allowing effective extraction of target-specific features. However, the high proportion of background in the search region can weaken target representations and reduce the tracker's ability to accurately discriminate target from the background. To address this limitation, we propose a Transformer-based tracking framework with a controlled attention mechanism that suppresses irrelevant attention from background search tokens and enhances target-focused feature learning. In early encoder layers of proposed tracker, attention from search tokens to template tokens is blocked to prevent background interference, while template-to-search attention is allowed for target-specific feature extraction. Following the early encoder layers, a learnable module is introduced to estimate the probability of each search token belonging to the target, using a multilayer perceptron. Based on the module's predictions, search tokens are divided into target and non-target groups, and attention from selected target tokens to the template tokens is enabled to enhance target-focused feature learning. Experiments on GOT-10k dataset demonstrate that our tracker outperforms recent state-of-the-art methods, achieving an average overlap of 73.8%, a success rate of 84.1% at a threshold of 0.5, and a success rate of 70.1% at a threshold of 0.75, while maintaining real-time processing of 39 frames per second.

Keywords: Transformer Tracking, Background Interference, Controlled Attention, Visual Object Tracking

Transformer Tracking with Controlled Attention for Background Suppression

Janani Kugarajeevan[1,2], Kokul Thanikasalam[3], Amirthalingam Ramanan[3], and Subha Fernando[1]

[1]Department of Computational Mathematics, University of Moratuwa, [2]Department of Interdisciplinary Studies, University of Jaffna,
[3]Department of Computer Science, University of Jaffna

Introduction

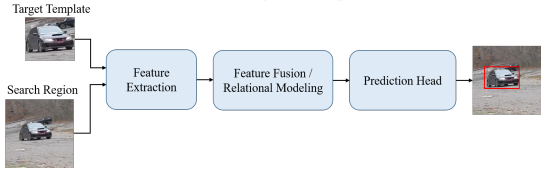
Visual Object Tracking (VOT) is the process of estimating a target object's state in a video sequence, beginning from its initial location in the first frame.



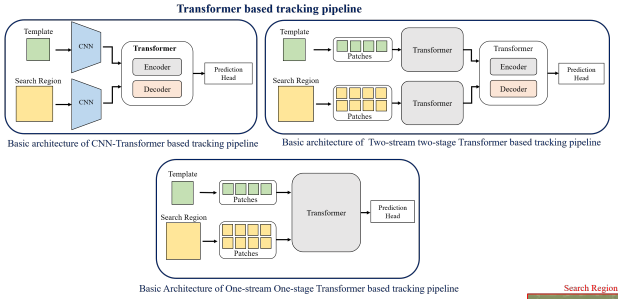
However, it remains far from human-level robustness due to appearance changes, occlusions, and background clutter.

Background

- Most deep learning-based trackers follow a pipeline of feature extraction, feature fusion or relational modeling, and a prediction head.



- Transformer-based tracking frameworks [1] have achieved superior performance in VOT due to their capability for global feature learning.

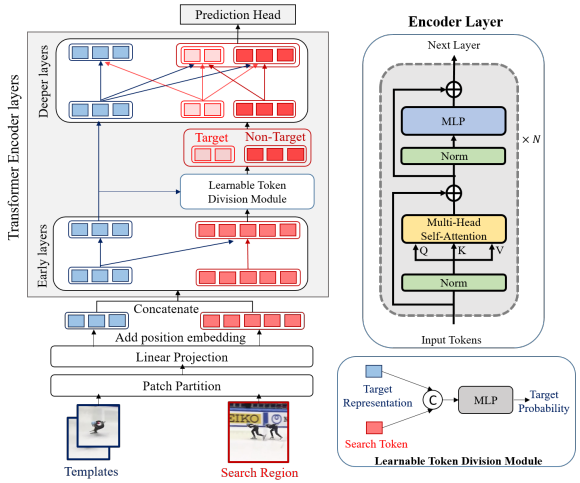


- Limitation of existing trackers: The large proportion of background within search region disturbs the actual target specific features and weakens the tracker's discrimination ability.



- We propose a one-stream Transformer tracking framework that controls the attention flow from background search tokens to template tokens to reduce background interference.

Methodology



- Templates and search regions are split into patches, projected into embeddings, and combined with positional encodings to generate initial tokens.
- Templates and search tokens are concatenated and processed together in encoder layers, enabling cross- and self-relations.
- In early layers, search-to-template attention is blocked to avoid background interference.
- A learnable module predicts each search token's target probability, dividing them into target and non-target groups, and attention from target search tokens to template tokens is enabled in later layers.
- Search tokens from the final encoder layers are reshaped into a 2D feature map and fed to the prediction head to locate the target.

Experimental Setup

Architecture

- **Backbone:** ViT-B encoder [2]
- **Template region size:** 128×128
- **Search region size:** 256×256
- **Patch size:** 16×16
- **#Template frames:** 2

Dataset: GOT-10k [3]

- 10k training and 180 testing videos.
- Testing ground truth is not publicly available.
- Zero class overlap between training and testing sets.

Test Results

Comparisons with SOTA trackers evaluated on GOT-10k.

Tracker	AO	$SR_{0.5}$	$SR_{0.75}$
OSTrack-256 [1]	71.0	80.4	68.2
TATrack-B [4]	73.0	83.3	68.5
ASAFormer [5]	71.3	80.0	67.6
TIT [6]	71.6	82.5	67.4
CTIFTrack [7]	71.3	81.5	66.7
DYTrack [8]	71.4	80.2	68.5
Our Tracker	**73.8**	**84.1**	**70.1**

Ablation Study

Attention from search to template.

Attention	AO	$SR_{0.5}$	$SR_{0.75}$
Entirely blocked	72.7	82.8	69.8
Ours	**73.8**	**84.1**	**70.1**

Ablation on search token division.

Based on	AO	$SR_{0.5}$	$SR_{0.75}$
Attention score	73.2	83.5	69.8
Learnable score	**73.8**	**84.1**	**70.1**

Conclusion

We introduced a one-stream Transformer tracker with controlled attention that suppresses background interference and enhances target-focused learning, achieving SOTA performance at 39 FPS.

References

[1] B. Ye, et al., Joint feature learning and relation modeling for tracking: A one-stream framework, ECCV, 2022.
[2] A. Vaswani, et al., Attention is all you need, NIPS, 2017.
[3] L. Huang, et al., GoT-10k: A large high-diversity benchmark for generic object tracking in the wild, TPAMI, 2019.
[4] K. He, et al., Target-aware tracking with longterm context attention, AAAI, 2023.
[5] X. Gong, et al., ASAFormer: Visual tracking with convolutional vision transformer and asymmetric selective attention, KBS, 2024.
[6] J. Yao, et al., Tracking in tracking: An efficient method to solve tracking distortion, EAAI, 2024.
[7] Z. Zhang, et al., CTIFTrack: Continuous temporal information fusion for object track, ESA, 2025.
[8] J. Zhu, et al., Exploring dynamic transformer for efficient object tracking, IEEE TNNLS, 2025.

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P25

Computer Networks and Mobile Computing

Bridging the GPS Divide: A Dual-Mode Real-Time Tracking App for Highway Buses in Sri Lanka

Tharindu Rajasinghe, S. A. A. H. Samaraweera*

Department of Computer Science, Faculty of Natural Sciences, The Open University of Sri Lanka

*saasa@ou.ac.lk

The efficiency of the public transport is a burning issue in Sri Lanka, especially in the case of highway buses which are not regularly scheduled and not seen in real-time. This project tries to address these challenges by introducing GoBus SL, a real-time bus tracking mobile application. The most notable aspect of this system is the dual-mode tracking ability. The buses fitted with GPS devices can send real-time information as well as the app enables conductors or drivers to share live locations via their smartphones which is a powerful backup in case of GPS non-availability. The literature review has shown that there is a significant research gap in this field because the existing solutions are based solely on built-in GPS and do not provide opportunities to update manually. To fill this gap, the proposed system has incorporated GPS tracking, validation that is controlled by the administrator and role-based access to drivers, conductors and passengers. The app was created with Flutter on the frontend, Spring Boot on the backend, PostgreSQL to handle data and Google Maps API for the location services. The included features are route search, estimated time of arrival and administrative approval processes. A controlled evaluation was conducted to examine the continuity of tracking and responsiveness of the dual-mode system. Simulating GPS failures, network fluctuations and fallback transitions were considered for the testing. The results showed that the manual fallback mechanism maintained uninterrupted location availability. Therefore, this application is a viable way of improving public transport in the developing world since it provides an effective alternative to GPS-based tracking.

Keywords: Real-time tracking, Mobile App, GPS Alternative, Flutter, Spring Boot

Bridging the GPS Divide: A Dual-Mode Real-Time Tracking App for Highway Buses in Sri Lanka

Rajasinghe T., and Samaraweera SAAH
Department of Computer Science, Faculty of Natural Sciences, The Open University of Sri Lanka
tharinduarunajith@gmail.com, *saasa@ou.ac.lk*

Introduction

- Efficient public transportation is a key aspect in the economic development in a country.
- In Sri Lanka, commuters face unreliable and uncoordinated bus services especially in highway bus routes.
- Existing global real-time bus tracking systems either rely entirely on GPS-equipped vehicles or are designed for different geographic and technological contexts [1]-[2].
- These systems fail when GPS devices malfunction or are not installed.
- Therefore, in this study the **GoBus SL** mobile application was developed to address these limitations.
- This app was designed for Sri Lankan highway buses that enables passengers to view real-time bus locations, estimated arrival times (ETA) and schedule details.
- It introduces a dual-mode tracking method that integrates both GPS devices and manual location updates by bus staff via their smartphones.
- Hence, it offers a more adaptable solution to real-world infrastructure constraints in Sri Lanka.

Objectives

- The aim of this project is to develop a mobile application that provides real-time bus tracking and timetable information, integrating GPS data and user updates to deliver accurate and reliable information to highway commuters.

- To achieve the above aim following three objectives were formulated.
  1) To Develop a real-time bus tracking app for Sri Lankan highway buses.
  2) To provide ETA using Distance Matrix API.
  3) To enable smartphone-based manual location sharing for non-GPS buses.

Methodology

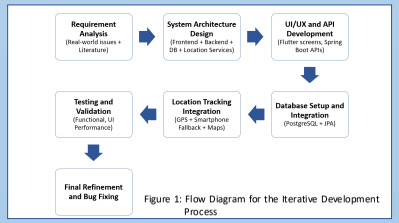


Figure 1: Flow Diagram for the Iterative Development Process

System Architecture

Component	Description
Database	• A PostgreSQL database stores all bus, staff, rout and location data.
Backend Server (Spring Boot API)	• Acts as the main controller of the system. • Receives location updates from buses. • Sends real-time location and ETA to users. • Handles admin requests and manages data.
Bus Location Sources	• **Bus with GPS device** Sends location directly to the backend. • **Bus without GPS** Driver/conductor uses their smartphone GPS to share location.
Admin Panel	• Admin adds buses, approves staff, updates schedules and manages system data. • All actions are sent to the backend server
User App	• Passengers request the current location of a bus. • The backend returns the latest live position. • The app displays the bus on Google Maps and shows ETA using Distance Matrix API.
Maps & ETA Services	• Google Maps API is used to show the bus on the map. • Google Distance Matrix API calculates real-time arrival predictions.

Solution

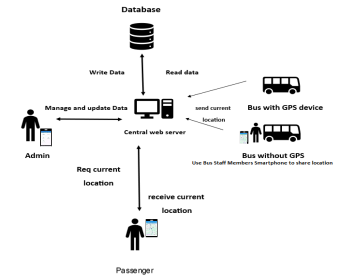


Figure 2: High-Level Architecture of the Proposed Bus Tracking System

Results

- The dual-mode model was tested under control conditions and demonstrated consistent accuracy in real-time location sharing.
- When GPS updates were intentionally interrupted, the manual fallback mode updated the location data to maintain a smooth flow without noticeable gaps.
- Unlike existing systems that depend solely on GPS hardware [2]-[4], GoBus SL's fallback mechanism via smartphone GPS ensures uninterrupted tracking even when GPS devices fail, a common issue in low-budget deployments.
- Responsiveness of the system showed that the system continuously handled location updates within short periods needed for real-time use.
- Compared to systems like College Bus Tracker [1], BUSNAV [4] and Smart Bus Tracking [5], GoBus SL offers improved adaptability to infrastructure limitations, supports admin control, and introduces a more practical passenger experience.
- The testing phase confirmed that passengers could reliably locate buses and plan their travel efficiently.

Conclusion

- The **GoBus SL** mobile application delivers a working and scalable real-time bus tracking solution for Sri Lankan highway transport.
- It fills notable research gaps by introducing a dual-mode GPS/manual tracking approach that improves service reliability in local conditions even under GPS hardware failure or signal loss.
- The controlled testing confirms the conceptual and technical viability of this mobile system.

Future Works

- Extension to normal route buses
- Excluding buses that have already passed
- Improve ETA calculation based on the specific bus route instead of general straight-line distances
- Advanced search features
- iOS Compatibility
- Security enhancements
- Improve user expérience (UX)

References

[1] Vijayalakshmi, N. Dhikshitha, S. Reshma, S. AkshayaPriya and T. Poushika, "College Bus Tracking System," International Journal for Research in Applied Science & Engineering Technology (IJRASET), vol. 11, no. VI, pp. 119-123, 2023.
[2] N. P. Patil, S. S. Divate, P. S. Patil, R. S. Kharche and L. H. Telgar, "ANDROID APP FOR BUS TRACKING USING GPS," International Research Journal of Modernization in Engineering Technology and Science, vol. 05, no. 02, pp. 1275-1277, 2023.
[3] A. S. Rao, P. Reena, K. Amaravathi, P. Akhila and K. Aravind, "GPS Based Real-Time Bus Positioning System," International Journal For Advanced Reasearch In Science & Technology, vol. 13, no. 12, pp. 133-140, 2023.
[4] K. Akhade, K. Dadge, R. Nikam, D. Satpute and K. Shinde, "BUSNAV: A Real Time Bus Tracking Application," International Journal of Innovative Science and Research Technology, vol. 9, no. 4, pp. 1932-1935, 2024.
[5] S. Redij, A. Rao, S. A. Kumar and A. Jadhav, "Smart Bus Tracking For Blind And Visually Impaired," ITM Web of Conferences, p. 44:02001, 2022.
[6] D. Singh, V. Santhanam, T. Punde, Srianjana and K. Bade, "Bus Tracker and Route Recommender using Google APIs," International Journal of Innovative Science and Research Technology, vol. 8, no. 5, pp. 2667-3672, 2023.

Decoding Phrasal Verbs with Context-Aware BERT Representations

Sujair Ibrahim[1], Abishethvarman Vadivel[2,*], Kumara T. G. S. Banage[1]

[1]Department of Software Engineering, Sabaragamuwa University of Sri Lanka,
[2]Department of Computing and Information Systems, Sabaragamuwa University of Sri Lanka

[*]abishethvarman@gmail.com

Phrasal verbs are an important component of English, with meanings that frequently go beyond the literal readings of their constituent words. Their context-dependent and polysemous nature poses significant challenges to natural language processing systems, making correct disambiguation critical for applications such as machine translation, sentiment analysis, and information retrieval. Traditional approaches, including rule-based and statistical methods, often fail to capture the varied meanings of phrasal verbs, especially in diverse or complex contexts. To address these challenges, we propose a methodology based on contextual BERT representations. This work is innovative in applying in-context editing directly to BERT's contextual embeddings in a supervised manner, allowing the model to refine its internal representations of phrasal verbs based on subtle contextual cues. Unlike previous transformer-based approaches, this strategy operates with lightweight and low computational requirements, making it practical for broader research settings and accessible with limited hardware. By guiding the model to adapt its embeddings of phrasal verbs and surrounding text toward precise semantic interpretations, the method reduces ambiguity that earlier techniques struggled to resolve. Empirically, the approach shows improvements across multiple metrics: Cosine Similarity (0.5889 to 0.6189), BLEU (0.2570 to 0.3150), ROUGE-L (0.4623 to 0.4901), Jaccard (0.3484 to 0.3648), and METEOR (0.3555 to 0.3607). Evaluation using semantic similarity and text-generation metrics confirms the approach's effectiveness for phrasal verb classification. Overall, the study highlights the potential of contextualized transformer models for fine-grained semantic understanding and demonstrates that integrating such models with in-context editing offers a practical, scalable, and effective contribution to semantic disambiguation research.

Keywords: Phrasal Verb Disambiguation, Contextualized Embeddings, BERT, In context Editing, Natural Language Processing

Decoding Phrasal Verbs with Context-Aware BERT Representations

Sujair Ibrahim[1], Abishethvarman Vadivel[2]*, Kumara BTGS[3]

[1,3]Department of Software Engineering, Sabaragamuwa University of Sri Lanka
[2]Department of Computing and Information Systems, Sabaragamuwa University of Sri Lanka

Introduction:

Guess the meaning of Phrasal Verb

"pick up"

- Is To collect or retrieve something
- Is To learn or acquire a skill gradually
- Is To improve or increase

Eg:

- **Can you pick up the package from the post office?**
- **She picked up Spanish while living in Mexico.**
- **Sales started to pick up after the promotion.**

What is Phrasal Verb?

Multi-word expressions with meanings often unrelated to individual words.

- Meaning **depends** heavily on **context** it is challenging for NLP models.
- NLP models, such as **BERT** and GPT, use self-attention mechanisms to capture contextual relationships between words

Problem Statement
- Traditional transformer models like BERT, while strong in syntax and token prediction, struggle with semantic disambiguation

Research Question
- Can in-context editing improve BERT's ability to understand contextual meaning of phrasal verbs?

Gap in Research

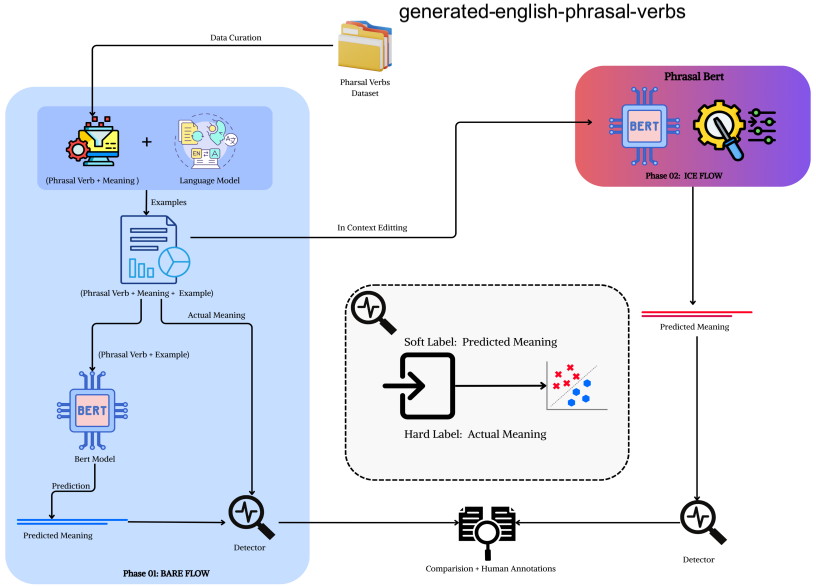
Limited research on **transformer-based phrasal verb disambiguation** with in-context editing methods.

Methodology:

Dataset - GitHub
generated-english-phrasal-verbs

3350 phrasal verbs
11,894 Instances

Training – 70%
Validation – 15%
Test – 15%



Results:

Benchmarks	Before In-context edit	After In-context edit
BLEU score	0.2570	**0.3150**
Cosine similarity	0.5889	0.6189
ROUGE-L	0.4623	0.4901
Jaccard Similarity	0.3484	0.3648
METEOR	0.3555	0.3607

Cosine & ROUGE-L gains: improved semantic + structural match.

Conclusions:

- **In-context editing improves performance:**
  - BERT achieves more accurate phrasal verb interpretation with in-context editing.
- **Better context understanding:**
  - In-context editing helps the model grasp contextual meanings better, resulting in more reliable and semantically rich NLP outputs.
- **Broader impact:**
  - These improvements strengthen the foundation for larger language models and support progress in idiom and multi-word expression processing.

Cybersecurity

Adaptive Adversarial Defense for Network Intrusion Detection Systems Using Dynamic GAN Ensemble

Vieronicka Kanesamoorthy*, Shriparen Sriskandarajah

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*vieronicka27@gmail.com

Adversarial Machine Learning (AML) has emerged as a critical threat to Network Intrusion Detection Systems (NIDS), enabling attackers to generate manipulated or synthetic network traffic that evades traditional machine learning–based defenses. The increasing dual use of Generative Adversarial Networks (GANs) further intensifies this challenge, as GANs can synthesize highly realistic malicious flows and adaptive perturbations that exploit structural vulnerabilities in NIDS models. However, existing defenses remain limited by severe dataset imbalance, poor minority-class representation, and insufficient robustness against evolving adversarial behaviours.This research proposes an Adaptive Adversarial Defense framework employing a Dynamic GAN Ensemble composed of Vanilla GAN, WGAN–GP, and CTGAN. The ensemble generates high quality synthetic samples and adversarially guided variants targeting the most vulnerable minority attack classes – Botnet and InfiltrationâĂŤin the CIC-IDS2017 dataset. The pipeline first mitigates imbalance through targeted synthetic augmentation, followed by adversarial guidance using Projected Gradient Descent, Carlini & Wagner, and HopSkipJump attacks generated on the training set to expose the model to diverse perturbation patterns. A defended Random Forest NIDS is trained using a combination of clean data, GAN-generated minority samples, and adversarially enriched samples. Experimental results demonstrate that the proposed approach significantly improves precision, recall, and F1-scores for minority classes compared to the baseline. Furthermore, the defended model exhibits substantially reduced adversarial attack success rates, achieving only 10.5% (PGD), 24.9% (C&W), and 80.9% (HSJ) on the test data, confirming stronger resilience under adaptive attack scenarios. Overall, the findings highlight the potential of dynamic GAN based augmentation combined with adversarial guidance as a promising and effective defensive strategy for modern NIDS.

Keywords: Network Intrusion Detection, Adversarial Machine Learning, GAN Ensemble, Data Augmentation, Cybersecurity

Adaptive Adversarial Defense for NIDS Using Dynamic GAN Ensemble

K.Vieronicka & S.Shriparen

Department of Computer Science, Faculty of Science, University of Jaffna

vieronicka27@gmail.com , shriparens@univ.jfn.ac.lk

Introduction

Network Intrusion Detection Systems (NIDS) increasingly face adversarial machine learning threats, where attackers craft malicious inputs to evade detection. Traditional machine learning-based NIDS struggle under adaptive adversarial attacks and are further weakened by class imbalance in datasets such as CIC-IDS2017. Recent studies show that Generative Adversarial Network (GAN)-based adversarial attack generators can degrade minority-class detection. This research proposes a Dynamic GAN Ensemble to enhance minority-class representation and adversarial robustness simultaneously.

Research Question- How can a defense mechanism based on GANs be employed to address the vulnerabilities of machine learning models against adversarial attacks thereby serving as a remedy for the prevailing issues of data imbalance & data scarcity?

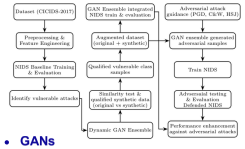
Contribution

- We propose a novel Dynamic GAN Ensemble to improve minority-class representation in CIC-IDS2017.
- We introduce an adversarial guidance strategy that trains GANs using Projected Gradient Descent, Carlini & Wagner, and HopSkipJump attacks to enhance robustness.
- We develop a hybrid defense training pipeline combining clean, synthetic, and adversarially perturbed samples.



Proposed Methodology

Detailed flow



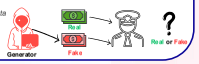
Key Strengths

- RF as baseline NIDS classifiers show strong performance on CIC-IDS2017 [6].
- GANs Vanilla GAN, WGAN-GP, CTGAN learn subtle patterns traditional models miss.
- Adversarial guidance prepares the model for real-world evasion attempts.
- Ensemble learning ensures diversity in synthetic attacks.
- Defended NIDS resists both noise and targeted manipulation.

- **GANs**

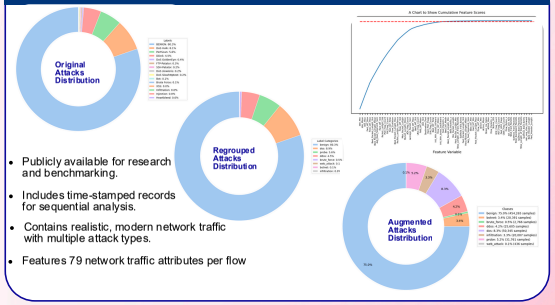


- *G – Generator (creates fake samples)*
- *D – Discriminator (distinguishes real vs fake)*
- *z – Random noise input to generator*
- *x – Real data sample*
- *BCE – Binary Cross Entropy is as*
- *Wasserstein - Distance between data distributions*
- *Gradient Penalty - Stabilizes WGAN-GP training*
- *MMD – Matches generated distribution to real data*



GANs share load equally and combine best from each to enhance performance & adversarial robustness.

Dataset



- Publicly available for research and benchmarking.
- Includes time-stamped records for sequential analysis.
- Contains realistic, modern network traffic with multiple attack types.
- Features 79 network traffic attributes per flow

Results & Discussion

- **Baseline NIDS performance evaluation**

Table 1: Baseline NIDS Classification Report (All Classes)

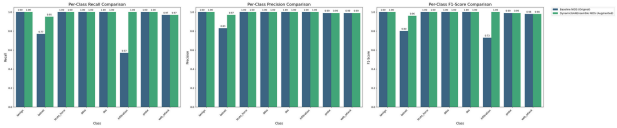
Class	Precision	Recall	F1-score	Support
benign	1.00	1.00	1.00	454264
botnet	0.83	0.77	0.80	391
brute_force	1.00	1.00	1.00	2767
ddos	1.00	1.00	1.00	25605
dos	1.00	1.00	1.00	50344
infiltration	1.00	0.57	0.73	7
probe	0.99	1.00	0.99	31761
web_attack	0.99	0.97	0.98	436

- **Performance comparison after synthetic data augmentation on vulnerable classes.**

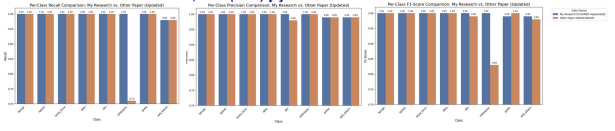
Table 2: Performance Comparison on after data augmentation

Method / Class	Precision	Recall	F1-Score
Keserwani et al. [8]			
Botnet	1.00	0.60	0.75
Lee et al. [7]			
Botnet	0.86	0.53	0.66
Zhao et al. [6]			
Botnet	1.00	1.00	1.00
Infiltration	1.00	0.71	0.83
Our Approach			
Botnet	1.00	1.00	1.00
Infiltration	1.00	1.00	1.00

- **BaselineNIDS performance evaluation & comparison after GAN ensemble training & synthetic data augmentation.**



- **GAN ensemble NIDS vs *Zhao, et al. (2024) [7]***



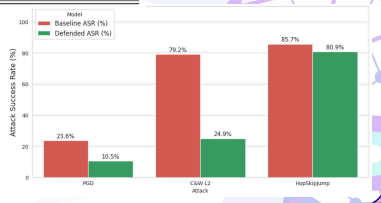
These results show better performance after addressing data imbalance problem.

- **Defended NIDS performance evaluation & comparison**

Table 3: Comparison with the state-of-the-art approaches on the CICIDS2017 dataset

Paper	PGD *ASR	C&W *ASR	HSJ *ASR
Alo et.al.[12] [GAN enhanced NIDS]	28.1	30.6	—
Sharma et. al.[10] (RF)	22.3	—	100.0
Zhang et. al.[11] (RF)	—	47.08	100.0
Ours (Dynamic GANEnsemble-RF)	10.5	24.9	80.9

*ASR = # successfully evaded attacks / # total attacks submitted



Lower the attack success rate (ASR) better the defense.

Conclusion

- This research demonstrates an adaptive adversarial defense strategy for NIDS using a Dynamic GAN Ensemble.
- The approach improves adversarial robustness, enhances minority-class detection, and provides insights into the evolution of adversarial signatures.
- Our approach achieves state-of-the-art results on CICIDS-2017 with a signature-based Random Forest using the same number of data.

References

[1] H. Jmila and M. I. Khedher, "Adversarial machine learning for network intrusion detection: A comparative study," Computer Networks, vol. 214, p. 109073, 2022.

[2] R. Panigrahi and S. Borah, "A detailed analysis of cicids2017 dataset for designing intrusion detection systems," International Journal of Engineering & Technology, vol. 7, no. 3.24, pp. 479–482, 2018.

[3] B. Kc, S. Sapkota, and A. Adhikari, "Generative adversarial networks in anomaly detection and malware detection: A comprehensive survey," Advances in Artificial Intelligence Research, vol. 4, no. 1, pp. 1835, 2024.

[4] S. O. Ale, A. S. Jamil, M. J. Hussein, M. K. Al-Dulaimi, S. W. Taha,and A. Khaporina, "Automated detection of cybersecurity threats using generative adversarial networks (gans)," in 2024 36th Confer- ence of Open Innovations Association (FRUCT), pp. 596–617, IEEE, 2024.

[5] D. Stiawan, M. Y. B. Idris, A. M. Bamhdi, R. Budiarto, et al., Cicids 2017 dataset feature analysis with information gain for anomaly detection, IEEE Access 8 (2020) 132911–132921.

[6] X. Zhao, K. W. Fok, and V. L. Thing, "Enhancing network intrusion detection performance using generative adversarial networks," arXiv preprint arXiv:2404.07464, 2024.

[7] J. Lee, K. Park, Gan-based imbalanced data intrusion detection system, Personal and Ubiquitous Computing 25 (2021) 121–128.

[8] P. K. Keserwani, M. C. Govil, E. S. Pilli, P. Govil, A smart anomaly-based intrusion detection system for the internet of things (iot) network using gwo–pso–rf model, Journal of Reliable Intelligent Environments 7 (2021) 3–21.

[9] Pujari, Medha, et al. "An approach to improve the robustness of machine learning based intrusion detection system models against the carlini-wagner attack." 2022 IEEE International Conference on Cyber Security and Resilience (CSR). IEEE, 2022.

[10] Sharma, Sandhya, and Zesheng Chen. "A Systematic Study of Adversarial Attacks Against Network Intrusion Detection Systems." Electronics 13.24 (2024): 5030.

[11] Zhang, Hangsheng, et al. "Explainable and transferable adversarial attack for ml-based network intrusion detectors." IEEE Transactions on Dependable and Secure Computing (2025).

[12] Alo, Salam Omar, et al. "Automated detection of cybersecurity threats using generative adversarial networks (gans)." 2024 36th Conference of Open Innovations Association (FRUCT). IEEE, 2024.

[13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, Ad vances in neural information processing systems 27 (2014).

[14] Adversarial Robustness Toolbox, GitHub. Available online: https://github.com/Trusted-AI/adversarial-robustness-toolbox (accessed on 18 December 2024).

[15] Ndayishukamiye, Thanoise, et al. "Adversarial Defense in Cybersecurity: A Systematic Review of GANs for Threat Detection and Mitigation." arXiv preprint arXiv:2509.20411 (2025).

Acknowledgement

Special thanks to Dr. Malka N. Halgamuge for her invaluable support by providing the initial proposal that shaped this research.

ARP Spoofing Prevention in SDN Using Host Agents and Cryptographically Verified DHCP IP-MAC Bindings

Sankalpa Fernando*, Shriparen Sriskandarajah

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*sankalpafernando2017@gmail.com

Address Resolution Protocol (ARP) spoofing enables attackers to intercept, modify, or block traffic by sending forged ARP messages. Traditional defences, including cryptographic ARP variants, kernel patches, and switch-based filtering, often introduce deployment complexity, compatibility issues, and scalability concerns. Detection-based mechanisms, such as machine learning classifiers or dedicated processing servers, provide only reactive protection and can add significant latency. We present a proactive ARP spoofing prevention mechanism designed for SDN environments, leveraging their centralized visibility and programmability to improve network security. The SDN controller observes DHCP transactions, generates verified IP-MAC mappings, and signs them cryptographically before publishing them to a Redis Stream. Lightweight host-side agents subscribe to the stream, verify digital signatures, and install static ARP entries locally, overriding any spoofed mappings. By using SDN's ability to monitor and control network-wide behavior programmatically, the system ensures spoofed packets are blocked at the host before they affect higher-layer applications. Implementation in Mininet demonstrates millisecond propagation of bindings and real-time blocking of spoofing attempts, even in networks of 40 hosts. Results show consistent 100% intrusion prevention once the DHCP server has assigned the IP addresses, with minimal computational overhead and excellent scalability. Compared to detection-heavy approaches, our architecture achieves stronger security guarantees while preserving low latency and ease of deployment. By combining SDN-based visibility, cryptographic verification, and efficient stream-based distribution, the system provides a practical and extensible solution for LAN and SDN environments, laying the foundation for future enhancements in resilience and recovery.

Keywords: ARP spoofing, SDN, DHCP, IPS

ARP Spoofing Prevention in SDN Using Host Agents and Cryptographically Verified DHCP IP–MAC Bindings

W.K.S Fernando and S. Shriparen
Department of Computer Science, University of Jaffna, Sri Lanka
sankalpafernando2017@gmail.com

INTRODUCTION

Address Resolution Protocol (ARP), defined in RFC 826, is stateless and lacks authentication. Attackers can easily broadcast forged ARP messages to associate their MAC address with a legitimate IP (ARP Spoofing), enabling Man-in-the-Middle (MitM) and DoS.
attacks

Limitations of Current Approaches:

- **Traditional Defence**: Cryptographic ARP variants and switch-based filtering introduce significant deployment complexity and hardware compatibility issues.
- **SDN Defence**: Existing SDN solutions (e.g., FICUR, SARP NAT, ML-based detection) are predominantly reactive. They detect attacks after they occur, often adding latency.

This research aims to design a proactive prevention mechanism that uses the SDN controller as the "source of truth" to enforce verified bindings at the host level.

METHODOLOGY

We propose a system integrated with dynamic IP addressing that requires no switch hardware modifications. It consists of three components:

1. SDN Controller (Ryu Application)

- **Monitoring:** Passively monitors DHCP traffic by capturing BOOTREPLY packets (UDP 67→68).
- **Verification:** Extracts the client IP (yiaddr) and MAC address from the offer.
- **Signing:** Signs the IP-MAC binding using an RSA private key.
- **Distribution:** Publishes the signed binding to a Redis Stream (dhcp_stream).

2. Redis Server (Message Broker)

- Operates on the same host as the controller for reliability.
- Uses Access Control Lists (ACLs) to ensure only the controller can write to the stream.
- Provides persistence, allowing agents to recover missed messages.

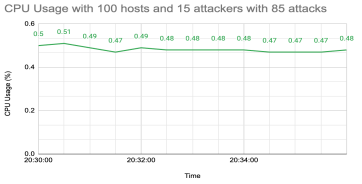
3. Host Agent (Enforcer)

- Pre-loaded with the Controller's Public Key.
- Subscribes to dhcp_stream updates.
- **Action:** Verifies the digital signature and installs a **Static ARP Entry**. Static entries strictly override dynamic ARP replies, rendering spoofing attempts futile

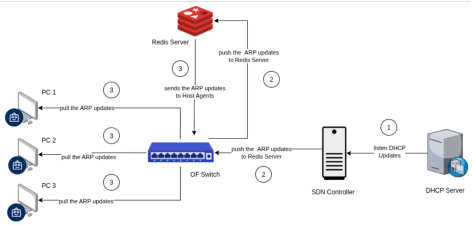
RESULTS

Evaluated in Mininet with network sizes from 3 to 40 hosts. Attacks were simulated using the arpspoof tool.

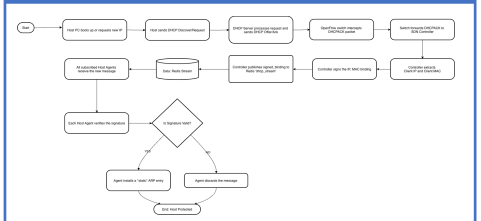
Hosts	Attackers	Prop. Delay (s)	IPR (%)
3	1	0.02	100
10	2	0.03	100
40	5	0.08	100



SYSTEM ARCHITECTURE



SYSTEM WORKFLOW



CONCLUSION

We presented a scalable defence against ARP spoofing that leverages SDN programmability and cryptographic verification. By eliminating reliance on insecure ARP broadcasts for verified hosts, we ensure robust SDN security.

Limitation: The system currently relies on the Redis server availability; there is no fallback mechanism if the Redis database is lost during an outage.

Future Work: Developing a distributed recovery strategy to restore bindings if the central store fails.

DISCUSSION

Proactive vs. Reactive: Unlike detection systems that alert administrators after cache pollution, our system proactively locks the ARP cache. Spoofed packets are blocked immediately at the host interface.

Scalability & Performance:

- Redis Streams allow efficient, asynchronous distribution of bindings without overwhelming the controller.
- The overhead on the host agent is negligible (signature verification only).

Security Analysis:

- **Integrity**: RSA signatures prevent attackers from injecting fake bindings into the stream.
- **Resilience**: Even if the controller is busy, Redis queues ensure no binding updates are lost.

Behavioural Analysis for Suspicious Transaction in Bitcoin Network

Raseswaran Kanistan[1], Jeyakumar Samantha Tharani

Department of Computer Science, University of Jaffna, Sri Lanka

*rkanistan2000@gmail.com

Blockchain networks enable decentralized digital asset exchange with transparency and trust, yet the pseudonymity of participants can be misused to conceal illicit activities such as money laundering, ransomware, and Ponzi schemes. In this study, we analyse behavioural and temporal features of Bitcoin transactions using the Elliptic++ dataset. The methodology involves five phases: data utilisation, preprocessing, feature selection, analysis, and validation. Transactions are chronologically split by Time step (0-30 for training/validation, 31-49 for testing), and duplicate or null entries are removed. Feature selection is performed using Pearson and Spearman correlation measures, identifying influential features such as first block appeared in. Three experimental setups are considered: (1) all 56 wallet address features, (2) 31 features selected via Pearson correlation, and (3) 28 features selected via Spearman correlation. Classification is performed using Random Forest and XGBoost, with evaluation via accuracy, precision, recall, and F1-score. Explainable AI using SHAP highlights key features affecting predictions. Results indicate that time-dependent and behavioural features play a crucial role in detecting suspicious activities, with XGBoost achieving the highest F1-score of approximately 71%. Limitations include the restricted dataset timeframe and lack of metadata on timestep generation and label derivation. Random Forest and XGBoost were chosen as standard methods, while future work will explore additional machine learning models such as Support Vector Machines, Neural Networks, and Long Short-Term Memory networks on multiple datasets to further validate the approach. Overall, this study demonstrates the effectiveness of temporal and behavioural features in identifying illicit Bitcoin transactions and contributes to the automated detection of money laundering and ransomware-related activities.

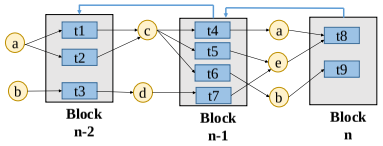
Keywords: Bitcoin, blockchain, money laundering, ransomware, machine learning

Behavioural Analysis for Suspicious Transaction in Bitcoin Network

Raseswaran Kanistan, Jeyakumar Samantha Tharani
Department of Computer Science, Faculty of Science, University of Jaffna
rkanistan2000@gmail.com

Introduction

Cryptocurrencies, such as Bitcoin, enable decentralized digital transactions with transparency and security. However, their pseudonymous nature can be exploited for illicit activities, including money laundering, ransomware, and darknet market transactions. Traditional studies focus on structural and transactional features, such as transaction counts, amounts, and network connectivity. Temporal characteristics of transactions(time-based patterns) have been underexplored, even though they can reveal suspicious behaviours. This study investigates time-dependent features of Bitcoin transactions using the Elliptic++ dataset. Machine learning models (Random Forest and XGBoost) and SHAP explainable AI are employed to identify which temporal and behavioural features most strongly indicate suspicious activity.



Address	Timestep
a	t1, t2, t4, t8
b	t3, t6, t9

Figure 1 : Timestep in the Bitcoin blockchain network

Objectives

- Identify behavioural patterns in Bitcoin transactions
- Classify suspicious activity based on transaction features
- Evaluate the impact of the time-based features using eXplainable AI

Methodology

➢Data Utilisation : Elliptic++ Dataset
➢Data Preprocessing : Remove duplicates & null rows, Scaling
➢Feature Analysis and Selection : Pearson & Spearman correlations
➢Experimental Setup :
- Data splitting
- Classification (RF and XGBoost)
➢Analysis : SHAP explainable AI
➢Validation : Precision, Recall, F1-score and Accuracy

Table 1 : Experimental data

Class	Training	Testing	Total
1	7755	6511	14266
2	147375	103713	251088
3	327603	229985	557588
Total	482733	340209	822942

Experiments

- **Setup 1:** All 56 features
- **Setup 2:** 31 features selected via Pearson correlation
- **Setup 3:** 28 features selected via Spearman correlation

Goal: Compare model performance and assess impact of temporal features

Results and Discussion

The classification results across the three experimental setups show that XGBoost consistently achieves higher precision, recall, and F1-score than Random Forest, with the best performance obtained using Pearson-selected features. The SHAP analysis further reveals that temporal attributes—such as *first block appeared in* and *last block appeared in*—along with fee-related features play the most influential role in identifying suspicious Bitcoin transactions.

Table 2 : Classification results for **Experiments I, II, and III**

Experiment	Random Forest				XGBoost			
	pre	recall	f1-score	acc	pre	recall	f1-score	acc
I	0.6748	0.6980	0.6778	0.6980	0.7063	0.6987	0.6963	0.6987
II	**0.6860**	**0.7096**	**0.6877**	**0.7096**	**0.7241**	**0.7132**	**0.7181**	**0.7132**
III	0.6696	0.6958	0.6728	0.6958	0.6942	0.6721	0.6761	0.6721

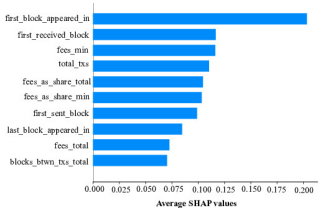


Figure 2 : SHAP Values for the features used in Experiment 3

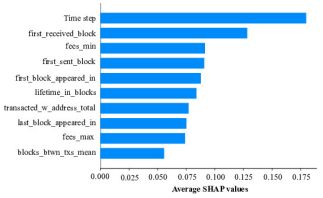


Figure 3 : SHAP Values for the features used in Experiment 1

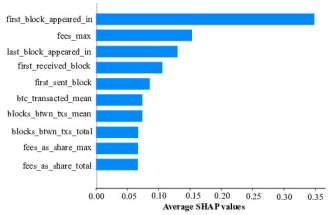


Figure 4 : SHAP Values for the features used in Experiment 2

Conclusion

- Behavioural and temporal features are crucial for detecting suspicious Bitcoin transactions.
- XGBoost achieved the best performance, confirming the importance of time-dependent features.
- Limitations include the dataset's restricted timeframe and lack of details on timestep generation and labels.
- Future work will combine graph-based and temporal deep learning models across multiple datasets to improve detection and reduce false positives.

References

[1] Elmougy, Y., & Liu, L. (2023). Demystifying fraudulent transactions and illicit nodes in the bitcoin network for financial forensics. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 3979-3990).
[2] Kute, D. V., Pradhan, B., Shukla, N., & Alamri, A. (2024). Explainable deep learning model for predicting money laundering transactions. *International Journal on Smart Sensing and Intelligent Systems*, *17*(1).
[3] Lorenz, J., Silva, M. I., Aparício, D., Ascensão, J. T., & Bizarro, P. (2020, October). Machine learning methods to detect money laundering in the bitcoin blockchain in the presence of label scarcity. In *Proceedings of the first ACM international conference on AI in finance* (pp. 1-8).
[4] Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., & Leiserson, C. E. (2019). Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *arXiv preprint arXiv:1908.02591*.
[5] Tharani, J. S., Charles, E. Y. A., Hóu, Z., Palaniswami, M., & Muthukkumarasamy, V. (2021). Graph based visualisation techniques for analysis of blockchain transactions. In *2021 IEEE 46th Conference on Local Computer Networks (LCN)* (pp. 427-430).

Zero Trust-based Network Traffic Anomaly Detection using Rule-based Approaches

S. W. Chamindu Lakshan Wijerathna*, Shriparen Sriskandarajah

Department of Computer Science, University of Jaffna, Sri Lanka

*swchamindulakshan99@gmail.com

Modern distributed networks face increasingly sophisticated cyber threats, yet conventional intrusion detection systems often struggle with real-time detection, interpretability, and resource efficiency. While machine learning–based approaches exist, they frequently incur high computational overhead, limited transparency, and insufficient alignment with Zero Trust security principles, leaving a gap for lightweight, adaptive, and explainable anomaly detection solutions. This research presents a rule-based anomaly detection system within a Zero Trust Architecture (ZTA) to enhance network security in distributed environments. Zero Trust enforces "never trust, always verify," requiring continuous inspection of all connections. Using the NSL-KDD (KDDTrain+) dataset, category-specific rules were developed to detect DoS, Probe, R2L, and U2R attacks. The methodology includes dataset preprocessing, attack mapping, statistical analysis of normal and attack traffic, and threshold extraction via mean, standard deviation, and interquartile ranges. Threshold updates are manual but support periodic re-evaluation of traffic, aligning with Zero Trust's continuous monitoring principle. Experiments show the system achieves high accuracy with low computational overhead, making it suitable for real-time detection. The model aligns with Zero Trust practices through constant verification, fine-grained rule enforcement, and evaluation of every connection rather than relying on perimeter security. The findings demonstrate that a structured rule-based approach can provide effective anomaly detection and practical defense against evolving cyber threats.

Keywords: Anomaly detection, Zero Trust, Rule-based detection, Network traffic analysis, Cybersecurity

Zero Trust-Based Network Traffic Anomaly Detection Using Rule-Based Approaches

Wijerathna, S W C L., Shriparen, S.
Department of Computer Science, Faculty of Science, University of Jaffna
2020csc058@univ.jfn.ac.lk , shriparens@univ.jfn.ac.lk

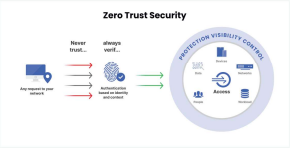
Abstract

Modern distributed networks face increasingly sophisticated cyber threats, while conventional intrusion detection systems struggle with real-time detection, interpretability, and efficiency. Machine-learning approaches exist but often require high computational power, lack transparency, and do not fully align with Zero Trust principles, creating a need for lightweight and explainable solutions. This research introduces a rule-based anomaly detection system integrated into Zero Trust Architecture (ZTA), which enforces "never trust, always verify" through continuous inspection of all connections. Using the NSL-KDD (KDDTrain+) dataset, category-specific rules were created for DoS, Probe, R2L, and U2R attacks through preprocessing, attack mapping, and statistical thresholding using mean, standard deviation, and IQR. Manual threshold updates support periodic traffic re-evaluation. The system achieves high accuracy with low computational cost, providing effective real-time anomaly detection aligned with Zero Trust requirements.

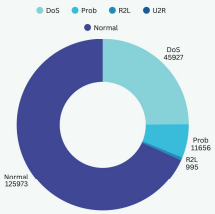
Introduction

Cyberattacks are becoming increasingly complex, making it difficult for organizations to secure their networks using traditional perimeter-based methods. As modern systems rely on distributed and cloud environments, Zero Trust Architecture (ZTA) has become essential, enforcing the principle of "never trust, always verify" through continuous inspection of every connection. Within this model, anomaly detection is critical for identifying abnormal or malicious network behavior in real time.

However, many existing anomaly detection systems struggle with scalability, high computational cost, and reduced accuracy when traffic patterns change. Traditional rule-based methods also become harder to maintain as rule sets grow more complex. This research introduces a lightweight rule-based anomaly detection system aligned with Zero Trust principles. Using the NSL-KDD (KDDTrain+) dataset, the system derives category-specific rules for detecting major attack types such as DoS, Probe, R2L, and U2R. By relying on statistically defined thresholds, the model remains simple, fast, and explainable, while supporting the continuous verification requirements central to Zero Trust security.



Dataset

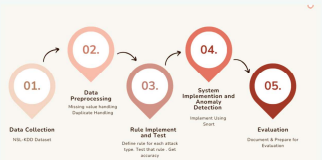


It includes normal traffic and four major attack categories:

- DoS – Denial of Service
- Probe – Information gathering/scanning
- R2L – Remote to Local attacks
- U2R – User to Root privilege escalation

The dataset is balanced, well-structured, and widely used for evaluating rule-based and anomaly detection approaches.
Its feature-rich structure allows extraction of meaningful statistical thresholds for rule creation.

Methodology



Data Preprocessing:

- Checked for missing values, duplicates, and inconsistent entries.
- Mapped all attack types into four main categories: DoS, Probe, R2L, U2R.
- Created a binary label (Normal = 0, Attack = 1) for overall anomaly detection.

Feature Selection:

- Analyzed key traffic features relevant to each attack type.
- Extracted statistical values (Mean, Standard Deviation, IQR) from normal traffic only.

Rule Creation:

- DoS → high traffic volume, high serror rate, repeated connections
- Probe → scanning-like behavior, high srv_count, unusual service access
- R2L → low src_bytes, suspicious login attempts, abnormal flags
- U2R → privilege escalation indicators, rare command patterns
- Each rule checks whether specific features exceed calculated thresholds.

```
alert udp any any -> any any (msg:"UDP_SMALL_SCAN"; dsize:<25; classtype:attempted-recon; sid:1000004; rev:1;)
```

Zero Trust Alignment:

- Every connection is evaluated individually following the "never trust, always verify" principle.
- Rules enforce continuous inspection, not perimeter-based assumptions.
- Supports lightweight, real-time validation suitable for Zero Trust environments.

Results & Discussion

```
Classification Report:
              precision    recall  f1-score   support

           0       0.93      0.97      0.95     67343
           1       0.96      0.92      0.94     58630

    accuracy                           0.9425         0
```

```
Accuracy: 0.94

Confusion Matrix
--------------------------
[[65049  2294]
 [ 4945 53685]]
```

```
Total:                35
------------------------------------------
Action Stats:
     Alerts:          0 (  0.000%)
------------------------------------------
Total:                6
------------------------------------------
Action Stats:
     Alerts:          4 ( 66.667%)
```

Strong Detection Performance:

- explainable thresholds that react reliably to key traffic features. It runs fast, requires very little computation, and maintains stable performance even when tested on unseen data.

Alignment with Zero Trust:

- Decisions rely on current behavior—not past trust—matching core Zero Trust principles.
- Each connection is evaluated independently, supporting continuous verification.

Key Limitations:

- Rules need periodic updates to handle new attack behaviors.
- NSL-KDD is older and less complex than modern traffic.
- Real networks may introduce noise that challenges hard thresholds.

Conclusion

This research demonstrates that a lightweight rule-based anomaly detection framework can effectively support Zero Trust security by continuously evaluating network traffic and applying clearly defined, category-specific rules. Using the NSL-KDD dataset, the system successfully identified major attack types—including DoS, Probe, R2L, and U2R—with strong accuracy and minimal computational cost. The transparent nature of rule-based decision-making ensures easy interpretation and policy alignment, which is essential in Zero Trust environments where continuous verification is required. Although the approach relies on manually updated thresholds, this process fits naturally within the periodic policy review cycles expected in Zero Trust operations.

Future work will focus on extending this system to handle more diverse and modern traffic patterns, improving adaptability to previously unseen attack behaviors, and integrating automated threshold tuning to minimize manual rule updates. Evaluating the model with larger datasets such as CICIDS2017 or UNSW-NB15—using distributed or high-performance computing environments—would further strengthen its real-world applicability.

References

1. Kindervag, J. (2010). Build Security Into Your Network's DNA: The Zero Trust Network Architecture. Forrester Research.
2. National Institute of Standards and Technology (NIST). (2020). Zero Trust Architecture (SP 800-207).
3. Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 dataset. IEEE Symposium on Computational Intelligence for Security and Defense Applications.
4. Kumar, J., & Gupta, D. (2019). A Review on KDD Cup99 and NSL-KDD Dataset. International Journal of Advanced Research in Computer Science.
5. Roesch, M. (1999). Snort: Lightweight Intrusion Detection for Networks. Proceedings of the 13th USENIX Conference on System Administration.

Natural Language Processing

Sinhala Handwritten Optical Character Recognition

P. N. Nipun Sadeepa Weerasinghe*, Kengatharaiyer Sarveswaran

Department of Computer Science, University of Jaffna, Sri Lanka

*nipunsw1999@gmail.com

This research presents the development of a Sinhala handwritten Optical Character Recognition (OCR) system. Sinhala, one of the official languages of Sri Lanka, features a highly complex script with numerous characters, diacritic modifiers (pillam), ligatures, and context-dependent variations, making handwritten recognition a particularly challenging task. To address this, we compiled a dataset comprising 315 classes selected based on contemporary Sinhala usage and expert consultation. Each class contains 134 handwritten samples collected from 138 participants (83 males and 55 females), representing school students (ages 14–18), university students (21–23), government employees, and school teachers. The 315 classes represent a comprehensive set derived from a 240M-word Sinhala corpus (∼1 billion characters), ensuring full coverage of all characters, ligatures, and modifiers used in contemporary Sinhala. This design enables accurate representation of contemporary Sinhala writing practices, particularly cases where modifiers appear before, after, or on both sides of a base character. Our approach focuses on non-touching, non-overlapping handwritten characters and lines, thereby simplifying segmentation and improving recognition reliability. The dataset forms the foundation for training convolutional and transfer-learning-based OCR models, including CNNs and EfficientNet architectures, aimed at enhancing the accuracy and robustness of Sinhala handwritten text recognition. Baseline experiments using a custom CNN achieved 30% test accuracy, while transfer learning with EfficientNetB0 improved performance to 65%, demonstrating the dataset's suitability for benchmarking Sinhala handwritten OCR. The complete dataset and preprocessing tools will be publicly released upon publication.

Keywords: Handwritten Sinhala character recognition, Data preprocessing, Multi-class classification, Unicode mapping

Sinhala Handwritten Optical Character Recognition

P.N.N.S. Weerasinghe*, K. Sarveswaran
Department of Computer Science, University of Jaffna
*nipunsw1999@gmail.com

Introduction

- Typing in Sinhala Unicode is difficult due to complex input methods.
- Sinhala has many visually similar character forms, making handwritten recognition challenging.
- This research develops a handwritten OCR system that recognises the full Sinhala character set for official documents.
- The resulting application converts handwritten Sinhala content into editable Unicode text.

Objectives

- Create a modern Sinhala handwritten dataset using samples from government and private-sector documents.
- Preprocess captured images to improve readability and dataset quality.
- Develop a handwritten OCR model to recognise modern Sinhala characters.

Background

- Sinhala letters include ligatures formed by combining multiple glyphs, including base characters, special characters and vowel modifiers.
- In these ligatures, glyphs may appear in front of, behind, above, or below the base character.



Representing using less classes

Defined types: F, L, B

Patterns

Methodology

Dataset Creation

- Collected Sinhala handwritten characters across 315 classes with 134 samples per class.
- Gathered data from diverse groups: school students, university students, government employees, and teachers.
- Used a custom Python annotation tool to label and organise samples into class-specific folders.

Preprocessing

- Converted images to grayscale and applied Otsu's thresholding for binarisation.
- Denoised images using the Non-Local Means algorithm.
- Performed edge detection and perspective correction to straighten angled scans.
- Applied sharpening to enhance character edges.
- Used morphological operations (opening/closing) to remove noise and fill gaps.
- Normalised backgrounds to reduce shadows and paper artefacts.

Segmentation

- Applied a pixel-based algorithm to segment lines and individual letters.
- Used morphological techniques to segment words from text lines.

Character Recognition

- Trained models using a Convolutional Neural Network (CNN).
- Applied transfer learning with EfficientNet-B0 for improved accuracy.

Unicode Generation

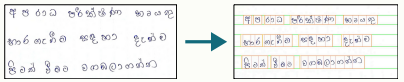
- Implemented a Unicode sequencing mechanism for Sinhala script.
- Developed a pattern-recognition module to output the correct Sinhala Unicode characters.

Results and Discussion

Image preprocessing to create clear image



Lines, words and characters segmentation



Models evaluation

Approach	Validation Accuracy	Testing accuracy	Evaluation
CNN	57%	30%	564 wrong predictions out of 805
Transfer Learning (EfficientNetB0)	89%	65%	241 wrong predictions out of 805

Conclusion

- Developed a full workflow for Sinhala handwritten character recognition across 315 classes.
- Created a diverse dataset and applied strong preprocessing and segmentation methods to enhance input quality.
- Achieved 65% accuracy using CNN and EfficientNet-B0 models.
- Demonstrated the feasibility of converting handwritten Sinhala into editable Unicode text.
- Provides a foundation for future improvements in Sinhala handwritten OCR and digital documentation.

Selected References

- Karunaratne M. L., Wijesiriwardana C. P., Nishantha K. M. I., & Kumara W. G. C. W. (2024). Efficiency and accuracy in Sinhala handwritten character recognition: A Gabor-initialized CNN perspective. Sri Lankan Journal of Technology, 2024(1).
- Karunanayaka M. L. M., Marasinghe C. A., & Kodikara N. D. (2005). Thresholding, noise reduction, and skew correction of Sinhala handwritten words. In Proceedings of the IAPR Conference on Machine Vision Applications (IAPR MVA 2005). Tsukuba Science City, Japan.
- Walawage K. S. A., & Ranathunga L. (2018). Segmentation of overlapping and touching Sinhala handwritten characters. In 2018 3rd International Conference on Information Technology Research (ICITR) (pp. 1–6). IEEE. https://doi.org/10.1109/ICITR.2018.8736129

Analysing Ambiguity in Anaphoric Reference: A Rule-Based Framework and a Proposal for Tamil

Lukas Beiske

Department of Linguistics, University of Konstanz, Germany

*lukas.beiske@uni-konstanz.de

We present a rule-based categorisation framework, that assigns descriptive labels to anaphorically ambiguous NPs annotated in the ARRAU and Phrase Detectives corpora. These corpora treat ambiguity as a meaningful feature rather than error or noise, since it is part of natural language use. Understanding ambiguity is of high importance for NLP applications like generative systems, machine translation, information extraction or text summarisation. We present a framework that assigns descriptive labels to ambiguous NPs by comparing the attribute values of two annotated alternative interpretations – such as referentiality, discourse status, bridging, semantic category, and genericity. The additional layer of annotation captures the underlying source of ambiguity, making the data more interpretable by summarizing diverging attributes. Thus, the framework identifies different types of anaphoric ambiguity. This supports both linguistic theory and the development of more robust NLP systems, while also enabling quantitative analysis: we identified 65 unique labels describing specific types of ambiguity that were assigned to 1,336 NPs, with just 16 accounting for over 89% of all cases. Finally, we point out avenues to extend annotation efforts and subsequent analysis from English to morphologically richer languages like Tamil. Among other phenomena, this would allow the study of zero-anaphora which adds further complexity and introduces more ambiguity.

Keywords: anaphoric ambiguity, ambiguity annotation, rule-based categorisation, Tamil

Analysing Ambiguity in Anaphoric Reference: A Rule-Based Framework and a Proposal for Tamil

Lukas Beiske
Department of Linguistics
University of Konstanz

Background and Motivation

What is Anaphora?

- **Anaphora** or **anaphoric expressions,** e.g. *pronouns*, make **reference** to preceding noun phrases (NPs), i.e. their **antecedent**

  (1) Markus was petting [the cat]$_i$. [It]$_i$ was cuddly.

- Enable discourse cohesion and brevity

Anaphoric Ambiguity

- Often it cannot be determined (precisely) to which NP an anaphor refers to:

  (2) Put the water in your bag and take **it** with you

  Does the pronoun *it* refer to *the water*, *your bag*, or both?

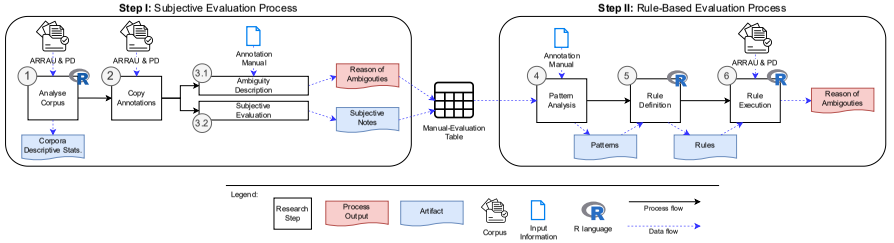
- Most annotation schemes force a single *"correct"* answer, ignoring the fact that ambiguity can be genuine

Ambiguity Matters:

Ambiguity in anaphoric reference is not an annotation error but an inherent property of natural language.

- Corpora like *ARRAU* and *Phrase Detectives* encode ambiguity either explicitly or implicitly (via **disagreement** between annotators), treating it as a meaningful feature
- Understanding anaphoric ambiguity is essential for NLP systems that rely on reference resolution, e.g.: translation, summarisation, information extraction

The Rule-Based Categorisation Framework: Methodology and Results



Objective:

Categorisation of every ambiguously annotated markable in the corpora and provide a label describing the ambiguity at hand.

Process:

Comparison of each attribute value of the 1st and 2nd interpretation, then assign a descriptive label based on comparison.

Compared Attributes:

- Discourse status:
  (3) [Markus]$_i$ was petting [the cat]$_k$. [It]$_k$ likes [him]$_j$.
  [Markus], [the cat] = DN; [him], [the animal], [It] = DO
- Diverging antecedents:
  (4) Put [the water]$_i$ in your [bag]$_j$ and take [it]$_{i\&?}$.
  [it] = [the water] or [bag]?
- Bridging relation:
  (5) We can't take [the car]$_k$. [The wheel]$_k$ is broken.
  The unusability of the car is inferred or "bridged" from the brokenness of the wheel
- Genericity:
  (6) [Ravens]$_k$ are black. [They]$_k$ are also smart.
  [They] refers to the kind of bird, and not specific ravens.
- Semantic category:
  (7) [The book] is nice.
  concrete = the physical book vs. abstract = its content.
- Referentiality:
  (7) [It]'s raining. / He is [a pirate].
  [It] = expletive, [a pirate] = predicate

Descriptive Statistics:

- 1.336 ambiguous NPs evaluated
- 65 unique labels assigned
- Small number accounts for large part of all NPs:

Table 4.3 Distribution of the 16 Most Frequent Labels in ARRAU and PD

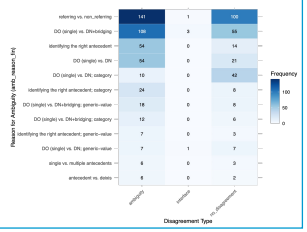
	Reason for Ambiguity (Label in amb_reason_fin)	Count	Percentage (%)	
1	referring vs. non_referring	282	20.43	
2	DO (single) vs. DN+bridging	259	18.77	→48.69
3	identifying the right antecedent	131	9.49	
4	DO (single) vs. DN	115	8.33	
5	DO (single) vs. DN; generic-value	84	6.09	
6	DO (single) vs. DN+bridging; generic-value	78	5.65	
7	DO (single) vs. DN; category	63	4.57	
8	identifying the right antecedent; category	43	3.12	
9	DO (single) vs. DN+bridging; category	39	2.83	→89.21
10	single vs. multiple antecedents	24	1.74	
11	identifying the right antecedent; generic-value	23	1.67	
12	antecedent vs. deixis	21	1.52	
13	identifying the right segment for deixis	21	1.52	
14	DN vs. DN, no related objects; generic-value	19	1.38	
15	antecedent vs. deixis; category	15	1.09	
16	DO (single) vs. DN+bridging; generic-value; category	14	1.01	

- Example:
  (8) "[…] why don't we couple [engine E2] to [the boxcar] and send **[it]** to Corning"

  Label for [it] = `identifying the right antecedent`

65 Diffrent Types of Ambiguity found in the Data

- Unique labels: 21
- Rare label (n ≤ 5): 21
- Assessment of disagreement:



Discussion and Conclusion

- Ambiguity is a **natural feature**, not a flaw of language or plain noise
- The framework **systematically categorises and labels ambiguous NPs** in the corpora: it identifies different **types of ambiguity**
- Categorised data becomes **more interpretable** for humans and machines by summarising diverging attribute features in a label
- Enables **quantitative analysis** of types of ambiguity (also together with disagreement)
- A small set of ambiguity types accounts for most cases found in the data
- Ambiguity reasons are diverse, but show clear distributional tendencies
- Findings align with linguistic theory (e.g. Sloppiness Hypotheses, Underspecification, Near-Identity relations)

Proposal for Tamil

- Existing resources mostly for English
- Tamil shows rich morphology and frequently omits NPs and pronouns (i.e. **zero-anaphora**)
- Adaptation requires **placeholder tokens** to represent omitted material
- Supported by existing Tamil NLP tools (morphological analyser, PoS tagger, chunker, clause boundary identifier and named entity recogniser)
- Would enable **ambiguity-aware annotation for underrepresented languages**
- Possible annotation scheme:



Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)
Poster ID: P33

Building A Baseline Sinhala Text-To-Speech System

Vigneswaran Wanith Nirman*, Ahrane Mahaganapathy, Kengatharaiyer Sarveswaran

Department of Computer Science,University of Jaffna, Sri Lanka

*wanithnirmanbc@gmail.com

This research proposes a baseline Sinhala Text-to-Speech (TTS) system that synthesizes natural and intelligible speech. Sinhala presents challenges for TTS due to its large inventory of consonant clusters, diphthongs, schwa detection rules and the scarcity of high-quality corpora. To address these, a five-hour single female-speaker audiobook dataset in the literary domain consisting of 7,664 unique words and covering the Sinhala phonetic inventory of 14 vowels and 40 consonants was compiled. A rule-based transcription-module was developed to transcribe text into the International Phonetic Alphabet (IPA) which was used to train a TTS system. A Tacotron-2 model was trained on the compiled dataset with a hop length of 256 and an Fast Fourier Transform size of 1,024 to generate mel-spectrograms. During inference, Sinhala text was first converted into IPA using the transcription-module, after which Tacotron-2 processed the input to generate mel-spectrograms. The Griffin–Lim algorithm (iterations:60, power:1.5) was subsequently applied for waveform reconstruction. The system was evaluated with 21 test sentences covering all phonemes including rare words and pronunciation-challenging words. Evaluation using Word Error Rate (WER) and Mean Opinion Scores (MOS) yielded a WER-5.55% and MOS naturalness:2.45, intelligibility:2.67, and speaker similarity:3.67 reflecting a moderate perceptual quality compared to TacoSi (MOS:4.39) and VITS-based system (MOS:4.18-4.62). The lower MOS is primarily attributed to the limited dataset. The baseline Sinhala TTS system, developed under limited computational resources, can be improved through larger corpora, neural vocoders, and enhanced text normalisation.

Keywords: Sinhala, Text-to-Speech, Tacotron-2, Grapheme-to-Phoneme, Griffin-Lim

BUILDING A BASELINE SINHALA TEXT-TO-SPEECH SYSTEM

Vigneswaran Wanith Nirman, Ahrane Mahaganapathy and Kengatharaiyer Sarveswaran
Department of Computer Science, University of Jaffna
wanithnirmanbc@gmail.com

SLSWCS Student Workshop on Computer Science

1. Introduction

- Sinhala is an Indo-Aryan language with complex phonology and challenging grapheme-to-phoneme (G2P) rules.



Figure 1: Sample G2P conversion in SInhala

- Challenges in developing Sinhala TTS
  - Limited high-quality speech corpora.
  - Grapheme-to-Phoneme (G2P) rules:
    schwa deletion, ZWJ usage,
    consonant clusters, diphthongs handling.
- Naturalness and intelligibility of the existing systems can be further improved to enhance the user experience. [1]

This research proposes a baseline Sinhala Text-to-Speech (TTS) system that generates natural and intelligible soeech with limited computational resources

2. Overview of Dataset

Source: Sera Audio Books, "Nidahas Wahallu"
domain cover: Literary
Format: WAV, 22.05 kHz, 16-bit PCM
Unique words used: 7,664
Phoneme coverage: 88.89% (In contemporary usage)

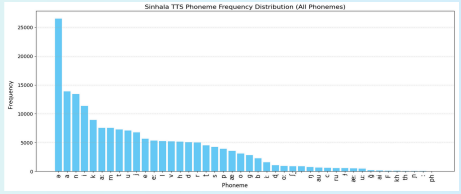


Figure 2: Phoneme frequency distribution of the dataset

3.Methodology

Methodology of the Text to Speech synthesis

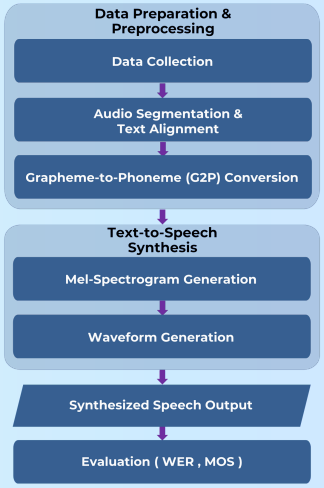


Figure 3: Methodology of the TTS

Modules and key parameters

Table 1-Modules and key parameters

Modules	Method	Key Parameters
G2P Conversion	G2P tool was developed for the phonetic transcription	ZWJ handling, Schwa deletion..etc
Mel-Spectrogram Generation	Tacotron 2	Epochs - 1,250 Batch size-32 Learning rate-0.0001 FFT-1,024 Hop length-256
Waveform Generation	Griffin-Lim	Iteration-60 Power-1.5

4.Results and Discussion

- A curated test set covering all the contexts of pronunciation rules, including rare and challenging words, was used to evaluate TTS accuracy.

Table 2-Results of WER and MOS

Metrics	Value		Interpretation
Word Error Rate (WER)	5.55%		Indicates strong word level accuracy.
Mean Opinion Score (MOS)	Naturalness-2.45 Inteligibility-2.67 Speaker Similarity-3.67	Max 5.00	Indicates moderate perceptual quality

- The moderate MOS results due the small speech corpus and limited phoneme frequency coverage.

5.Conclusion

- The baseline Sinhala TTS system successfully integrates rule-based phonetic transcription module we developed with neural synthesis.
- Limitations : Low phoneme-frequency coverage, Small speech corpus
- Future works : Expand dataset, Add text-normalization module, Use improved neural vocoders (HiFi-GAN)

Selected References

[1] T.V.R.Weerakoon, K.K.S.Nayanathara, and L.I.L.Harischandra, "Sinhala Speech-to-Speech Chatbot Using Deep Learning Approaches,"University ofColombo School of Computing,2025.[Online]. Available: https://dl.ucsc.cmb.ac.lk/jspui/handle/123456789/4905
[2] H. Liu, T. Baoueb, M. Fontaine, J. Le Roux, and G. Richard, "GLA-Grad: A Griffin-Lim extended waveform generation diffusion model," in *Proc. IEEE Int. Conf. Acoust., Speech & Signal Process. (ICASSP)*, 2024.
[3] T. Nadungodage, C. Liyanage, A. Perera, R. Pushpananda, and R. Weerasinghe, "Sinhala G2P Conversion for Speech Processing," in *Proc. 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages*, Colombo, Sri Lanka, 2018, pp. 1–6. [Online]. Available:https://www.researchgate.net/publication/328072699_Sinhala_G2P_Conversion_for_Speech_Processing

Poster ID: P34

Dialect Divergence in Sri Lankan Tamil: A Corpus-Based Computational Analysis

Gobalakrishnar Nilusana*, Kengathaiyar Sarveswaran

Department of Computer Science, University of Jaffna, Sri Lanka

*nilusana07gobalakrishnar@gmail.com

Sri Lankan Tamil exhibits substantial dialectal variation across regions including Jaffna, Batticaloa, Malayagam, Moor, and Trincomalee, each distinguished by phonetic, lexical, and syntactic features. Despite this diversity, computational resources for these dialects remain scarce, creating significant challenges for automated processing and NLP applications. In this study, we develop a Sri Lankan Tamil dialect corpus and conduct divergence analysis using both sentence embeddings and lexical measures.The corpus comprises 54,650 sentences totaling 501,394 words, distributed across six dialectal varieties: Jaffna (10,009 sentences, 70,367words),Malayagam(10,105; 86,562), Trincomalee (5,800; 54,617), Batticaloa (8,716; 74,864), Moor (10,020; 78,378), and Standard Tamil (10,000;136606). Cosine similarity analysis using TF-IDF and LaBSE embeddings reveals that Batticaloa, Malayagam, Moor, and Trincomalee form a closely related cluster (similarity range: 0.94–0.98), Standard Tamil is the most distinct, and Jaffna appears between the cluster and Standard Tamil. In addition, we constructed a set of 100 parallel sentences across all dialects with native speaker support, which were evaluated using LaBSE to examine sentence-level divergence. Visualization techniques such as heatmaps and network graphs were used to compare dialectal relationships and highlight clustering tendencies. The results confirm a dialect continuum within Sri Lankan Tamil and provide empirical data useful for automated dialect classification, resource development, and computational modeling of under-resourced regional varieties.

Keywords: Sri Lankan Tamil, Dialect Classification, Sentence Embeddings, Lexical Analysis, Semantic Similarity

Dialect Divergence In Sri Lankan Tamil - A Corpus-based Computational Analysis

Gobalakrishnar Nilusana and Kengatharaiyar Sarveswaran
Department of Computer Science,University of Jaffna
nilusana07gobalakrishnar@gmail.com

1. Introduction

- A dialect is defined as theregionally or locally based varieties of language [1].
- There are mainly five types of major geographic dialect variation in Sri Lankan Tamil:
  - Jaffna Tamil
  - Batticaloa Tamil
  - Malayaga Tamil
  - Moor Tamil
  - Trinco Tamil
- Dialect identification is essential for localization, because systems can respond more naturally when they recognize the user's dialect.
- No previous computational study exists for Sri Lankan Tamil dialects, making this the first dialect-based computational analysis of its kind.

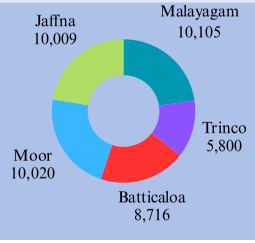
Objective

- Built a dialect classifier capable of identifying major geographic dialect variations of Sri Lankan Tamil using text data from multiple sources:

Direct contributions from native speakers representing each dialectal region,Shortstories from books

2.Dataset Creation

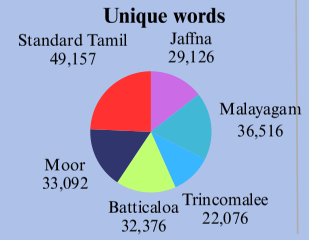
Sentence counts

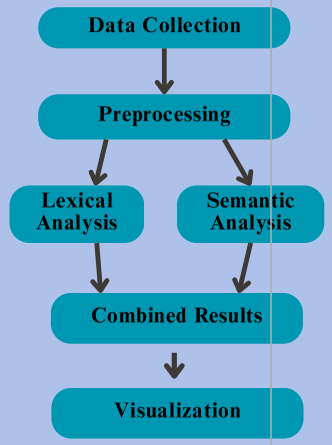


3. Methodology

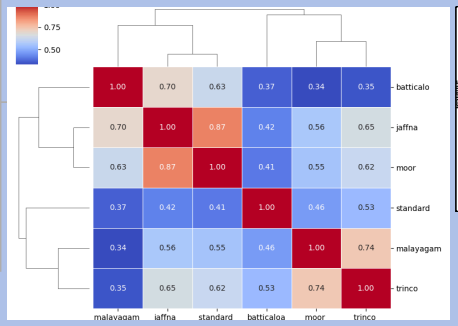
Data Collection

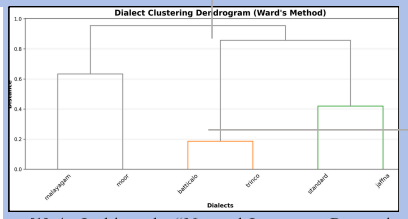
- Data was collected as text from the Noolaham Foundation.
- For the purpose of performing lexical analysis, we created unique tokens for each dialect.
- The main idea is to distinguish the differences between each dialect and Standard Tamil.





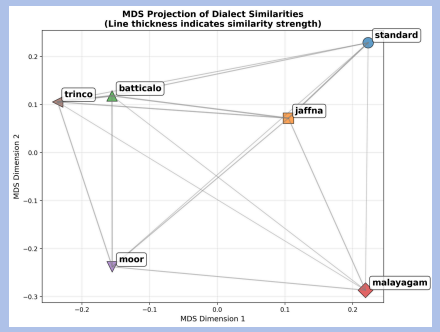
- **Lexical Analysis**: We constructed a vocabulary set for each dialect and examined the resulting lexical variation across the dialects.
- **Semantic Analysis**: We employed the **LaBSE** transformer model to generate normalized sentence embeddings for evaluating semantic similarity and divergence.





4.Analysis

1. **Jaffna as a Linguistic Bridge**
   - Positions Jaffna as a central, mediating variety in the dialect continuum
2. **Eastern & Up-Country Dialect Cluster**
   - Malayagam, Moor, Trincomalee, and Batticaloa form a highly cohesive group
3. **Standard Tamil's Unique Position**
   - Lexically distinct despite cosine similarity scores



5. Conclusion

Our analysis confirms that Sri Lankan Tamil dialects have diverged, each exhibiting distinct and consistent linguistic features.

6. Selected References

[1] A. Joshi et al., "Natural Language Processing for Dialects of a Language: A Survey," ACM Computing Surveys, vol. XX, no. Y, pp. ZZZ–AAA, Jan. 2024.
[2] K. Prasanna and A. Arora, "IruMozhi: Automatically Classifying Diglossia in Tamil," arXiv.org, [Online]. Available: arXiv:2301.07639. [Accessed: Nov. 7, 2025], 2023.
[3] S. Bonta, "Negombo Fishermen's Tamil: A Case of Indo-Aryan Contact-Induced Change in a Dravidian Dialect," Anthropological Linguistics, vol. 52, no. 3, pp. 310–343, 2010.

Poster ID: P35

Improving Label Consistency in Tamil Emotion Datasets through a Multi-Method Pipeline

Sankavi Mohan*, Andrew Charles Eugene Yugarajah

Department of Computer Science, University of Jaffna, Sri Lanka

*2020csc033@univ.jfn.ac.lk

Emotion recognition in text is an important task in Natural Language Processing (NLP). However, in low-resource languages such as Tamil, progress is limited due to the lack of reliable and consistently annotated datasets. Existing corpora often contain label noise, which reduces their usefulness for training robust models. This study presents a multi-method relabeling pipeline aimed at improving the quality of a publicly available Tamil emotion corpus by identifying misclassified samples and reassigning more consistent labels. The proposed workflow integrates four techniques: a Support Vector Machine (SVM) classifier with Indic-BERT embeddings, Improved Deep Embedded Clustering (IDEC), cross-lingual transfer learning with XLM-RoBERTa, and OpenAI GPT-4o mini. These approaches capture different perspectives – statistical classification boundaries, cluster-level structure, multilingual contextual understanding, and semantic reasoning. Predictions from all methods are combined using an ensemble majority-voting strategy to generate stable and high-confidence labels. Experimental results demonstrate clear improvements in dataset consistency. The baseline SVM achieved an accuracy of 0.37, XLM-RoBERTa fine-tuning reached 0.48, clustering (k = 32) obtained 0.76, and ensemble voting stabilized performance at 0.53. Overall, this pipeline provides a practical and scalable approach for improving Tamil emotion datasets and supports the development of more reliable emotion classification models in low-resource NLP settings.

Keywords: Emotion recognition, Data quality, Transformers, Generative AI, Clustering

Improving Label Consistency in Tamil Emotion Datasets through a Multi-Method Pipeline

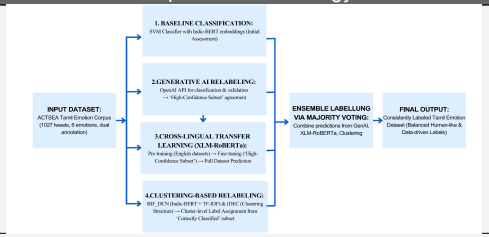
M. Sankavi and E.Y.A. Charles
Department of Computer Science, University of Jaffna
2020csc033@univ.jfn.ac.lk, charles.ey@univ.jfn.ac.lk

Introduction

Emotion recognition in text is widely used in Natural Language Processing (NLP). However, for low-resource languages like Tamil, progress is limited due to small datasets and inconsistent manual annotations. Label noise reduces model performance and weakens reproducibility.

The primary contribution of this work is the development of a hybrid, multi-method relabeling pipeline that significantly improves label consistency and establishes a more reliable foundation for Tamil emotion classification in low-resource settings.

Proposed Methodology



Dataset

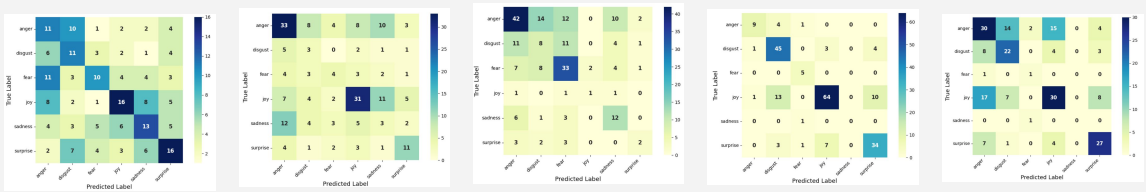
ACTSEA Dataset (MERCon 2019)

- 1,027 Tamil tweets
- Six emotions: anger, disgust, fear, joy, sadness, surprise
- Includes "correctly classified/misclassified" expert tags

	Anger	Disgust	Fear	Joy	Sadness	Surprise	Total
Correctly Classified	69	86	80	146	85	159	625
Misclassified	60	83	63	36	106	54	402
Total	129	169	143	182	191	213	1027

Experiment and Results

No	Method	Accuracy
1	SVM (Baseline)	0.37
2	Generative AI	0.41
3	XLM-RoBERTa Fine-tuned	048
4	IDEC Clustering (k=32)	0.76
5	Ensemble Voting	0.53



Conclusion

- Clustering revealed strong internal structures aligned with emotion categories.
- Generative AI improved semantic relabeling.
- XLM-RoBERTa brought cross-lingual knowledge to Tamil.
- Ensemble voting balanced strengths of all methods.
- **The proposed hybrid pipeline significantly improves label consistency in Tamil emotion datasets.**

References

[1] A. A. Maruf et al., "Challenges and Opportunities of Text-Based Emotion Detection: A Survey," *IEEE Access*, 2024.
[2] M. Garg and C. Saxena, "Emotion Detection from Text using Machine Learning," *Computational Intelligence Methods for Sentiment Analysis*, 2024.
[3] R. Tamilkodi et al., "Emotion Detection in Text: Advances in Sentiment Analysis," *Int. J. System Assurance Engineering & Management*, 2024.
[4] W. S. Ismail, "Emotion Detection in Text using Deep Learning," *Journal of Wireless Mobile Networks*, 2024.
[5] H. V. Khurdula et al., "Advanced Text-Based Emotion Detection using Bi-LSTM, BERT & Palm V2," *IEEE SoutheastCon*, 2024.
[6] M. S. U. Sourav et al., "Transformer-based Text Classification on Bangla Emotion Corpus," *arXiv preprint*, 2022.
[7] O. Gokhale et al., "Emotion Analysis in Tamil," *DravidianLangTech Workshop (ACL)*, 2022.
[8] P. Ram N. et al., "Multi-Class Emotion Classification on Tamil–Tulu Code-Mixed Text," *ICACCS*, 2024.
[9] V. O. Yenumulapalli et al., "Emotion Classification in Hindi–English Code-Mixed Dialogue," *SemEval*, 2024.
[10] R. Jenarthanan et al., "ACTSEA: Annotated Corpus for Tamil & Sinhala Emotion Analysis," *MERCon*, 2019.
[11] C. Vasantharajan et al., "TamilEmo: Fine-Grained Emotion Detection Dataset for Tamil," *SLTU–Low Resource Languages*, 2022.
[12] Q. Zhang et al., "BIF-DCN for Clustering and Sentiment Analysis," *PeerJ Computer Science*, 2025.

Modeling Classical Nahuatl Verbal Nuclear Clauses with a Finite State Transducer

Francesca Hoffmeier*, Natalia Villanueva

University of Konstanz, Germany

*francesca.hoffmeier@uni-konstanz.de

This project encodes Classical Nahuatl's rich morphology with the use of a Xerox finite state (XFST) programming language. We present a finite state morphological analyzer of the verbal nuclear clause (VNC) in Classical Nahuatl within the XFST environment. Following Andrew's (2003) analysis of Nahuatl as a polysynthetic language, whose basic syntactic units are "nuclear clauses", we implement the VNC as a single morphosyntactic string that obligatorily encodes subject and predicate information. The verbal predicate, which consists of a stem and suffixes, encodes the verb's transitivity, the stem's lexical information, and a tense specification, thereby being able to function as a full sentence. We organize the lexicon in order that verb stems branch into intransitive vs. transitive paths, with transitive stems further divided into monadic and dyadic valence. We use tags to distinguish object person/number and deploy flags to block valence morphology with intransitives, enforce agreement between Valence 1 and Valence 2, and maintain subject number agreement. Phonological post-lexical rules handle attested alternations. The coverage matches the portion of the paradigm, and the flag architecture prevents common overgeneration errors. The transducer licenses the VNC templates described in the reference grammars and blocks illicit combinations. The main limitation of this project lies in its restriction to a small set of verbs. Near-term work includes expanding tense/mood/aspect, enriching stem inventories, and extending the analyser to nominal nuclear clauses. Longer-term goals include gold-standard evaluation against annotated corpora and integration into pedagogical tools for Nahuatl language documentation and instruction.

Keywords: Nahuatl, verbal nuclear clause, finite-state morphology, XFST

Modeling Classical Nahuatl Verbal Nuclear Clauses with a Finite State Transducer

Francesca Hoffmeier[1] and Natalia Villanueva[2]

[1 2]*University of Konstanz*

Introduction

- Classical Nahuatl is a language indigenous to **México**
- Variations spoken by over 1 million speakers all over México
- **Polysynthetic language** that employs a larger number of morphemes to form "words"
- Andrews (2003) introduces the concept of "**nuclear clauses**" , morphosyntactic unit that contains both subject and predicate

Verbal Nuclear Clause							
Empty valence							
pers[1]	pers[2]	(STEM)		tns	num[1]	num[2]	
Monadic valence							
pers[1]	pers[2]	va	(STEM)		tns	num[1]	num[2]
Dyadic valence							
pers[1]	pers[2]	va[1]	va[2]	(STEM)	tns	num[1]	num[2]

VNC example → *timitztlazotla*							
t(i)	ø	m	itz	tlazotla	ø	ø	ø
+2sgS	+nomS	+Sg	+2P_Obj	+Verb	+prs	+cm	+sg

Reference: J. R. Andrews Rev.ed., Introduction to Classical Nahuatl, 2nd ed. Austin, TX: University of Oklahoma Press: Norman, 2003.

Methodology

- Xerox finite state (XFST) morphological analyzer
- Morphological analysis & generation
- Lexicon mirrors VNC template

Lexicon fragment

```
LEXICON Pers2
nomS+:0                 Intransitive;
nomS+:0                 Transitive;


LEXICON Transitive
@U.val.yes@             Monadic;
@U.val.yes@             Dyadic;


LEXICON Monadic
Spec+:ne                Stem;
Hum+:tE                 Stem;
NonHum+:tla             Stem;


LEXICON Dyadic
Sg+:0                   Person_Sg;
Pl+:0                   Person_Pl;

LEXICON Intransitive
@U.val.no@              Stem;
```

Post-lexical REGEX rules

```
n -> n i || .#. _ Con ?+
.o.
t -> t i || .#. _ Con ?+
.o.
t i -> t || _ l a ?+
```

Results and Discussion

- Transducer licenses the VNC templates and blocks illicit combinations
  - Intransitive stems bypass object morphology
  - Monadic vs. dyadic transitive stems admit the appropriate object person/number paradigms
  - Person-number agreement holds across circumfix
  - Flag architecture prevents overgeneration and enable agreement
  - Limited to a small set of verb stems

```
timitztlazotla
2sgS+nomS+Sg+2P_Obj+tlazotla+Verb+prs+cm+sg

tiquintlazotla
2sgS+nomS+Pl+3P_Obj+tlazotla+Verb+prs+cm+sg

nitEtlazotlani
1sgS+nomS+Hum+tlazotla+Verb+cprs+cm+sg
```

Conclusion

A finite state architecture captures the core morphology of Classical Nahuatl when coupled with targeted feature tags, agreement-enforcing flags, and a small set of phonological rules.

- Long-term goals:
  - Expand tense/mood/aspect/verb stems
  - Extend to nominal nuclear clauses
  - Gold-standard evaluation against annotated corpora
  - Integration into pedagogical tools for Nahuatl language documentation and instruction.

Perplexity-Based Evaluation of Open-Source LLM Adaptation for Tamil

Jeyarajalingam Varsha*, Tharmaraj Thayatharsan, Kengatharaiyer Sarveswaran

Department of Computer Science, University of Jaffna, Sri Lanka

*varshajeyaraj@gmail.com

This study examines how multilingual large language models can be adapted to Tamil, which remains a low-resource language despite its long literary heritage and limited availability of high-quality Unicode text. Three open-weight models, Llama-2-7B, Qwen2-7B, and Mistral-7B, were evaluated using the Unsloth optimisation framework, with the 7B size selected as the largest model that could be trained within the available computational resources. Attempts to use larger models required A100 GPUs and resulted in rapid credit usage, making them unaffordable. Continued pretraining was carried out on 25% of the CulturaX-ta dataset (827,480 tokens) using Google Colab Pro with T4 and L4 GPUs in 4-bit quantisation. Perplexity was measured on 5,000 held-out tokens before and after pretraining, as it directly reflects predictive uncertainty and is a widely used intrinsic evaluation method that does not require task-specific fine-tuning. The results show that Llama-2-7B and Qwen2-7B achieve clear reductions in perplexity, indicating effective adaptation to Tamil even with limited data and modest hardware. Qwen2-7B obtained the best performance, while Llama-2-7B showed the significant relative improvement. Mistral-7B, however, recorded a degradation in performance, possibly due to catastrophic forgetting when exposed to unfamiliar data. Overall, the findings showed that multilingual models can be adapted efficiently to low-resource languages using lightweight methods such as Unsloth pipeline, highlighting the importance of model choice, multilingual coverage, and pretraining history. Future work will explore larger datasets, additional models, and alternative quantisation or training strategies.

Keywords: Perplexity, Language Models, Tamil, Continued Pretraining, Low-resource languages

Perplexity-based Evaluation of Open-Source LLM Adaptation for Tamil

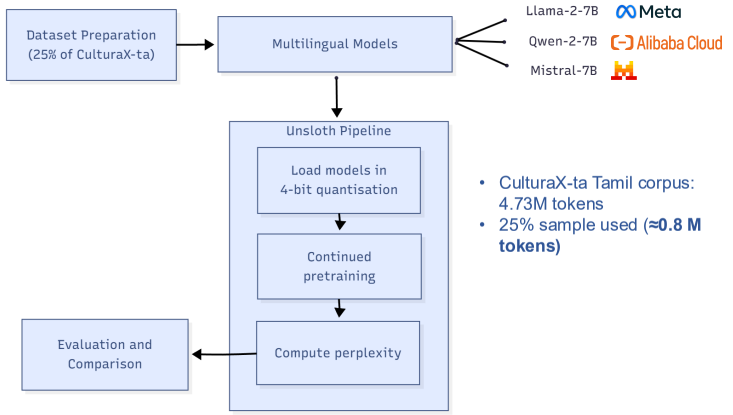
Jeyarajalingam **Varsha**, Tharmaraj **Thayatharsan**, Kengatharaiyer **Sarveswaran**
Department of Computer Science, University of Jaffna.
varshajeyaraj@gmail.com

Scan to read the full paper.

1. Introduction

- Tamil remains a low-resource language due to limited digitised text and the continued use of legacy, non-Unicode fonts.
- Many multilingual LLMs lack sufficient Tamil exposure in their pretraining corpora [1].
- Despite recent progress, meaningful and reliable Tamil generation is still limited [1].
- In this study, Llama-2-7B, Qwen2-7B and Mistral-7B are adapted using the Unsloth framework under limited data and constrained compute [2,3,4].
- Perplexity is used to assess predictive confidence; lower scores indicate better next-token prediction [2].
- **Objective: Identify an open-weight model that adapts efficiently to Tamil using modest data and hardware.**

2. Methodology



- CulturaX-ta Tamil corpus: 4.73M tokens
- 25% sample used (**≈0.8 M tokens**)

3. Results and Discussion

Model	Baseline	After
Llama-2-7B	121.20	54.30
Qwen-2-7B	62.49	47.03
Mistral-7B	50.81	67.34

Table 1: Perplexity scores of models before and after continued pretraining

- Continued pretraining improved performance for Llama-2-7B and Qwen2-7B.
- Llama-2-7B showed the significant improvement, while Qwen2-7B achieved the best final score.

Results support that large multilingual models trained on broad corpora can still be efficiently adapted to low-resource languages with limited data and hardware.

4. Conclusion

- Unsloth-based continued pretraining improved perplexity score for Llama-2-7B and Qwen2-7B, showing effective Tamil adaptation even with limited data and computational power.
- Mistral-7B showed degraded performance, likely due to catastrophic forgetting when exposed to unfamiliar data.
- Future work will explore larger datasets and alternative training methods to enhance performance.

5. Selected References

[1] A. Balachandran, "Tamil-LLaMA: A new Tamil language model based on LLaMA 2," *arXiv preprint* arXiv:2311.05845, Nov. 2023. [Online]. Available: https://arxiv.org/abs/2311.05845

[2] "Shared Task – ICIIT Conclave 2025: Building Compact Sinhala & Tamil LLMs." [Online]. Available: https://iciit.iit.ac.lk/conclave/shared-task/

[3] Unsloth AI, "Continued Pretraining." [Online]. Available: https://docs.unsloth.ai/basics/continued-pretraining

[4] K. Singhapoo, A. Inthanil, and A. Pillai, "Fine-tuning AI models with limited resources," in *Proc. 11th Int. Conf. Eng., Appl. Sci., Technol. (ICEAST)*, 2025, pp. 148–151, doi: 10.1109/ICEAST64767.2025.11088202.

Poster ID: P38

Target Identification in Tamil Hate Speech

Sathasivam Nerujan*, Kengatharaiyar Sarveswaran

Department of Computer Science, University of Jaffna, Sri Lanka

*sathasivamnerujan35@gmail.com

Hate speech detection identifies harmful language online to create safer digital spaces. Target identification determines whether the hate is aimed at individuals, organizations, or communities, which is crucial to prevent emotional harm and social or economic impact. For Tamil, it bridges a significant linguistic gap and improves digital safety for Tamil-speaking users. Using Named Entity Recognition (NER), hate speech targets can be automatically detected and classified for precise moderation. This research introduces and validates a novel, multi-stage pipeline for automated target identification in the low-resource Tamil language. The pipeline employs customized models across three phases: first, MuRIL is fine-tuned for hate speech detection on a dataset of 12,396 sentences, equally split into hate and non-hate. Second, XLM-RoBERTa is adapted for NER with modified entity labels (PER, ORG, COMM). Third, the extracted entities are classified as targets using a binary 0/1 scheme. Final results demonstrate the efficacy of this approach: MuRIL achieved 0.73 accuracy for hate speech detection, outperforming the traditional 0.71 TF-IDF baseline. Crucially, the final target identification accuracy reached 0.72, superior to the upstream NER Macro F1-score of 0.69 and competitive with related literature benchmarks (0.68 F1-score). This validates that MuRIL's focused binary intent classification is highly effective at refining extracted entities. As a future direction, we plan to leverage larger annotated datasets and lightweight architectural improvements to further enhance model robustness.

Keywords: Hate speech, Target identification, Tamil NLP, MuRIL, XLM-RoBERTa, TF-IDF

Target Identification In Tamil Hate Speech

Sathasivam Nerujan, Kengatharaiyar Sarveswaran
Department of Computer Science, University of Jaffna, Sri Lanka
sathasivamnerujan35@gmail.com

SLSWCS Student Workshop on Computer Science

Introduction

- Hate speech is a growing online harm that threatens the dignity and safety of **individuals, organizations,** and **communities**.
- **Identifying who the hate targets** are not just detecting hate speech is essential for protecting these groups effectively.
- **Tamil's low-resource** context lacks tools for identifying hate targets, creating a critical gap in digital safety.
- We address this challenge by developing a **multi-stage pipeline** that detects hate speech and extracts its targets using MuRIL and XLM-RoBERTa.

Methodology

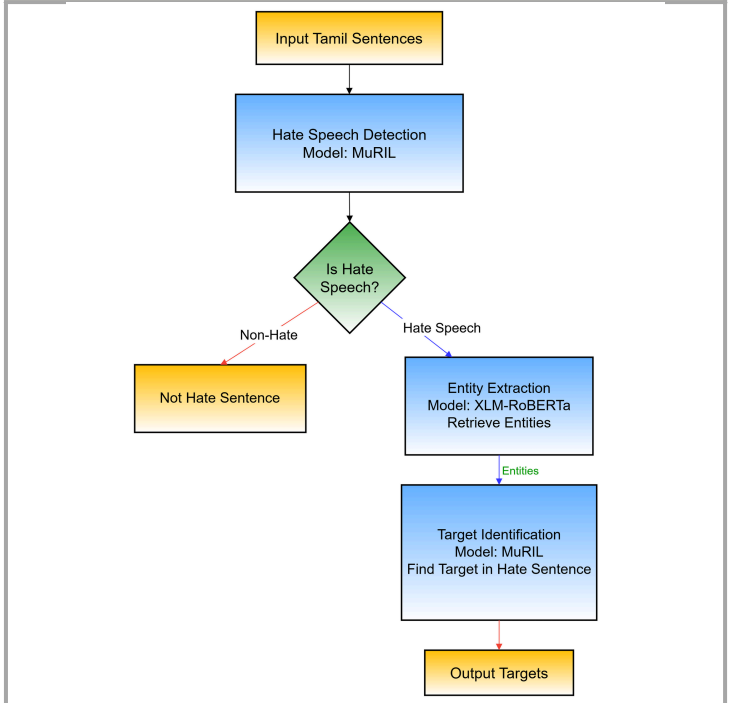


Figure: proposed methodology

Results

Table: experimental result

Size	Tasks	Dataset Source	Results	Model used
12,394 Sentences	Hate Detection	FIRE2020 & Dravidian CodeMix	0.73 Accuracy	MuRIL
2,070 Sentences	Entity Recognition	Custom Multi-way Parallel Corpus	0.68 F1-score	XLM-RoBERTa
600 Sentences	Target Identification	Manually Annotated Corpus	0.72 Accuracy	MuRIL

Challenges

- **Linguistic & Code-Mixed Complexity**: Handling Tamil-English text and Tanglish text and cultural nuances that standard models often miss.
- **Limited Resources**: Overcoming the scarcity of annotated Tamil datasets for hate speech detection.
- **Contextual Ambiguity**: Distinguishing between neutral mentions of communities and actual hate speech targets.

Discussion

- MuRIL (0.73 accuracy) outperformed the TF-IDF baseline (0.71) and matched state-of-the-art Tamil hate speech results.
- Target Identification: 0.72 accuracy, surpassing upstream NER (0.68) and ensemble benchmarks (0.68 F1).
- Binary classification with MuRIL improved entity extraction and more reliably identified malicious intent.

Conclusion

- Our pipeline improves Tamil hate speech target detection beyond NER alone.
- Binary intent classification enhances accuracy over pure entity extraction.
- MuRIL + XLM-RoBERTa offers a stronger approach for Tamil target linking.

Selected References

[1] S. Swaminathan, H. Ganesan, and R. Pandiyarajan, "HRS-TECHIE@Dravidian-CodeMix and HASOC-FIRE2020: Sentiment analysis and hate speech identification using machine learning, deep learning and ensemble models," 2020.
[2] S. Ranathunga, A. Ranasinghea, J. Shamala, A. Dandeniya, R. Galappaththi, and M. Samaraweera, "A Multi-way Parallel Named Entity Annotated Corpus for English, Tamil and Sinhala," *arXiv preprint arXiv:2412.02056*, 2024.
[3] A. Ganguly, A. N. B. Emran, S. S. C. Puspo, M. N. Raihan, D. Goswami, and M. Zampieri, "MasonPerplexity at Multimodal Hate Speech Event Detection 2024: Hate speech and target detection using transformer ensembles," *arXiv preprint arXiv:2402.01967*, 2024.
[4] B. R. Chakravarthi et al., "Detecting caste and migration hate speech in low-resource Tamil language," *Language Resources & Evaluation*, vol. 59, pp. 3051–3086, 2025, doi: 10.1007/s10579-025-09848-x.

Web Technologies

A Performance Comparison Between Rest API and GraphQL in Microservice Architecture

Thulashika Balachandran*, Kengatharaiyar Sarveswaran, Samantha Tharani Jeyakumar

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*lasibala24@gmail.com

Microservice architectures are increasingly adopted for building scalable and maintainable software systems, where selecting an efficient communication protocol is crucial for achieving optimal performance. REST and GraphQL are widely used in microservices, yet most existing studies evaluate them primarily through GET requests, offering limited understanding of their behaviour in data-creation scenarios. This study addresses this gap by examining POST operations, which more accurately reflect real-world transactional workloads. Using an e-commerce application composed of four microservices connected to a MongoDB database, REST and GraphQL endpoints were implemented and tested under simulated concurrent workloads ranging from 100 to 500 users. Apache JMeter was used to measure response time and throughput, which were chosen as primary metrics because they directly represent API performance, while CPU usage was excluded due to its dependency on hardware and runtime variability. The results show that REST consistently outperforms GraphQL in both metrics, recording significantly lower response times and higher throughput across all concurrency levels. GraphQL performs competitively under lighter loads but experiences a steeper increase in latency as concurrency rises, indicating that its flexible query execution introduces additional processing overhead during write-heavy operations. These findings suggest that REST is more suitable for systems requiring fast, high-volume request handling, whereas GraphQL may still be advantageous in applications prioritising flexible data retrieval despite potential performance trade-offs. The study provides practical guidance for developers and architects in selecting API protocols for microservices and highlights opportunities for future research involving broader workloads, hybrid API strategies, and extended performance indicators.

Keywords: Microservices, API, REST, GraphQL, Performance Testing, JMeter

A Performance Comparison Between REST API and GraphQL in Microservice Architecture

Thulashika Balachandran, Kengatharaiyar Sarveswaran and **Samantha Tharani Jeyakumar**

Department of Computer Science, Faculty of Science, University of Jaffna

lasibala24@gmail.com

Abstract

Microservices rely on efficient communication protocols for handling real-world workloads. This study compares REST and GraphQL using POST operations in an e-commerce system built with four microservices and MongoDB. Performance was evaluated under 100–500 concurrent users using Apache JMeter, focusing on response time and throughput. Results show that REST consistently delivers lower latency and higher throughput than GraphQL, which performs well under light loads but degrades as concurrency increases. The findings highlight REST as the more suitable choice for high-volume transactional workloads, while GraphQL remains useful when flexible data retrieval is prioritised.

Introduction

Microservices split applications into independent services, making communication efficiency critical. REST uses multiple endpoints, while GraphQL provides a single flexible endpoint. This study evaluates their performance in terms of response time and throughput for flat and nested data under varying user loads and offers practical recommendations for protocol selection.

Objective

- Measure and evaluate the performance of REST and GraphQL
- Analyze the comparative efficiency of both communication styles
- Provide evidence-based recommendations

Methodology

1. Microservice Architecture Setup
   - Independent services implemented using a microservice architecture.
   - Each service exposes parallel REST and GraphQL endpoints with identical CRUD operations.
   - Services interact with a shared MongoDB database under consistent schema and indexing.
2. Request Handling and Logging
   - Client requests are directed simultaneously to REST and GraphQL API gateways to ensure identical workloads.
   - Requests are logged at gateway and service levels, capturing timestamps for precise response time and throughput measurement.
3. REST and GraphQL Implementation
   - REST endpoints follow standard HTTP methods (POST, GET, PUT, DELETE).
   - GraphQL endpoints follow a schema-first approach, allowing flexible client queries.
4. Database Interaction
   - Both protocols use identical data access logic and collections in MongoDB.

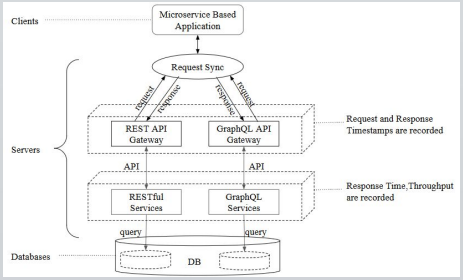


Figure 1: An Illustration of the evaluated system architecture.

Experiment

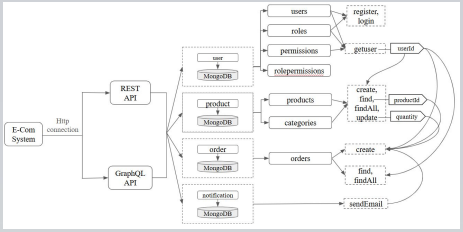


Figure 2: Microservices architecture for an E-Commerce system using REST API and GraphQL

Performance Testing Setup

- POST operations were tested using Apache JMeter to simulate concurrent users.
- Request loads: 100, 200, 300, 400, 500
- Metrics recorded: Response Time(sec) and Throughput (requests/sec).
- Each test was repeated 5 times, averages and standard deviations were computed.
- A t-test was used to validate the statistical significance of performance differences.

Action	Required Permission	Flow of Communication
Place Order	manage_orders	User (checkPermission) → Product (verify stock) → User (fetch info) → Order (create) → Product (update quantity) → Notification (send email)

Table 1: Microservice Interaction Flow for the Place Order Operation

Literature Review

- The REST performs faster than GraphQL in terms of response time [3,5].
- gRPC often outperforms both REST & GraphQL, especially in high-load microservice systems [4].
- GraphQL improves data efficiency, reducing over-fetching, and allows flexible querying [7].
- CPU usage tends to be higher for GraphQL due to query parsing overhead [8].

Existing studies mainly use GET requests and rarely explore MongoDB-based microservice environments.

Results

GraphQL shows slightly slower response time due to query parsing and nested field resolution. Throughput increases with higher request volumes, with REST demonstrating better scalability under load. Figures 3 and 4 present the response time and throughput measurements for both approaches.

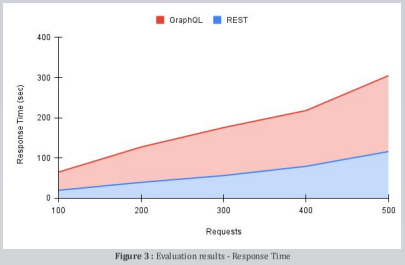


Figure 3 : Evaluation results - Response Time

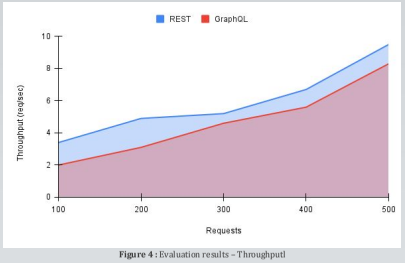


Figure 4 : Evaluation results – Throughputl

Discussion & Conclusion

- The REST outperformed GraphQL in both response time and throughput across all tested scales.
- GraphQL offers greater flexibility in data retrieval, its additional processing overhead leads to slower performance in microservice environments.
- REST is more efficient and scalable for high-load microservice communication in this system.
- GraphQL may be preferred when flexible, fine-grained data access is more important than raw performance.

Future Work

- Comparison by integrating gRPC under the same microservice architecture to evaluate all three protocols.
- Conduct experiments using real-time production traffic instead of simulated JMeter loads for more realistic performance insights.

References

[1] L. Costa and A. N. Ribeiro, "Performance evaluation of microservices featuring different implementation patterns," in International Conference on Intelligent Systems Design and Applications. Springer, 2021, pp. 165–176.

[2] Ramu, V. (2023). Performance Impact of Microservices Architecture. Rev. Contemp. Sci. Acad. Stud, 3, 1-6.

[3] . Lawi, B. L. Panggabean, and T. Yoshida, "Evaluating graphql and rest api services performance in a massive and intensive accessible information system," Computers, vol. 10, no. 11, p. 138, 2021.

[4] M.Niswar, R. A. Safruddin, A. Bustamin, and I. Aswad, "Performance evaluation of microservices communication with rest, graphql, and grpc," International Journal of Electronics and Telecommunications, vol. 70, no. 02, pp. 429–436, 2024.

[5] N. Vohra and I. B. K. Manuaba, "Implementation of rest api vs graphql in microservice architecture," In 2022 International Conference on Information Management and Technology (ICIMTech), pp. 45–50, 2022.

[6] S. L. Vadlamani, B. Emdon, J. Arts, and O. Baysal, "Can graphql replace rest? a study of their efficiency and viability," in 2021 IEEE/ACM 8th International Workshop on Software Engineering Research and Industrial Practice (SER&IP). IEEE, 2021, pp. 10–17.

[7] G. S. M. Diyasa, G. S. Budiwitjaksono, H. A. Ma'rufi, and I. A. W. Sampurno, "Comparative analysis of rest and graphql technology on nodejs-based api development," Nusantara Science and Technology Proceedings, pp. 43–52, 2021.

[8] E. Frigård, "GraphQL vs. REST: A comparison of runtime performance," 2022.

CertiRide: A Tamper-Free Vehicle Service Record Ledger

Naveen Hettiwaththa*, Samantha Jeyakumar, Haritha Senadheera, Hiroshan Hettiarachchi, Sandaru Gunarathna, Tharushi Adhikari

Department of Computer Science, Faculty of Science, University of Jaffna, Sri Lanka

*naveenhettiwaththa@gmail.com

Centralised maintenance and service record systems in Sri Lanka suffer from inadequate preventive maintenance practices, inaccurate vehicle records, manual data handling, and delays in repair approvals. These systemic issues result in higher operational costs, reduced vehicle reliability, and inefficient fleet management across government institutions. Addressing these challenges, this research proposes CertiRide, a decentralised vehicle maintenance and service record system built on the Hyperledger Fabric private blockchain. By leveraging blockchain technology, CertiRide ensures an immutable record of vehicle maintenance histories while preserving user privacy through pseudonymity. Additionally, the integration of AI-powered tools: including OpenAI models and Azure Document Intelligence, enables automated extraction and validation of structured data from scanned documents related to both the vehicle and its owner. The components of CertiRide support secure digital storage, management, and retrieval of tamper-proof service records and vehicle condition information. This system significantly mitigates the limitations of traditional approaches and enhances trust among buyers, sellers, and service providers in the aftermarket vehicle industry.

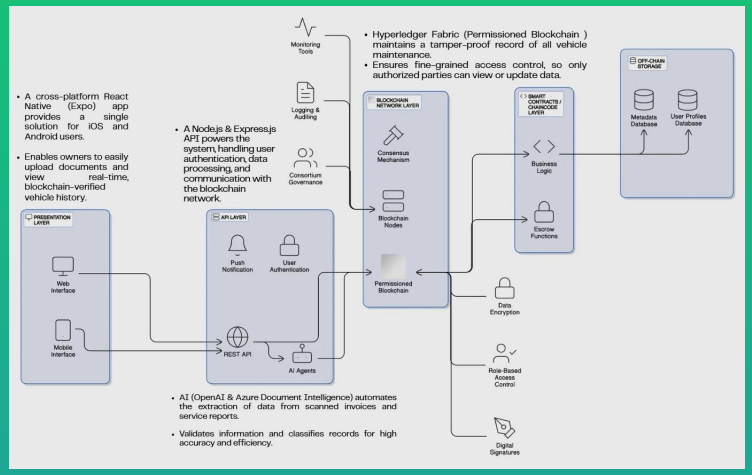
Keywords: Tamper-proof vehicle record, Decentralised vehicle maintenance, Hyperledger, OpenAI

CertiRide: A Tamper-Free Vehicle Service Record Ledger

Naveen Hettiwaththa, Haritha Senadheera, Sandaru Gunarathna, Hiroshan Hettiarachchi, Tharushi Adhikari, Samantha Tharani Jeyakumar
naveenhettiwaththa@gmail.com
Department of Computer Science, Faculty of Science, University of Jaffna

Introduction

- Paper-based vehicle maintenance records are prone to loss, tampering, and limited accessibility.
- **CertiRide** is a blockchain-powered Vehicle Maintenance and Service Record System
- Designed to offer transparent, tamper-proof, and easily retrievable vehicle histories.
- Provides digitizing service records and storing them on a distributed ledger.
- Enhance trust in the vehicle history, reducing fraud and improving long-term vehicle reliability.

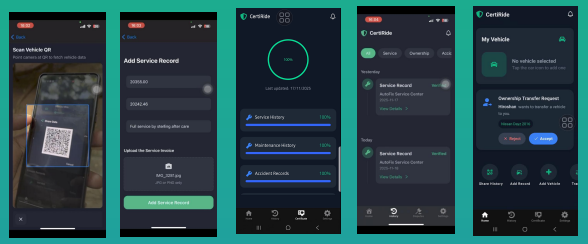


Methodology

- *Architecture:* Node.js, Prisma ORM, and blockchain integration for secure record storage.
- *Data:* Vehicle profiles, service details, mileage, workshop information, and documents are collected from verified users.
- *Blockchain Layer:* Each service event is hashed and anchored to a blockchain network to ensure immutability and integrity.
- *API & Backend:* A RESTful API handles certificate generation, user authentication, and CRUD operations.
- *Authentication:* Generated certificates include a QR code allowing buyers and stakeholders to instantly validate authenticity.
- *User Interface*: Clean, responsive dashboards allow owners and workshops to manage records seamlessly.

Results and Discussion

- **CertiRide** significantly enhances the trustworthiness of vehicle service histories.
- Integration of blockchain eliminates unauthorized alterations and supports secure data sharing across workshops and owners.
- Certificate generation feature further promotes transparency by offering a portable, verifiable document summarizing the vehicle's complete maintenance lifecycle.



Scan to download android apps

Conclusion

CertiRide demonstrates how blockchain technology can modernize vehicle record management by ensuring permanent, verifiable, and transparent maintenance histories. The system effectively addresses issues in traditional documentation and opens a pathway to more reliable second-hand markets, informed decision-making, and improved long-term vehicle care.

*We extend our gratitude to our supervisor, collaborators, and the participating vehicle owners and workshops whose feedback shaped the development of **CertiRide**.*

Sri Lanka Student Workshop on Computer Science (SL-SWCS'25)

List of Reviewers

Prof. (Mrs.) Amalka Pinidiyaarachchi
University of Peradeniya, Sri Lanka

Prof. Amirthalingam Ramanan
University of Jaffna, Sri Lanka

Ms. Ann Sinthusha A. V.
University of Vavuniya, Sri Lanka

Dr. (Mrs.) Barathy Mayurathan
University of Jaffna, Sri Lanka

Dr. Damayanthi K. Herath
University of Peradeniya, Sri Lanka

Dr. E. Y. Andrew Charles
University of Jaffna, Sri Lanka

Prof. Gihan Dias
University of Moratuwa, Sri Lanka

Dr. Jabed M. Chowdhury
La Trobe University, Australia

Dr. (Mrs.) Jananie Jarachanthan
University of Jaffna, Sri Lanka

Dr. Kathiravelu Thabotharan
University of Jaffna, Sri Lanka

Dr. Kengatharaiyer Sarveswaran
University of Jaffna, Sri Lanka

Prof. Kitsuchart Pasupa
King Mongkut's Institute of Technology Ladkrabang, Thailand

Prof. (Mrs.) Maheshi B. Dissanayake
University of Peradeniya, Sri Lanka

Prof. Mahesan Niranjan
University of Southampton, England

Dr. Manjusri Wickramasinghe
University of Colombo, Sri Lanka

Dr. Nandakishor Desai
The University of Melbourne, Australia

Dr. (Ms.) Nirthika Rajendran
University of Jaffna, Sri Lanka

Mr. Ratnasingam Sakuntharaj
University of Vavuniya, Sri Lanka

Mrs. Sajeetha Thavareesan
Eastern University, Sri Lanka

Prof. Saman Halgamuge
University of Melbourne, Australia

Dr. Shalini Banerjee
Karlsruhe Institute of Technology, Germany

Dr. Sinnathamby Mahesan
University of Jaffna, Sri Lanka

Mr. Sittampalam Sotheeswaran
Eastern University, Sri Lanka

Prof. Siyamalan Manivannan
University of Jaffna, Sri Lanka

Mr. Somaskandan Suthakar
University of Jaffna, Sri Lanka

Dr. Sriskandarajah Shriparen
University of Jaffna, Sri Lanka

Dr. Sugandima M. Vidanagamachchi
University of Ruhuna, Sri Lanka

Dr. Thanikasalam Kokul
University of Jaffna, Sri Lanka

Dr. Tharmarajah Thiruvaran
University of Jaffna, Sri Lanka

Mr. Thirukumaran Sanmugasundaram
University of Vavuniya, Sri Lanka

Dr. Thirusittampalam Ketheesan
University of Jaffna, Sri Lanka

Dr. Uthayasanker Thayasivam
University of Moratuwa, Sri Lanka

Mr. Veerapathirapillai Vinoharan
University of Vavuniya, Sri Lanka

Dr. Vishwas Patil
Indian Institute of Technology Bombay, India

Mrs. Yasotha Ram Ramanan
University of Vavuniya, Sri Lanka

Organising Committee

General Chair	Dr. (Mrs). B. Mayurathan
Programme Chairs	Dr. E. Y. A. Charles
	Dr. K. Sarveswaran
	Dr. (Ms.) R. Nirthika
Finance Chairs	Dr. (Ms.) J. Samantha Tharani
	Mrs. N. Kesavi
Session Chairs	Dr. K. Thabotharan
	Prof. A. Ramanan
	Dr. S. Shriparen
Editors	Prof. M. Siyamalan
	Mr. S. Suthakar
Web & Publicity Chairs	Dr. S. Shriparen
	Mr. V. Visithan
	Mr. N.Thileepan
Logistic Chair	Dr. T. Kokul

Achievers of Past SL-SWCS

SL-SWCS'23

Best Presentation	**Mrs. Janani Kugarajeevan** *Transformer Tracking Using Spatiotemporal Features*
Best Poster	**Mr. Janith Somaweera** *Car Damage Severity Classification with Convolutional Neural Networks*

SL-SWCS'21

Best Presentation (Joint)	**Ms. Nirthika Rajendran** *Deep Learning for Diabetic Retinopathy Grading*
	Mrs. Shajini Majuran *A Landmark-free Attention Network for Fashion Clothes Classification and Attribute prediction*
Best Poster (Joint)	**Ms. M. Mayuravaani** *Classification of Surface Defects using Semi-Supervised Deep Learning*
	Ms. R. Priyadarshani *Underwater Image Enhancement using Convolution Neural Network*

SL-SWCS'19

Best Presentation (Joint)	**Ms. Nirthika Rajendran** *Is Soft Pooling Better than Max and Average Pooling? A Comparative Study on HEp-2 Cells and Retinal Image Classification Tasks*
	Mrs. Shajini Majuran *A Multiscale Contextual Technique for Fashion Clothes Landmark Localisation*
Best Poster (Joint)	**Ms. Kasthuriarachchi, T.C** *Deep Learning Approach to Detect Plagiarism in Sinhala Text*
	Ms. S. Meruja *Analysis of Methods to Handle Medical Sensor Data Towards Health Disorder Identification*

SL-SWCS'17

Best Presentation	**Mrs. Shajini Majuran** *A Feature-Driven Hierarchical Classification Approach to Emotions in Speeches using SVMs*
Best Poster	**Mr. Rajendran Miller** *A psychological based analysis of Marketing Email Subject Lines*

SL-SWCS'15

Best Presentation	**Mr. T. Kokul** *Multi-Person Tracking Using Generic Person Detector and Particle Filter*
Best Poster	**Ms. S. Charini, Ms. K. Anusha and Ms. W.M.B.K. Weerasooriya** *Automated Leaf Recognition System for Plant Classification*

SL-SWCS'13

Best Presentation	**Mr. K. Kanarupan**
Best Poster	**Mr. N. Nilashan, Mr .S. Thenuzan, Mr. S. Mayooran and Ms. D. Lukshica**

SL-SWCS'11

Best Presentation	**Ms. Yawwani Gunawardana** *Docking Studies and High Throughput Virtual Screening for Breast Cancer Drug Discovery using a Distributed Environment*
Best Poster	**Mr. Sivamayam Sivasuthan** *A Hybrid Reactive Routing Protocol for MANETs*

Sponsors

We sincerely thank our sponsors for their generous support of the SL–SWCS'25. Your continued commitment helps us recognise student excellence and create meaningful opportunities for young researchers across the country. The "Best Presentation" and "Best Poster" prizes of SL-SWCS'25 are sponsored by Prof. Saman K. Halgamuge (University of Melbourne, Australia) and Dr. Malka N. Halgamuge (RMIT University, Australia). Your support plays a vital role in making this workshop a success, and we are deeply grateful for your partnership.

Platinum Sponsor



Gold Sponsor



Silver Sponsors





Bronze Sponsors





Thank you for being a part of

SL-SWCS'25

https://www.slswcs.jfn.ac.lk

@ SL-SWCS'25 All rights reserved

Platinum Sponsor

Gold Sponsor

Silver Sponsors

Bronze Sponsors