



Matching RFID and image-based vehicle identification data to detect illegal road use

MS Fiorita

 **orcid.org/ 0000-0001-6810-0460**

Dissertation accepted in fulfilment of the requirements for
the degree *Master of Engineering in Computer and
Electronic Engineering* at the North-West University

Supervisor: Prof AJ Hoffman

Graduation: June 2025

Abstract

Effective road law enforcement is a vital part of transportation systems. Escalating traffic congestion and widespread illegal road usage pose fresh hurdles for law enforcement systems. An intelligent law enforcement system that aims to reduce illegal road use must enable immediate intervention on a tactical level without disrupting the flow of legal vehicles. To achieve a fast tactical response, strategic-level information, such as location and identity, must ensure law enforcement officials are in the right place at the right time. This thesis presents a system that equips authorities with the correct information to act against illegal road users while maintaining normal traffic flow. The proposed system improves existing solutions by combining Radio Frequency Identification (RFID) and Computer Vision (CV) data to detect and identify illegal vehicles. The authorities can leverage this information to remove such vehicles from the road without affecting the legal road user's experience. This thesis will investigate the performance criteria for such a system within a real-world environment and will determine which combination of data fields stored on the license plate tag can enable accurate vehicle identification while still being readable within a practical setup. Practical measurements will verify if the theoretical requirements can be achieved in practice.

Keywords: *Intelligent Transportation Systems, RFID, In-road detection, Computer Vision, ANPR.*

Acknowledgements

I extend my deepest gratitude to everyone who contributed to the completion of this dissertation. This journey has been both challenging and rewarding.

First and foremost, I express my deepest gratitude to the Lord Almighty. In times of challenge and triumph, I am profoundly thankful for the guidance, strength, and blessings bestowed upon me throughout this academic journey. It is with a humble heart that I acknowledge the divine grace that has illuminated my path and sustained me in moments of uncertainty.

Secondly, I express my sincere appreciation to my study leader, Prof. Alwyn Hoffman. Thank you for your invaluable guidance, unwavering support, and insightful feedback. Your expertise and commitment have been instrumental in shaping the trajectory of this research.

Special thanks to my family and my wife for their unwavering support and understanding during this academic endeavor. Their encouragement provided the motivation needed to persevere through the challenges.

Table of Contents

Abstract.....	i
Acknowledgements.....	iii
Table of Contents.....	iv
List of Figures	vii
List of Tables	ix
Abbreviations.....	xii
1 Problem Statement	1
1.1 Introduction	1
1.2 Problem definition	2
1.3 Proposed solution	3
1.4 Project aim and objectives.....	3
1.4.1 Primary objective	3
1.4.2 Secondary objective.....	4
1.4.3 Research questions	4
1.5 Anticipated benefits.....	4
1.6 Research Scope	5
1.7 Deliverables.....	5
1.8 Research Methodology and Design	6
1.9 Ethical Considerations.....	7
2 Literature Review	8
2.1 Electronic Vehicle Identification	8
2.1.1 Computer vision systems	8
2.1.2 Applications of Computer Vision in EVI systems	9
2.1.3 RFID systems	12
2.1.4 Applications of RFID in EVI systems	13
2.2 Computer Vision for Intelligent Transportation Systems	15
2.2.1 Visual Feature extraction	15
2.2.2 Deep learning.....	16
2.2.3 Factors affecting the Computer Vision EVI system	28
2.3 RFID for Intelligent Transportation Systems	35
2.3.1 Active or Passive RFID	35
2.3.2 Tagged vehicles	35
2.3.3 RFID readers.....	38
2.3.4 Factors affecting the EVI RFID system	42

2.3.5	Tag information.....	49
2.3.6	Market acceptance	51
2.3.7	Combining RFID and Computer vision for EVI	53
2.4	Chapter conclusion	54
3	Preliminary System Design	55
3.1	System requirements.....	55
3.1.1	Functional requirements.....	56
3.1.2	Performance requirements.....	56
3.1.3	Physical requirements.....	56
3.1.4	Interface requirements	57
3.1.5	Environmental requirements.....	57
3.1.6	Design requirements.....	57
3.1.7	Legal requirements	57
3.2	Context diagram.....	57
3.3	Functional flow.....	58
3.3.1	Level 0	58
3.3.2	Level 1	59
3.3.3	Level 2	60
3.4	System architecture	61
3.5	Component Selection.....	62
3.5.1	Edge controller	63
3.5.2	RFID reader	63
3.5.3	RFID tag.....	65
3.5.4	RFID reader antenna	67
3.5.5	CV Camera.....	71
3.6	Chapter conclusion	73
4	Initial RFID testing and experimentation.....	74
4.1	Testing Conditions.....	74
4.1.1	Tags to be tested.....	74
4.1.2	Tag initialization	74
4.1.3	Test environment.....	74
4.2	Waveform	74
4.3	Protocol.....	74
4.4	Frequency.....	75
4.5	Physical Setup	75

4.6	Transmit Power	75
4.7	Experimental equipment	76
4.8	Test methods	77
4.8.1	Tag Read Range.....	77
4.8.2	Tag orientation tolerance	82
4.8.3	Tag spacing from metal background.....	86
4.8.4	In-road reader antenna radiation pattern (tag facing reader)	89
4.8.5	RFID signal Illumination area	93
4.9	Chapter conclusion	97
5	Detail Design.....	99
5.1	System architecture	99
5.2	Hardware	100
5.2.1	RFID system hardware	100
5.2.2	Camera system hardware	103
5.3	Software	104
5.3.1	RFID system software	106
5.3.2	Digital signature	106
5.3.3	Computer vision system software	106
5.3.4	Data synchronisation	110
5.3.5	Integration software	111
5.4	Chapter conclusion	111
6	Testing and Evaluation.....	112
6.1	RFID tests	112
6.1.1	Experimental parameters	112
6.1.2	Test environment.....	112
6.1.3	Reader mode test.....	114
6.1.4	Multi Tag test	121
6.1.5	Tag Height test	126
6.1.6	Vehicle lateral offset test.....	132
6.1.7	Multi reader test	137
6.1.8	Worst case test	145
6.2	Computer vision testing.....	147
6.2.1	ANPR	147
6.2.2	ANPR road testing.....	149
6.2.3	VMM	152

6.2.4	VMMR road testing.....	155
6.2.5	Integration testing	157
6.2.6	Integration testing conclusion	168
6.3	Chapter conclusion	168
7	Conclusion	170
7.1	Research questions answered	170
7.1.1	Can RFID and Computer Vision technology be combined to increase the effectiveness of illegal vehicles detection?	170
7.1.2	What are the performance criteria of the proposed in-road RFID system?.....	171
7.1.3	Does the system meet the specific theoretical requirements as stated in literature?	172
7.2	Summary	173
7.3	Future research.....	174
8	References.....	175
9	ANNEX A	188
9.1	Additional information on CNNs.....	188
9.1.1	Convolutional layer	188
9.1.2	Activation layer	189
9.1.3	Pooling layer.....	190
9.1.4	Fully connected layer	190
9.1.5	The Back propagation algorithm.....	191
9.1.6	Loss function	191
9.1.7	Dropout	192

List of Figures

Figure 1: Deep learning model structure [39].....	9
Figure 2: Computer vision for EVI [90].....	10
Figure 3: Vehicle model classification.....	11
Figure 4: Passive RFID communication [71].....	13
Figure 5: RFID EVI system [72]	14
Figure 6: Spatial hierarchy of the visual world [73]	16
Figure 7: Transfer learning applied to vehicle classification [74]	17
Figure 8: Example of a model overfitting the data	18
Figure 9: Generating cat images using data augmentation [75].....	19
Figure 10: Bilinear CNN for image classification [76].....	21
Figure 11: License plate recognition [77].....	23
Figure 12: AOLP vehicle dataset	25
Figure 13: multi-stream vehicle recognition model proposed in [40]	27

Figure 14: Sample of the CompCars dataset.....	28
Figure 15: Image without(left) and with(right) infrared [78].....	30
Figure 16: Ideal vertical camera installation [78]	31
Figure 17: Ideal horizontal camera installation [78].....	31
Figure 18: Camera focal point [78]	32
Figure 19: ANPR misread due to screw fixing [18].....	33
Figure 20: Slotted metal numberplate tag [131]	36
Figure 21: Hollow plastic numberplate tag [130]	37
Figure 22: On metal numberplate tag [132]	38
Figure 23: Linear vs circular polarization [79].....	40
Figure 24: Desired radiation pattern [8]	41
Figure 25: Effects of metal on RFID performance [14]	44
Figure 26: Evenly spaced reader interference [1].....	46
Figure 27: Effective reading zone [69]	47
Figure 28: SANRAL Gantry locations [80].....	48
Figure 29: System context diagram	58
Figure 30: System functional flow	59
Figure 31: Operational context	59
Figure 32: Extract vehicle features context	60
Figure 33: Capture tag data context	60
Figure 34: High level architecture.....	61
Figure 35: Full system architecture.....	62
Figure 36: ThingMagic M6E micro RFID reader module [115].....	65
Figure 37: TÖNNJES IDePLATE [81]	67
Figure 38: RCSensor developed by Licensys [69].....	70
Figure 39: HiLook IP bullet camera [143].....	72
Figure 40: Tag read range measuring points	78
Figure 41: Read range test setup (vehicle mounted tag).....	78
Figure 42: RFID number plate read range.....	81
Figure 43: Orientation tolerance test illustration	82
Figure 44: Tag orientation tolerance test setup	83
Figure 45: Tag orientation tolerance	85
Figure 46: Tag measuring points.....	86
Figure 47: Metal background test setup.....	87
Figure 48: Parallel setup	88
Figure 49: In-road reader antenna test setup.....	90
Figure 50: In-road antenna test setup	91
Figure 51: In-road antenna radiation pattern.....	92
Figure 52: Tag illumination area test setup	94
Figure 53: Illumination area test setup.....	95
Figure 54: RFID signal illumination area	96
Figure 55: Max tag read range vs vehicle lateral position	96
Figure 56: Detailed system architecture.....	100
Figure 57: Assembled in-road RFID reader	101
Figure 58: Required RFID antenna ground plane.....	102
Figure 59: Camera system hardware	104
Figure 60: System software flow diagram	106
Figure 61: Vehicle detection pipeline	110

Figure 62: Low speed test environment	113
Figure 63: High speed test environment	113
Figure 64: RFID in-road setup	115
Figure 65: Reader mode test results for 96 bits	116
Figure 66: Reader mode test results for 192 bits	117
Figure 67: Reader mode test results for 480 bits	118
Figure 68: Reader mode test results for 640 bits	119
Figure 69: Multi-tag test setup	122
Figure 70: Multi-tag test (High throughput mode)	123
Figure 71: Multi-tag test (Dense M4 mode)	124
Figure 72: Multi-tag test (Dense M8 mode)	125
Figure 73: Tag height test results (96 bits)	128
Figure 74: Tag height test results (192 bits)	129
Figure 75: Tag height test results (480 bits)	130
Figure 76: Tag height test results (640 bits)	131
Figure 77: Lateral offset test results (96 bits)	133
Figure 78: Lateral offset test results (192 bits)	134
Figure 79: Lateral offset test results (480 bits)	135
Figure 80: Lateral offset test results (640 bits)	136
Figure 81: Multi reader test results (96 bits)	138
Figure 82: Multi reader test results (640 bits)	139
Figure 83: Multi reader test results (96 bits)	140
Figure 84: Multi reader test results (640 bits)	141
Figure 85: Multi reader test results (96 bits)	142
Figure 86: Multi reader test results (640 bits)	143
Figure 87: Worst case test results.....	146
Figure 88: Custom ANPR dataset example	148
Figure 89: Computer vision vehicle setup.....	151
Figure 90: Stanford Cars dataset sample [142].....	153
Figure 91: VMMR training accuracy.....	154
Figure 92: Computer vision vehicle setup (early afternoon)	159
Figure 93: Multi vehicle, single tag test setup	164
Figure 94: Image convolution	188
Figure 95: Layer representations learned by a CNN	189
Figure 96: Average and max pooling operations	190
Figure 97: Fully connected layer	191
Figure 98: The feedback signal is used to adjust the weights.....	192

List of Tables

Table I: Comparison between active and passive tags	12
Table II: Comparison between license plate recognition algorithms	25
Table III: Comparison between vehicle make and model recognition algorithms.	28
Table IV: In-Road and Overhead Comparison.....	53
Table V: Comparison between different RFID readers	64
Table VI: ThingMagic M6E micro specifications	65
Table VII: Comparison between RFID tags.....	66
Table VIII: Toenjes IDEPLATE specifications.....	67

Table IX: Comparison between RFID antennas.....	68
Table X: Licensys RCSensor specifications	70
Table XI: Comparison between CV cameras	71
Table XII: Reading protocol	74
Table XIII: Experimental equipment	76
Table XIV: Read range calculation at 0 degrees (free space)	80
Table XV: Read range calculation at 0 degrees (vehicle mounted)	80
Table XVI: Tag orientation test result	83
Table XVII: Tag orientation test result	84
Table XVIII: YOLOv7 model architecture.....	107
Table XIX: VGG16 model architecture	108
Table XX: PaddleOCR model architecture.....	109
Table XXI: ThingMagic reader modes	114
Table XXII: 96-bit Reader mode test	116
Table XXIII: 192-bit Reader mode test results	117
Table XXIV: 480-bit Reader mode test results	117
Table XXV: 640-bit Reader mode test results	118
Table XXVI: Multi-tag test (High throughput)	122
Table XXVII: Multi-tag test (Dense M4)	124
Table XXVIII: Multi-tag test (Dense M8)	125
Table XXIX: Tag height test (96 bits)	127
Table XXX: Tag height test (192 bits)	128
Table XXXI: Tag height test (480 bits)	129
Table XXXII: Tag height test (640 bits)	130
Table XXXIII: Vehicle offset test results (96 bits)	133
Table XXXIV: Vehicle offset test results (192 bits)	133
Table XXXV: Vehicle offset test results (480 bits)	134
Table XXXVI: Vehicle offset test results (640 bits)	135
Table XXXVII: Multi reader test (High throughput mode - 96 bits).....	138
Table XXXVIII: Multi reader test (High throughput mode - 640 bits).....	138
Table XXXIX: Multi reader test (Dense mode M4 - 96 bits)	140
Table XL: Multi reader test (Dense mode M4 - 640 bits).....	141
Table XLI: Multi reader test (Dense mode M8 - 96 bits).....	142
Table XLII: Multi reader test (Dense mode M8 - 640 bits).....	143
Table XLIII: Worst case test.....	146
Table XLIV: ANPR model accuracies.....	149
Table XLV: ANPR test results (vehicle 1)	151
Table XLVI: ANPR test results (vehicle 2)	152
Table XLVII: Summary of Computer Vision testing accuracies	154
Table XLVIII: VMMR test results (vehicle 1)	156
Table XLIX: VMMR test results (vehicle 2)	156
Table L: Ideal integration test results (early afternoon)	160
Table LI: Ideal integration test results (early evening)	160
Table LII: Single vehicle, multi-tag test results (vehicle 1)	162
Table LIII: Single vehicle, multi-tag test results (manually introduced tag).....	162
Table LIV: Multi-vehicle, single tag test results (vehicle 1).....	165
Table LV: Multi-vehicle, single tag test results (vehicle 2).....	165
Table LVI: Multi-vehicle, multi tag test results (vehicle 1).....	167

Table LVII: Multi-vehicle, multi tag test results (vehicle 2).....	167
--	-----

Abbreviations

ANPR – Automatic Number Plate Recognition
API – Application Programming Interface
ASK – Amplitude Shift Keying
BCNN – Bilinear Convolutional Neural Network
BLF – Band Link Frequency
BS – Base Station
CNN – Convolutional Neural Network
CRCNN – Cascade Residual Convolutional Neural Network
CV – Computer Vision
DR – Divide Ratio
DSB-ASK – Double Side Band Amplitude Shift Keying
ECC – Elliptic Curve Cryptography
EIRP - Equivalent isotropic radiated power
EPC – Electronic Product Code
ETSI - European Telecommunications Standards Institute
EVI – Electronic Vehicle identification
FHSS – Frequency Hopping Spread Spectrum
FPS – Frames Per Second
HOG – Histogram Oriented Gradients
IEC – International Electrotechnical Commission
IP – Ingress Protection
IR – Infrared
ISO – International Organization for Standardization
ITS – Intelligent Transportation Systems
LP – License Plate
MDC – Micro Data Centre
MLC – Magnetic Loop Counter
NFC – Near Field Communication
OCR – Optical Character Recognition
PR-ASK – Phase Reversal Amplitude Shift Keying
RF – Radio Frequency
RFID – Radio Frequency Identification
RSSI – Received Signal Strength Indicator
RSU – Roadside Unit
SANRAL – South African National Road Agency Limited
SDK – Software Development Kit
SGD – Stochastic Gradient Decent
SIFT – Scale Invariant Feature Transform
SINR – Signal to Interference and Noise Ratio
SURF – Speeded Up Robust Features
SWP - Spatial Weighted Pooling
TARI – Type A Reference Interval
TID – Tag ID
TUT – Tag Under Test
UART – Universal Asynchronous Receiver Transmitter
UHF – Ultra High Frequency
UII – Unique Item Identifier
US – United States
VIN – Vehicle Identification Number
VMMR – Vehicle Make and Model Recognition
VTP – Vehicle Test Point
YOLO – You Only Look Once

1 Problem Statement

This chapter explores the complex engineering problem and the solution to solve this problem. It also documents the context in which the problem exists, introducing the project objectives, benefits, requirements, scope, and deliverables.

1.1 Introduction

Road transport systems support sustainable growth across society and the economy. However, increasingly congested traffic conditions and the illegal use of roads pose challenges to optimally managing road networks [8]. Intelligent Transportation Systems (ITS) aim to improve the user experience by increasing traffic safety and reducing congestion and illegal road use [10]. To achieve this, ITS systems must focus on effective law enforcement that can be reconciled with streamlined traffic flows within a traffic management system.

Over the past decades, traffic authorities have incorporated numerous technological approaches into law enforcement activities to enhance operational efficiency. Unfortunately, although these approaches may have some success in reducing illegal road use, they often come at the cost of negatively impacting the experience of law-abiding road users.

The primary means of identifying suspicious vehicles is through static checkpoints. Ad hoc static checks on a selected number of vehicles would mean offenders have a better chance of escaping law enforcement. The variability of human nature also exposes the approach to bribery and dishonest behaviour from law enforcement officials [19].

A common approach requires legal vehicles to be identified by a license plate. However, the license plates used in most countries are relatively easy to clone, and criminals buy these illegal plates from dealers or re-use scrapped legal plates [35].

Traffic cameras present an effective way to monitor traffic flow and enforce speed laws without disrupting traffic flow. Traffic cameras use license plates to identify vehicles; however, these systems cannot distinguish between legal and illegal license plates. Road users can also remove, obscure, and falsify plates to evade camera detection [19].

Adel Ammar et al. [84] identified several key challenges in license plate detection techniques. First, images used for vehicle identification are captured in unconstrained environments, where factors like varying lighting conditions and camera angles often reduce detection accuracy. Second, license plate designs vary by region and even across different production years, adding complexity to accurate identification. Third, the infrastructure needed to manage the detection system is demanding: license plate information must be extracted in real time, synchronised with

the connected management system, and supported by significant resources to analyse and monitor tracked vehicles effectively.

Research presented in [82] introduces a method for detecting false license plates by analysing spatio-temporal inconsistencies in license plate data. The approach involves collecting vehicle footage from multiple traffic cameras positioned at different locations. By establishing spatio-temporal relationships between records with matching license plates, the system identifies potential discrepancies based on vehicle movement patterns, enabling the detection of false plates. While this technique shows promise, it relies on an externally connected system and requires a substantial database of vehicle records for effective implementation.

Another technique proposed in [83] uses RFID technology to uniquely identify and track stolen vehicles. The system uses passive tags installed on vehicles to act as unique identifiers. These tags can be read by RFID readers positioned along traffic routes and capture the information on a centralised server. When a stolen vehicle is reported, the system uses a dedicated algorithm to search for the last known location of the tagged vehicle. Once the vehicle is located, necessary measures can be implemented to secure it. The main pitfall of this system is that the RFID tags may be prone to tampering or removal, rendering the system ineffective. Unauthorised individuals could disable or replace tags without the system noticing.

Furthermore, there is also a need for immediate intervention versus post-detection intervention. While an unmanned automated system can detect a potential suspect and initiate an intervening action, there is a possibility that the vehicle could be beyond the reach of authorities by the time that action is taken. If such suspects use cloned identities of legal vehicles, efforts towards law enforcement will be associated with increased harassment of law-abiding citizens. A system that aims to ensure that all vehicles on the road are legally identifiable must enable immediate intervention on a tactical level. Access to strategic-level information such as location and identity is crucial to allow law enforcement officials to achieve a prompt tactical response, ensuring they are positioned correctly and at the appropriate time. The officials may be roaming police vehicles with their own detection sensors or stationary police vehicles connected to roadside infrastructure or a local command centre.

1.2 Problem definition

The increase in traffic volume is proportional to the increase in illegal road use, i.e., more traffic would mean even higher percentages of illegal road users. Illegal road use primarily includes the cloning of number plates, migration of a legal license plate between different vehicles and intentionally avoiding detection by obscuring/damaging the number plate. To effectively combat the illegal use of roads, authorities must be able to reliably verify the identity of the vehicles.

From the background discussion it is clear that current approaches to traffic law enforcement are not effectively combatting illegal road use. Consequently, there is a pressing need for a fool-proof system that can accurately identify illegal vehicles and respond to illegal road use without compromising the experience of legal road users.

1.3 Proposed solution

This thesis investigates a system for traffic law enforcement that combines two different technologies to overcome some of the challenges faced by existing approaches. By embedding RFID tags in vehicle license plates, a unique identifier can be assigned to each vehicle, creating a "fingerprint" for the vehicle. This information can be used to uniquely identify vehicles and track them as they move. Computer Vision is used to extract vehicle identification information from traffic surveillance data. If a tagged number plate moves through the detection zone, it can be reliably detected by comparing the features saved on the tag with the features extracted from the traffic camera system. A cryptographic suite will be used to protect the data stored on the tag.

Combining these technologies will enable multi-factor verification of the true identities of suspicious vehicles. It will also provide a means of detecting vehicles that do not conform to the regulatory marking requirements. The system will equip authorities with strategic-level information to achieve immediate intervention on a tactical level.

The identities of vehicles can be verified automatically when moving over in-road RFID readers, or manually with handheld RFID readers at dedicated roadblocks to enable effective and efficient intervention, ensuring that a road-user knows that if you drive with a false vehicle identity you will be caught, sooner rather than later.

Previous research [35], [8] described an intelligent vehicle monitoring system based on in-road RFID readers. It defined the technical requirements for such readers to detect RFID-enabled license plates mounted on in-traffic vehicles successfully. Building on this foundation, this thesis aims to develop a more robust and reliable vehicle monitoring system.

1.4 Project aim and objectives

1.4.1 Primary objective

The primary objective of the project is to investigate a novel traffic law enforcement system that combines RFID technology with Computer Vision for vehicle identification and verification, within a real-world environment. This includes determining the combination of data fields that can enable accurate vehicle identification within a practical setup. Practical measurements will verify if the theoretical requirements can be achieved in practice.

A specific focus will be placed on data verification from the tag, in contrast to verification from an online source.

1.4.2 Secondary objective

The secondary objectives include:

1. To investigate the effect of certain controllable parameters on the performance of the system, both in a controlled laboratory and in a practical traffic environment
2. To assess the accuracy and reliability of the RFID and Computer Vision system in identifying vehicles equipped with RFID-enabled license plates in controlled traffic scenarios.
3. To compare the performance of the proposed system with the technical requirements defined in previous research [35], [8] and identify any areas for improvement.

These objectives involve the development of the system, its testing in real-world conditions, and the assessment of its effectiveness for vehicle identification to support law enforcement.

1.4.3 Research questions

The following questions should be answered at the end of the study:

- Can RFID and Computer Vision technology be combined to increase the effectiveness of illegal vehicles detection?
- What are the performance criteria of the proposed system?
- Does the system meet the specific theoretical requirements as stated in literature?

1.5 Anticipated benefits

The following are some of the anticipated benefits of the solution:

- It will help authorities to regulate traffic in dense traffic environments. The RFID system can be used to generate traffic data and can automatically alert authorities of congestion and help regulate the flow of traffic by using alternative routes.
- It will help authorities to enforce road laws, such as speed limits and road access. The RFID system can be used to measure the average speed of vehicles, which can be used to identify vehicles that are exceeding the speed limit.
- It will prevent criminals from migrating legal plates between different vehicles. As previously mentioned, the solution will provide unique vehicle identities. Any manipulation of the vehicle identity will be detected by the system and reported to authorities. However, the solution may struggle to automatically detect plate migration between vehicles of the same make and model.

- It will prevent criminals from illegally cloning or duplicating legal number plates. The secure vehicle fingerprint will make it nearly impossible for criminals to clone plates.
- The solution can be used to identify priority vehicles such as ambulances. Priority vehicles can be detected by the RFID system, which can provide them with prioritised right of way.
- The system can be used in a smart toll collection setup. Toll fees can automatically be added to a registered road user's account when passing a dedicated RFID reading point.
- The solution can be used to assist autonomous driving in specific contexts. Autonomous vehicles equipped with UHF RFID readers move within an environment where specified routes are marked with RFID tags that are fixed to the driving surface to create a grid. As the vehicle moves through the marked environment, it scans the area to search for tags, using the landmark coordinates to map its position [62]. This may be extended to public road use for future autonomous traffic networks.

1.6 Research Scope

The aim of this study is not to develop a system that automatically detects and deters criminal activity by issuing tactical and strategic interventions to law enforcement, but rather to identify the critical technical components and investigate their potential within such a system. Therefore, the scope of the project includes the development of a combined RFID and camera-based system to uniquely identify vehicles and to investigate the performance of such a system in a real-world environment. This entails the extraction of vehicle features from traffic cameras to serve as a digital fingerprint, storing this data on an RFID tag embedded in the license plate of a vehicle, and encrypting the data for security purposes. Furthermore, it includes deploying an operational RFID/camera system using in-road RFID readers and tagged number plates and testing the system to evaluate its potential in a traffic environment.

1.7 Deliverables

Completion of the project will satisfy the following deliverables:

- A system that combines RFID technology and Computer Vision for vehicle identification and verification
- Documentation detailing the system's architecture, design, and setup instructions for easy replication.
- A report summarizing the results of real-world testing, including reliability and performance metrics.
- Detailed analysis of the system's performance compared to predefined technical requirements and benchmarks.
- A technique that can uniquely identify vehicles in-motion.
- A technique that enables the offline verification of vehicle identities.

- A technique that can be used to detect and prevent the cloning and migrating of number plates.

1.8 Research Methodology and Design

To accomplish the objectives outlined in this study, it is essential to devise a well-structured research strategy. The initial phase involves a meticulous examination of existing studies related to the topic, aiming to pinpoint their limitations. This critical analysis forms the foundation for enhancing future research designs. The primary focus of the study will be on advancing the system-level design of a vehicle identification system.

Employing a Quantitative research methodology, the following sequential steps will guide the development of the proposed system:

Literature Review and Gap Analysis: Conduct a comprehensive literature review to identify gaps and shortcomings in current research on the subject. This critical analysis will serve as a basis for refining and improving future study designs.

Preliminary System Design: Focus on the system-level design of the vehicle identification system. Clearly outline the architecture, components, and functionalities of the proposed system based on the insights gained from the research.

Initial testing and experimentation: Following the selection of components, the next pivotal phase involves the establishment and testing of an initial prototype installed within a semi-controlled environment. This preliminary testing serves as a foundational step, setting the stage for the subsequent detailed design and testing phase. The primary objectives during this stage are to assess the capabilities and limitations of the system and to ascertain whether the predefined specifications have been successfully met.

Detail Design: Building upon the insights gained from the initial testing and system validation, the detailed design phase emerges as a crucial stage where necessary adjustments are made to enhance the system's overall functionality. This phase encompasses the construction of components as required by the in-road RFID system, including a suitable 3D printed casing, and the software design as required for both the RFID and computer vision subsystems.

Implementation and Testing: Validate the designed system through rigorous testing procedures. This step ensures that the proposed solution meets the predefined objectives and is effective in addressing the identified limitations.

Discussion: This phase involves a thorough examination of the test results obtained during the research, aiming to extract meaningful insights and draw informed conclusions. The discussion serves as a platform for a detailed analysis, providing context to the findings and shedding light on their implications.

Conclusion and recommendations: Summarise the findings and provide actionable recommendations for implementing improvements in future vehicle identification system designs.

By meticulously following these steps, the research endeavours to contribute valuable insights to the field and foster advancements in the system-level design of vehicle identification systems.

1.9 Ethical Considerations

There are no ethical considerations or conflict of interest to consider for this study.

2 Literature Review

This chapter delves deeper into the theoretical work and previous studies related to the envisaged system. This comprehensive examination is crucial for establishing a robust understanding of the problem domain, thereby laying the groundwork for informed and effective problem-solving strategies.

2.1 Electronic Vehicle Identification

EVI (Electronic vehicle identification) is a sub-system of ITS that uniquely identifies a vehicle electronically and stores identification information digitally. In recent years two distinct methods have been developed for electronic identification namely image-based techniques and sensor-based techniques [40].

Image-based techniques rely on cameras to take snapshots of reality and require sophisticated processing techniques to turn this snapshot into useful digital information. Image-based techniques have advanced to the point of near-human level classification and identification. Sensor-based techniques make use of electronic devices to detect, monitor and control the traffic environment. RFID (Radio Frequency Identification) is an example of a sensor-based technique that has been introduced to the traffic environment [85]. Recent developments in RFID technology have made it a suitable solution for EVI. RFID systems offer storage capabilities, while being cost-effective.

2.1.1 Computer vision systems

Computer vision, an integral component of artificial intelligence, involves the extraction of meaningful information from digital images, videos, and various visual inputs. Enabling computers to perceive, observe, and comprehend their surroundings, this field mirrors the human ability to train minds through environmental observation over time. Instead of relying on biological components like retinas, optic nerves, and visual cortex, computer vision employs mathematical frameworks, cameras, data, and processing algorithms [86]. The training process involves exposing the system to extensive data rapidly, as computer vision systems process thousands of training samples through successive layers. These layers facilitate the learning of new data representations, allowing the system to recognise and distinguish images [87]. To illustrate, training a computer to identify diverse vehicle types necessitates the supply of thousands of images representing different vehicle classes for learning distinctions and recognition.

Facilitating these human-like tasks in computer vision are two crucial subfields of machine learning: Deep Learning and Convolutional Neural Networks (CNN). The term "deep" in deep learning signifies the incorporation of recurrent layers of representations. The depth of a deep learning model is determined by the number of contributing layers, with most models consisting

of hundreds of layers automatically learned from exposure to extensive training samples [88]. Figure 1 depicts the structure of a deep learning model as an example of this intricate process.




Figure 1: Deep learning model structure [39]

In deep learning applications tailored to computer vision, CNNs have emerged as a foundational technique [97]. In the era preceding the prominence of convolutional networks, data scientists heavily relied on feature engineering to derive meaningful information from images. Commonly, these techniques involved hardcoded features such as counting loops or assessing object heights in images, demanding a profound understanding of the problem. In contrast, CNNs streamline this process by automatically extracting valuable features from the data [89].

Convolutional neural networks employ convolutional layers to grasp local patterns within windows of the input image. These layers learn by applying convolutional transformations to localised patches extracted from the input image [39]. For instance, a CNN initially identifies edges and shapes and progressively acquires additional information as it iterates over the training data. Notably, while these networks offer automated feature extraction, they come with a substantial computational resource demand and a requisite for extensive datasets for effective training. Chapter 5 provides a more in-depth exploration of CNNs.

2.1.2 Applications of Computer Vision in EVI systems

Computer vision is playing a key role in the evolution of the transportation system. Computer vision is largely manifesting itself in the traffic environment through video-based imaging. Figure 2 below shows an example of a vision-based traffic system. These systems usually consist of traffic cameras, data processing centres and databases. The traffic cameras monitor certain traffic/ vehicle parameters and report data to the processing centres. The processing centres analyse the data to make predictions or act on a specific situation [90].




Figure 2: Computer vision for EVI [90]

Automatic Number Plate recognition (ANPR)

In the realm of image-based systems, Automatic Number Plate Recognition (ANPR) systems live up to their name by autonomously reading vehicles' number plates without human intervention [91]. The ANPR algorithm orchestrates a sequence of steps to accomplish this task [99]. Initially, the input image undergoes capture and pre-processing to ensure it aligns with the correct format, exhibits quality, and maintains the appropriate size. The subsequent step involves number plate detection, where an algorithm or machine learning model adeptly localises the number plate within the image. Following detection, the process moves into license plate recognition, encompassing a series of steps to identify, extract, and recognise characters from the license plate, a procedure known as Optical Character Recognition (OCR). Character detection and extraction typically involve segmenting characters and creating bounding boxes around each, with the final recognition performed by a trained convolutional neural network [40].

As a vehicle traverses an ANPR camera, its registration number is read and cross-referenced against records stored in a database. ANPR systems have found widespread applications in e-tolling, discouraging criminal activities, and enforcing traffic laws. Presently, these systems achieve accuracies exceeding 90% [92], which means multiple undetected vehicles. This limitation poses challenges, particularly in areas with high traffic volumes and increased criminal

activity on roads. Consequently, it becomes evident that ANPR alone falls short, necessitating supplementary solutions to enhance overall vehicle identification accuracy.

Vehicle classification

Another application of computer vision in a traffic scenario is identifying and classifying different vehicles. Other than identifying the vehicle by its number plate, the physical characteristics such as make, and model can be used. Vehicle make classification is less complex than model classification due to the higher similarity between vehicles of the same make. Vehicle manufacturers copy certain traits from previous models into new models [93]. Figure 3 below shows an example of how difficult it can be to distinguish between two different vehicle models.



Figure 3: Vehicle model classification

The heightened resemblance among vehicle parts amplifies the intricacies involved in model classification, given that algorithms must discern distinctions based on subtle differences. In response to this challenge, researchers have devised multi-stream classification architectures. These architectures deploy feature extractors to each individual part as well as to the comprehensive ensemble of parts [40]. Section 5.2.6 furnishes an in-depth exploration of these methodologies, offering a comprehensive understanding of the strategies employed to successfully identify and leverage the nuanced variations within vehicle components.

Vehicle classification can be used in conjunction with ANPR systems to accurately detect vehicle anomalies. Number plate cloning and migration can more easily be detected through an automatic system that captures the number plate and links this information to other vehicle identity features.

2.1.3 RFID systems

RFID, an automatic wireless data collection technology, consists of essential components such as the RFID reader, RFID tag, and antenna [94].

There are mainly two types of RFID systems:

1. Active (Tags are energised by a battery)
2. Passive (Tags are energised by the received Electromagnetic waves from the reader. [94])

Table I compares the active and passive topologies [95]:

Table I: Comparison between active and passive tags

Active		Passive	
Advantages	Disadvantages	Advantages	Disadvantages
Longer read range	Expensive	Increased security	Short read range
Larger memory	Bulky	Small	Lack of backup power
Accurate capturing	Hard to hide	Less expensive	Less accurate capturing
Backup power supply	Less secure	More secure	Limited memory

This study will exclusively concentrate on passive RFID systems. In a passive RFID system, the process unfolds as follows: an RFID reader emits a modulated RF signal towards the RFID tag, which comprises an antenna and an integrated circuit chip. This chip draws power from the antenna and, in return, adjusts its input impedance, leading to the modulation of the backscattered signal. Amplitude shift keying (ASK) is a prevalent modulation technique in RFID, where the chip's impedance alternates between two states. One state aligns with the antenna, facilitating power collection, while the other remains distinctly mismatched [95]. Figure 4 visually outlines the passive RFID communication process.




Figure 4: Passive RFID communication [71]

Passive systems are designed to operate in three main frequency bands, namely:

- 125 – 134 kHz (low frequency) – These systems typically have a very short read range, ranging from 1cm to 10cm, and are widely used for livestock tracking [95].
- 13.56 MHz (High Frequency) – These systems have a longer read range than the low frequency devices, between 1 centimetre and 1 meter. They are commonly used for access control and data transmission (NFC) [95].
- 865 – 960 MHz (Ultra High Frequency) – These systems make up the majority of the RFID population and have read ranges between 2 meter and 15 meter. These systems are commonly used for supply chain management and race timing [95].

Increasing the RFID system's frequency extends its read range yet simultaneously increases its susceptibility to reflections and interference [96]. To achieve effective vehicle detection using RFID, a read range of typically more than 2 meters is essential. Consequently, our research will emphasise UHF RFID systems.

2.1.4 Applications of RFID in EVI systems

A typical UHF (Ultra High Frequency) RFID based vehicle identification system is depicted in Figure 5. The vehicles in the scenario are fitted with RFID tags. When the vehicles pass a reader, the data from the tags are read wirelessly [10]. This data is typically communicated to a roadside unit which processes the data and communicates the information to the cloud.




Figure 5: RFID EVI system [72]

Below are some of the applications of RFID-based EVI systems:

- Navigation system – Vehicles can estimate their position by obtaining position data from an online cloud based navigation system that is linked to the in-road RFID reader network. When the vehicle drives over the reader, the tag is detected, and the identity and location information is logged in the cloud and can be retrieved by a cloud connected vehicle. The cloud-based reader network can also provide the user with data such as traffic direction, speed limit, lane destination and event notifications [33][34].
- Access control – The system can automatically control the access of vehicle to select locations like parking garages or property estates [100].
- ITS (Intelligent traffic system) – The system can identify priority vehicles, smart tolling of vehicles, track stolen vehicle, generate traffic data [100].
- Unmanned patrol systems – Violation of traffic rules can be detected 24/7 and authorities can respond based on the received information [100].
- Unmanned vehicles – RFID systems can aid in the development of autonomous driving systems [101].
- Estimation of Vehicle Distance: The real-time estimation of distances between vehicles on the road is crucial for providing drivers with essential information to ensure safe distances and prompt actions. Yan Huo et al. [63] introduced a method for maintaining road distances to establish a secure spacing requirement between vehicles. In this system, vehicles are equipped with RFID readers, and RFID tags are affixed to the pavement. The system records the speed and timestamp of each vehicle as it passes over an RFID tag. The information stored on the tags is read to determine the status of the vehicle in front. To uphold a safe following distance, the system adjusts the vehicle speed based on the calculated distance between vehicles.

2.2 Computer Vision for Intelligent Transportation Systems

The widespread deployment of video surveillance cameras in traffic scenarios serves various purposes, including monitoring traffic conditions, implementing electronic tolling, and overseeing criminal activities [40]. While surveillance data accurately reflects reality, the challenge lies in the sheer volume of data, making manual analysis nearly impossible on a human scale. Trained programs conducting automatic analysis of video streams prove invaluable in simplifying these efforts. Substantial focus has been directed toward surveillance tasks like vehicle counting and speed estimation [33,36,37], utilizing sensor-based or computer vision techniques. Visual approaches excel in recognizing and analysing patterns for advanced tasks such as vehicle identification and Automatic License Plate Recognition (ANPR) [38], whereas sensor-based methods like Magnetic Loop Counters (MLC) and radar techniques are limited to detecting vehicles in specific areas.

Until recently, none of these techniques could match human proficiency in vehicle identification and classification. However, the landscape changed with the advent of CNNs. CNNs, mimicking the human brain in analysing visual data, enable computers to perform visual data manipulation tasks at a level approaching human capability. This transformative capability has elevated the field, allowing computers to achieve near-human proficiency in visual data analysis and manipulation [89].

2.2.1 Visual Feature extraction

In the early stages of visual feature extraction, emphasis was placed on manually extracting low-level features. One of the pioneering algorithms in this domain was SIFT (Scale Invariant Feature Transform), designed to detect and describe features in an image [42]. This method involved extracting select points in an object, storing them, and subsequently comparing the newly extracted features with the stored ones to recognise the object. While effective, this approach suffered from computational inefficiency. Subsequent advancements led to the introduction of SURF (Speeded Up Robust Features) and HOG (Histogram Oriented Gradients) to enhance processing speed. However, these methods still required manual engineering of layers of representations in a process known as feature engineering—a task often challenging and prone to brittleness [102].

In contrast, deep learning revolutionised this paradigm by automating the feature engineering step. Instead of manually engineering features, deep learning learns all the relevant features in a single pass [40]. The upcoming section will delve into the theories and methods of the deep learning approach, highlighting its capacity to streamline and automate the feature extraction process for more efficient and robust visual analysis.

2.2.2 Deep learning

2.2.2.1 Convolutional Neural Networks

A CNN stands out as a subtype of deep learning neural networks with a primary focus on applications related to vision. The architecture of CNNs is characterised by layers stacked sequentially. In this arrangement, the initial convolutional layer excels at extracting fundamental local features, such as edges and curves. Subsequent layers build upon these extracted features, learning larger patterns that comprise the features from the preceding layer, and this hierarchical process continues. Consequently, CNNs demonstrate the ability to extract increasingly complex visual concepts [97].

The efficacy of convolutional layers in CNNs stems from the inherent spatial hierarchy present in the visual world. Local edges come together to form local objects, which, in turn, combine to create high-level concepts, as illustrated by the cat image in Figure 6. This hierarchical feature extraction mirrors the natural organization of visual information, making CNNs particularly effective in capturing and understanding the hierarchical structure of visual data [39].




Figure 6: Spatial hierarchy of the visual world [73]

A CNN normally consists of convolutional layers, activation layers, pooling layers, fully connected layers, and normalization layers, where each layer learns a different representation of the original input [98]. Additional information regarding the functioning of the different layers is provided in the appendix. A disadvantage of the layered architecture is that it can become computationally expensive when layers are added [40].

2.2.2.2 Machine Learning Paradigms

Machine learning paradigms are broadly categorised into two primary types: supervised learning and unsupervised learning. Supervised learning, the more prevalent of the two, involves labelled training data, where the algorithm is trained using data with predefined labels. In contrast, unsupervised learning lacks labelled training data, and algorithms within this paradigm aim to discover inherent patterns or groupings in the data. Clustering is a typical form of unsupervised learning, focusing on automatically identifying groupings in the dataset [41].

An interesting hybrid of supervised and unsupervised learning is found in the realm of transfer learning. The primary goal of transfer learning is to reduce the amount of data required for training a network, especially when faced with constraints such as limited resources and time. This approach entails training a model on one task and evaluating it on a different but related task. A pre-trained network, for example, is one that has undergone training on a substantial dataset, such as a large-scale image classification task like ImageNet. If these pre-trained networks are sufficiently large and general, they can be repurposed as versatile models for various computer vision problems. As an illustration, Figure 7 demonstrates how transfer learning can be employed to classify vehicles by repurposing a network trained on ImageNet for a different but related task.




Figure 7: Transfer learning applied to vehicle classification [74]

2.2.2.3 Overfitting

In machine learning, the ultimate objective is to create models that generalise well, meaning they perform effectively on new, unseen data. Overfitting represents a significant challenge in achieving this goal. It occurs when a network models the training data too meticulously, capturing not only the underlying patterns but also the noise present in the data, which impairs the model's generalization capability. Essentially, overfitting arises when a model learns the idiosyncrasies of the training data to an extent that hinders its performance on broader datasets.

Figure 8 illustrates an example of overfitting, where the model's accuracy on the validation set plateaus after a certain number of iterations (epochs), indicating the model is starting to overly tailor itself to the training data and losing generalization ability. At the initial stages of training, models often underfit the data. However, with continued iterations, they may begin to overfit, learning patterns specific to the training set but irrelevant to new data.




Figure 8: Example of a model overfitting the data

While exposing the model to more training data is the ideal way to prevent overfitting, this is not always feasible due to limited data availability. To address this, techniques such as data augmentation are employed, particularly in the context of convolutional neural networks (CNNs). Data augmentation introduces diversity to the training dataset through transformations like rotation, scaling, or flipping, preventing the model from memorizing specific patterns and enhancing its ability to generalise effectively.

2.2.2.4 Data Augmentation

Overfitting often arises from a shortage of training samples, making it difficult for a model to generalise to new data. While the ideal scenario would involve an infinite amount of training data to prevent overfitting, this is rarely feasible. To address this limitation, data augmentation is employed. This technique generates additional training data by applying transformations to existing samples, creating diverse representations of the original data and enhancing the model's generalization capability [103].

Common data augmentation methods include rotation, shifting, shearing, zooming, flipping, and filling transformations [39]. These techniques ensure the model is exposed to varied data, reducing the likelihood of encountering the same sample twice and encouraging better generalization. Figure 9 demonstrates how data augmentation can be used to create additional training images for cat classification.



Figure 9: Generating cat images using data augmentation [75]

Data augmentation only modifies existing information to generate new samples. This means that the new samples are still highly correlated with the few original samples. The model is not exposed to entirely new input data; thus, it is not enough to eliminate overfitting.

2.2.2.5 Machine learning architectures

The architecture of a machine learning algorithm, particularly in the context of deep learning, significantly influences its performance.

Effect of the number of layers on the model performance

The number of layers in a neural network determines its capacity to learn complex representations from the data. Shallow networks, with fewer layers, are effective for simpler tasks but may struggle with complex datasets due to their limited ability to capture intricate patterns [136]. In contrast, deeper networks can model more complex features, as each successive layer extracts higher-level abstractions. However, deeper architectures are prone to challenges such as vanishing or exploding gradients, which can hinder training. Techniques such as batch normalization and residual connections, as introduced in ResNet, mitigate these issues and enable effective training of very deep networks [137].

Effect of the type of layers on the model performance

The type of layers also plays a crucial role. For instance, convolutional layers are specifically designed for spatial data, such as images, and are designed for tasks like object detection and image classification [138]. Fully connected layers are versatile but can be computationally intensive for large-scale problems. The inclusion of specialised layers like recurrent layers for sequential data or attention mechanisms for feature weighting further enhances performance for specific tasks [139]. Thus, the architecture's design, including the choice and arrangement of layers, directly impacts the algorithm's ability to generalise, its computational efficiency, and its overall accuracy.

2.2.2.6 Bilinear Convolutional Neural Networks (BCNNs)

Fine-grained computer vision tasks, such as identifying specific vehicle makes and models or distinguishing between bird species, are particularly challenging due to subtle category differences and confounding factors like pose variations and object location [68].

Traditional methods often rely on localizing parts of the image and framing their appearances based on location. However, this requires manual annotation and supervised training, which can be labor-intensive and impractical for certain tasks [68].

To address these limitations, Bilinear Convolutional Neural Networks (BCNNs) offer an alternative solution. A BCNN integrates two feature extractors, combining their outputs through a bilinear pooling layer. This layer computes the outer product of the features at each location and aggregates them across all locations to create a robust descriptor [68]. By leveraging bilinear pooling, BCNNs improve the efficiency and accuracy of fine-grained recognition tasks without requiring manual part annotations.

.




Figure 10: Bilinear CNN for image classification [76]

Figure 10 demonstrates the operation of a BCNN. An image is passed through two feature extracting CNNs, and their outputs are multiplied using the outer product. The result is then pooled and passed through a classification layer to deliver the predictions.

2.2.2.7 Visualising CNNs

Frequently researchers develop deep learning models without understanding the focus of the model or how the model makes decisions. This can be problematic as models can learn false representations that are far from the intended representation. Visualisation can help to fine-tune and improve a model based on the observed weaknesses.

Most often deep learning models are seen as ‘black boxes’, they learn representations from samples that can be difficult to present in human-readable form. However, because CNNs learn representation of visual concepts, they can easily be explored visually. Several techniques have been developed for visualising representations learned by CNNs.

- Visualising intermediate activations - This helps to get an idea of how the different convolutional layers transform the input [104].
- Visualising filters – Provides an understanding of what visual patterns each filter in the CNN captures [104].
- Visualising class activation maps – This allows you to visualise what image regions are used by a CNN to identify a specific class in the image, i.e. which regions were relevant to which classes [105].

2.2.2.8 One stage vs two stage detectors

One-stage detectors, such as YOLO and SSD (Single Shot Multibox Detector), are designed for speed and efficiency. They simplify the detection process by simultaneously predicting the class and location of objects in a single step. This streamlined architecture results in lower computational requirements and faster inference times compared to two-stage detectors. As a result, one-stage detectors are particularly suitable for real-time applications, where maintaining high frame rates and low latency is essential [140]. Recent advancements in one-stage detectors have significantly improved their accuracy, narrowing the performance gap with two-stage models. [141].

On the other hand, two-stage detectors like Faster R-CNN (Region-based Convolutional Neural Networks) and R-FCN (Region-based Fully Convolutional Networks) offer superior accuracy by breaking down the detection process into two stages. The first stage generates potential object proposals (regions of interest), and the second stage classifies these proposals and refines their bounding boxes. This additional step enables two-stage detectors to achieve higher precision, particularly for complex scenes or small objects. However, the trade-off is that two-stage detectors typically require more computational resources and time, making them less suited for real-time applications where speed is paramount [141].

2.2.2.9 Recognising Number Plates

Number plate recognition has evolved into a pivotal component of Intelligent Transportation Systems (ITS), attracting considerable attention with advancements in digital cameras and increased data processing capacity. Automatic Number Plate Recognition (ANPR) systems play a crucial role in uniquely identifying vehicles based on their license plates. These systems typically consist of cameras designed to detect and recognise number plates within a real-world scene. The camera captures an image of the vehicle, serving as input for an algorithm [40].

The algorithm, in turn, undergoes a series of steps to analyse, locate, and extract the number plate regions from the image background [40]. Subsequently, the algorithm engages in plate recognition by segmenting the characters within the detected region, enabling the identification of the vehicle through its number plate. Figure 11 visually encapsulates the sequential process involved in ANPR, illustrating how these systems effectively capture, analyse, and recognise number plates in practical scenarios.




Figure 11: License plate recognition [77]

License Plate Detection

The initial phase in license plate detection involves precisely localizing the plate within the input image. Various CNN methods have been developed for accurate object detection, including YOLO [44], YOLOv2 [45], and SSD [43]. In the detection process, the goal is to determine the bounding box that encapsulates the license plate within the input image. While traditional machine learning techniques with handcrafted features have been employed for plate detection, they often suffer from lower accuracy due to the insufficient features learned from characters [40].

In contrast, deep learning approaches have emerged as powerful alternatives. For instance, one approach [48] employs a 4-layer CNN to detect text zones within the image. Another study [49] utilises a deep CNN pipeline for license plate detection, accommodating different lighting conditions and a variety of license plates using traffic camera image data. These deep learning methods leverage the capacity of neural networks to automatically learn relevant features, enhancing the accuracy and robustness of license plate detection tasks. Moreover, other methodologies, as outlined in [47], consider the license plate as an area characterised by a dense concentration of key points, employing SIFT descriptors for detection tasks.

License Plate recognition

In the realm of license plate recognition, two main methods exist: segmentation-based and segmentation-free.

Segmentation-Based Methods:

1. *Projection-Based Approach*: This method leverages the colour difference between characters and the background of the license plate. Vertical and horizontal projections in a binary image help segment character. However, it can be sensitive to plate rotation.
2. *Connected Component-Based Methods*: These methods label connected pixels in a binary image for segmentation. While less sensitive to rotation, they struggle when characters are joined.

After segmentation, character recognition involves two algorithms:

- *Template Matching*: Compares the similarity of a character with existing templates. Best matching template is chosen.
- *Learning-Based Approaches*: Immune to rotation, illumination, and character font variations. Multiple features such as gradient and SIFT are extracted for character recognition.

Segmentation-Free Methods: In these methods, the entire license plate image is used for character recognition without segmentation. A sliding window generates possible characters, and a recogniser analyses the output characters incrementally. Deep learning approaches include:

- *CNN-Based Approaches*: [52] uses a CNN to recognise license plate characters extracted through segmentation-based methods.
- *YOLO-Based Approaches*: [54] employs a YOLO-based network for license plate character detection and recognition.

End-to-End Approach: Described in [40], an end-to-end methodology consists of two main stages: plate detection and plate recognition. For plate detection, YOLOv2 is employed, treating object detection as a unified regression problem. The output of the YOLOv2 network comprises cropped plate images that are subsequently fed into the recognition stage. In this stage, a collaboration between a CNN and a Recurrent Neural Network (RNN) occurs, eliminating the need for character segmentation.

These techniques underwent evaluation using the AOLP dataset, deliberately designed to assess the resilience of license plate detection and recognition algorithms across diverse conditions. The

AOLP dataset comprises images taken with cameras of varying quality, from different viewing angles, and in diverse lighting conditions as shown in Figure 12. This variety in data challenges the adaptability and reliability of the algorithms, providing a comprehensive testbed for their performance.



Figure 12: AOLP vehicle dataset

The AOLP dataset was used to train and test some of the approaches mentioned above. The accuracy and precision results from the tests are presented in the table below. The accuracy is the percentage measure of successfully recognised number plates. A correctly recognised number plate means all the characters of the plate are recognised correctly. The accuracy is measured by comparing the actual labels of the validation set with the predicted labels of set.

Table II: Comparison between license plate recognition algorithms

Method	Accuracy (%)
YOLOv2 used in [40]	91.48
CRCNN used in [40]	98.97
4 layer CNN used in [48]	97.62
SIFT Used in [47]	81.33

From the Table II above we can see that the hand-crafted feature extractors like SIFT [47] performed the worst. The YOLOv2 and CRCNN methods, which incorporates deep convolutional networks, provided the best accuracy results.

2.2.2.10 Vehicle make and model classification

Vehicle model identification has witnessed a transition from early recognition models using computationally expensive low-level feature representations like SIFT to more recent adoption of CNN models for speed, robustness, and enhanced feature extraction. While SURF and HOG provided improvements, CNN techniques with multiple hidden layers have become prominent, enabling the extraction of relevant high-level features resilient to noise.

Different approaches have been proposed for vehicle model identification:

1. **Part-Based Approach:** [55] proposed a part-based approach to detect components such as the license plate, lights, and grill using a CNN for global classification. However, this method was limited to frontal images.
2. **Relevance of Parts:** [56] focused on finding the most relevant parts in each class using a cascade of classifiers. [57] introduced a Spatial Weighted Pooling layer for magnifying features corresponding to discriminative parts, but it faced challenges with large variations in scale and position.
3. **Data Enhancement:** Boxcars [58] boosted recognition performance by extracting additional data from surveillance footage and incorporating it into the CNN. [59] utilised YOLOv3 to crop vehicle input images, reducing background noise and employing a multi-layer CNN with data augmentation and dropout functions.
4. **Global and Parts-Based Representations:** [40] proposed a robust methodology for make and model recognition, extracting both local and global feature representations. The system employed separate modules for local vehicle parts and global images, using a fusion layer for final prediction. Challenges included varying numbers of parts for different vehicles, leading to different-sized output feature vectors. The architecture employed YOLO for parts detection and a multi-stream approach for subtle changes in car models of the same make, utilizing a VGG16 network for recognition/classification. This model is depicted in Figure 13 below.

Overall, the shift from traditional feature extraction to CNN models has significantly enhanced vehicle model identification, addressed challenges and enabled more accurate and robust recognition.




Figure 13: multi-stream vehicle recognition model proposed in [40]

The methods for make and model detection described earlier have been evaluated using the CompCars [60] dataset. This dataset is comprehensive and diverse, encompassing images of vehicles from various angles, under different lighting and weather conditions. The CompCars dataset is publicly available and includes:

- Web-data: Consisting of 136,727 images capturing the entire car and an additional 27,618 images focusing on car parts. This part of the database covers 163 car makes and 1716 car models.
- Surveillance data: Comprising 50,000 car images captured in a front view. The surveillance data includes most of the makes and models present in the Web-data.

Figure 14 shows a sample of the CompCars dataset.



Figure 14: Sample of the CompCars dataset

Table III: Comparison between vehicle make and model recognition algorithms.

Approach	Accuracy (%)
Yang [60]	76.7
Boxcars [58]	84.8
Baseline VGG16 [40]	92.66
Multi stream VGG16 [40]	95.07
SWP-CNN [57]	97.6

Table III shows results on the CompCars dataset where approaches with no deep convolutional networks achieve the worst result, such as Yang [60]. The Baseline VGG16 and BoxCars deep networks rely on global feature representations, which achieved better results than the hand-crafted feature extractors, such as Yang [60]. However, the best results were achieved with the multi-stream VGG16 and SWP-CNN deep networks, which used parts-based approaches.

2.2.3 Factors affecting the Computer Vision EVI system

Most of the computer vision models developed in the examples above used datasets of controlled images taken from environments with ideal weather, lighting conditions and viewing angles. In a real-world setting this is not the case; the weather and lighting cannot always be controlled, and the different viewing angles play an enormous role in the successful detection of a vehicle. When deploying such a system you also need to consider the human factor. The driver's goal may be to avoid camera detection. In this section, several factors that affect the reliability and accuracy of the Computer Vision system will be discussed.

2.2.3.1 Natural

Natural factors like light and weather have a large impact on the reliable performance of the system. These factors cannot be controlled, and the system must be designed in such a way to compensate for these factors.

Weather conditions

Adverse weather conditions such as rain, fog, or snow can significantly impact the performance of computer vision systems, particularly in traffic monitoring and Automatic Number Plate Recognition (ANPR) applications. These conditions introduce visual noise when captured in 2D images, caused by suspended particles forming random patterns. This weather-related noise mimics the disruptive effects of digital or electric noise, making feature extraction and character recognition more challenging [20].

The interference from these patterns complicates the accurate discernment and interpretation of license plate information, reducing system effectiveness. Addressing the impact of weather-induced noise is thus crucial in designing robust computer vision systems for outdoor applications, ensuring reliable performance even in adverse conditions [20][134]

.

Lighting

Lighting conditions play a crucial role in the performance of camera-based systems, especially in scenarios like license plate recognition. Various factors related to light can impact the system's effectiveness:

1. **Glare:** Headlights, sunlight, and wet roads can lead to glare on license plates and vehicle features, potentially hindering the camera system's ability to capture clear images. Glare can obscure important details, making it challenging for the system to accurately identify and recognise vehicles [135].
2. **Night Conditions:** Low-light or night conditions present challenges for camera-based systems. Additional glare from headlights can affect image quality. However, the use of Infrared (IR) light capabilities, as shown in Figure 15, can be instrumental in capturing license plates at night. IR light is not visible to the human eye but can enhance visibility for cameras, enabling them to operate effectively in low-light or nighttime scenarios [135].

In summary, managing and adapting to various lighting conditions is crucial for the optimal performance of camera-based systems, particularly in applications like license plate recognition. Controlling glare, ensuring proper illumination, and leveraging technologies like IR capabilities are strategies employed to address the challenges posed by diverse lighting conditions



Figure 15: Image without(left) and with(right) infrared [78]

2.2.3.2 Operational

Operational considerations can be controlled to improve the performance of the system. For camera-based systems these considerations include the type of image stream, the placement of the cameras, and the image focal point.

Type of image stream

The image stream is largely dependent on the framerate of the camera, while the framerate is based on the vehicle's speed. When setting up the camera for fast moving vehicles, it is necessary to take a stream of images instead of just a snapshot. When a rapidly moving vehicle passes by, capturing a single snapshot is likely to result in a blurred image. For vehicles with speeds lower than 50km/h, the FPS should be set to 5. For vehicles traveling between 50km/h and 100km/h, the FPS should be set to 10. For vehicles traveling faster than that, the FPS should be set to 12. The vehicle speed is directly proportional to the framerate. This allows the recognition engine to have more opportunities to process the vehicle features as it passes a certain point [78].

Placement

The setup and positioning of cameras are critical factors for the effective performance of camera-based systems, especially in applications like Automatic Number Plate Recognition (ANPR). Several considerations should be taken into account:

1. **Lane Coverage:** Ideally, a camera should be set up to cover a single lane. While some cameras may cover multiple lanes, doing so can reduce the resolution per lane, impacting the system's accuracy. Single-lane coverage is generally recommended for optimal results.

2. **Placement Factors:** The placement of cameras is influenced by several factors, such as camera height, road layout, and the positioning of vehicles. Optimal performance for Automatic Number Plate Recognition (ANPR) cameras is generally achieved when they are positioned perpendicular to and at the same level as the vehicle's number plate. While achieving perfect perpendicular angles in most installations may be challenging, software solutions can compensate for any declination or skew in the camera's positioning relative to the number plate when processing raw images [61].
3. **Mounting Angles:** For optimal results, the horizontal mounting angle should be as small as possible. It is recommended to keep the vertical angle within the camera's supplier limits. Figures 16 and 17 below illustrate the ideal horizontal and vertical mounting angles. Additionally, mounting the camera higher than the license plate is advised to minimise glare from headlights [135].

Taking these factors into consideration ensures that the camera is positioned in a way that maximises its effectiveness in capturing clear and accurate images, particularly in challenging scenarios such as low-light conditions, glare, or adverse weather. Proper camera setup is crucial for the successful implementation of camera-based systems like ANPR.




Figure 16: Ideal vertical camera installation [78]




Figure 17: Ideal horizontal camera installation [78]

Focal point

The image focal point of a camera is a crucial aspect that determines the effective area in the scene for extracting vehicle features, particularly in applications like Automatic Number Plate Recognition (ANPR). Considerations for the focal point include:

1. **Focal Point Width:** The width of the focal point is critical. If it is too wide, number plates and features may appear too small for recognition. On the other hand, if it is too narrow, important features may not be captured. Striking the right balance in focal point width is essential for optimal recognition [78].
2. **Visibility Duration:** It's important to ensure that vehicles remain visible within the field of view for an adequate duration. This ensures that the camera system has sufficient time to capture and successfully read all the relevant vehicle attributes. Figure 18 illustrates an example of the focal point of an ANPR camera, emphasizing the importance of capturing features when the vehicle crosses this focal point [78].
3. **Multiple Frames:** In many instances, multiple frames of each vehicle may be needed for accurate identification. If this is the case, it's crucial to ensure that the desired features remain in focus across the entire field of capture. This may involve adjusting the focal point width and positioning to accommodate the variability in vehicle movement and attributes [78].

Achieving an optimal focal plane is vital for the camera system to reliably extract features and ensure accurate identification. Balancing the width of the focal plane, considering the visibility duration, and accounting for the need for multiple frames contribute to the effective operation of camera-based systems, particularly in ANPR applications [78].




Figure 18: Camera focal point [78]

2.2.3.3 Driver

We need to develop a system that can stay reliable even when drivers try to manipulate it. Drivers often go to great lengths to avoid being detected by authorities, especially when they are involved in criminal activities. Drivers intentionally obscure/mask the identities of their vehicles, and very often take on the vehicle identities of law-abiding citizens.

Intentional driver behaviour

- Drivers can make cosmetic and on-purpose changes to a vehicle form which may hide its identity. Criminals can use these vehicles to easily manipulate the detection system when it is classifying the vehicle based on make or model. Many vehicle manufacturing companies use this method to prevent the pre-publication of the looks of vehicles prior to launch. As most vehicle recognition models are trained on thousands of existing vehicle images from different classes, it will be difficult for the model to classify a vehicle with a hidden make or model identity (or a vehicle class it has never seen before) [135].
- Obscuring the number plate can also cause significant flaws in the camera-based detection system:
 - Fixing screws – The position of the fixing screw of the number plate can influence the accuracy of the reading. Experiments done in [18] shows that 72.6 percent of misreads are caused by the screw cap. Figure 19 shows an example of how the fixing screw can influence character recognition. The 'U' may be misinterpreted as a 'B' or a 'E'.



Figure 19: ANPR misread due to screw fixing [18]

- Obscured number plate – Applying adhesives to the plate can affect the retro-reflective properties of the plate. This can cause a glare, which can affect the recognition performance of the camera [18].
 - Marks – As for the case of the screw cap, any marks on the number plate can lead to a false reading. In the study [17], this accounted for 23 percent of the unsuccessful read attempts.
- Number plate migration is a significant problem for vehicle identification systems. It has been reported that drivers can swop number plates more than seven times in a stretch of road to avoid detection.

- Number plate cloning is also widely used by criminals to avoid detection. Usually law abiding citizens also suffer the consequences of cloned number plates. It is the mission of law enforcement systems to correctly identify cloned plates and avoid the harassment of law abiding citizens [106].

Unintentional driver behaviour

- Driver may accidentally modify the form of the vehicle, for example in a car crash. This can make it difficult for the recognition system to identify the vehicle based on body parts. These drivers may unintentionally avoid detection.
- Drivers can obscure the number plate or other vehicle features when they are towing a trailer or another vehicle. The towed object blocks the view of the camera and may cause a failed read. Therefore, care should be taken that camera-based systems have a view of both the front and back of the vehicle to mitigate the effects of an item obscuring the view of the license plate or other vehicle features.

2.3 RFID for Intelligent Transportation Systems

2.3.1 Active or Passive RFID

As discussed in section 2.1.3, active RFID systems use an external source of energy to power tags [94]. However, this extended range may not be suitable for reliable identity verification, as it increases the likelihood of detecting unintended objects. Therefore, passive RFID tags are ideal for identifying vehicles in motion, such as those in traffic.

2.3.2 Tagged vehicles

The RFID tag is a multifaceted device, featuring a transceiver, a programmable memory, and a distinct tag identifier (TID). RFID tags are available in various forms, each tailored to its intended application. Nevertheless, the integration of an RFID tag with a vehicle introduces a set of unique challenges. Standard off-the-shelf RFID tag labels often prove unsuitable for demanding environmental conditions, and their reading performance degrades near metal surfaces. Ensuring consistent and reliable performance in proximity to such surfaces requires the implementation of a specialised antenna and tag design.

2.3.2.1 On vehicle placement

The placement of RFID tags on vehicles for electronic identification plays an important role in the reliability of the system.

On-vehicle tag placements are:

Windscreen – Windscreen tags are usually placed on the inside of the vehicle, protected from environmental conditions and tampering. These tags can only be read by overhead Gantry readers. Certain vehicle windscreen tints shield radio waves, preventing tags from being successfully read. Windscreen tags are not practical for motorcycle use, seeing that most motorcycles do not have windscreens [2].

Headlamp – These tags address the shortcoming of windscreen tags when used on motorcycles. Headlamp tags can be interrogated by Gantry readers. Headlamp tags are also more prone to tampering as the tag is on the outside of the vehicle [2].

Licence plate – RFID tags embedded in license plates provide a long lifecycle and can be easily managed through regulatory license plate life cycle processes. These tags can also easily be adapted to comply with country specific requirements and registration systems. [2]. However, variations in numberplate characteristics and materials necessitate careful consideration during tag installation.

2.3.2.2 Tagged License Plate Types

Several RFID plate configurations exist, each with their respective advantages and disadvantages. This section specifically focusses on the physical number plates themselves. These configuration are listed below:

1. Slotted metal plates

Slotted metal number plates make use of slot antennas for RF communication. These antennas consist of a metal numberplate with one or more slots cut out. When a radio frequency current is applied to the numberplate, the slot radiates electromagnetic waves much like a dipole antenna.

A disadvantage of the slotted metal numberplate is the degraded performance when near metal surfaces. This problem can be overcome by using spacers between the metal surface and numberplate. This means that the reliability of the system is determined by the correct fitment of the numberplate, which is uncontrollable in a worldwide roll-out. One advantage of the slotted metal numberplate is that it cannot be tampered with or removed without visible damage or without compromising its intended functionality. Figure 20 shows an example of the slotted metal numberplate tag.



Figure 20: Slotted metal numberplate tag [131]

2. Hollow plastic

Hollow plastic number plates address the stand-off problem associated with slotted metal number plates. These number plates are also cheaper and easier to manufacture. However, the plastic numberplate is impractical to emboss. Embossing is the method by which symbols, numbers, borders, and letters are depressed or elevated on number plates making them tough to replicate. Embossing is the commonly accepted method of manufacturing a numberplate; therefore plastic number plates will see lower market acceptance [21]. RFID enabled plastic number plates provide a security related

workaround without the need to emboss the numberplate. Figure 21 shows an example of the hollow plastic numberplate tag.



Figure 21: Hollow plastic numberplate tag [130]

3. On metal tags

Mounting or embedding RFID tags on metal can pose challenging. Metal surfaces reflect energy emitted from RFID readers. However, there are specially designed RFID tags that will work around metal. Commercially available on-metal tags are RF inefficient. Tags that provide the same performance as slotted metal number plates and plastic number plates are usually bulky, which pose a problem on number plates with limited space. The tags are also more expensive compared to metal and plastic numberplate tags [21]. Figure 22 shows an example of the on metal numberplate tag.



Figure 22: On metal numberplate tag [132]

2.3.3 RFID readers

The RFID reader is undeniably a pivotal component of any RFID system. The individual reader's performance characteristics wield a more significant influence on the overall system performance compared to the RFID tags. Consequently, in this section, we will explore the components and functionalities of RFID readers.

2.3.3.1 Components of an RFID reader

An RFID reader is a device used to communicate with RFID tags, collect data from them, and transmit that data to a computer or control system.

The RFID reader's microcontroller or processor processes the data received from the tag, which can be used for various purposes, such as tracking inventory, monitoring assets, or recording the location of objects. The processed data is typically transmitted from the RFID reader to an external system through communication ports like Ethernet, USB, RS-232, or Wi-Fi [94].

At a lower level we have the reader's processor, transceiver, and antennae. The transceiver is a critical component of an RFID reader that connects to the antenna. The antenna is responsible for sending out radio frequency (RF) signals and receiving signals from RFID tags. The transceiver, which is short for "transmitter-receiver," converts RF signals into digital data that can be processed by the reader's processor. It acts as an interface between the antenna and the reader's processing components, such as the processor [94].

RFID readers are often accompanied by software that allows users to configure settings, collect data, and integrate the reader into larger systems. The correct configuration of the reader settings is one of the most important parts of any RFID installation [107].

2.3.3.2 Reader setup

By scanning their designated RFID tags, RFID reader units can detect moving vehicles and determine their identities. There are three options for reader placement: overhead, roadside, and in-road.

Overhead - Gantry readers are typically installed on structures spanning the roadway, which requires a substantial infrastructure setup. These readers are typically used for toll collection and access control. Gantry readers can effectively detect tags on vehicle windshields but struggle to read tags on number plates.

Roadside – Roadside readers are typically situated alongside the road, commonly mounted on poles or stands. In contrast to gantries, these readers demand less infrastructure, rendering them a more adaptable solution. While they are capable of detecting both windshield and number plate tags, their coverage is constrained to a single lane due to limitations in distance and potential obstructions from vehicles in distant lanes.

In-road – Embedded directly into the road surface, in-road readers can be strategically placed in every lane, providing an unobstructed view of number plate tags. Additionally, the deployment of in-road readers incurs only a fraction of the costs associated with overhead installations, as indicated by [21]. Consequently, it is evident that in-road reader placement has the potential to address most challenges associated with overhead and roadside placements. Therefore, for the purposes of this study, our exclusive focus will be on in-road RFID readers.

2.3.3.3 Types of antennas

Antennas in RFID systems are broadly categorised into two types: linear polarised and circular polarised. Linear polarised antennas emit radio waves in a single direction, either horizontally or vertically. These antennas are suitable for applications where tags are consistently oriented in a particular direction. For optimal performance, the tag antennas must have the same polarization as the reader antenna. On the other hand, circular polarised antennas emit RF signals in a circular pattern, providing more flexibility in tag orientation. This allows for effective communication even when tags are not consistently oriented. However, one drawback of circular polarization is a reduced reading range, as energy is distributed in various directions. Linear polarised antennas are commonly used when tags have a fixed orientation, ensuring reliable communication. In contrast, circular polarised antennas find applications in scenarios where tags may have variable orientations, such as in retail or logistics settings. Figure 23 illustrates the difference between linear and circular polarization [108].




Figure 23: Linear vs circular polarization [79]

Since tagged number plates typically maintain a fixed orientation when affixed to a vehicle, a linear polarization is the more ideal choice for this application to maximise the achievable read range.

2.3.3.4 Antenna radiation patterns

Antenna radiation patterns describe how an antenna radiates or receives power in relation to direction, playing a crucial role in achieving the desired coverage and performance in RFID or wireless communication systems. The choice of antenna radiation pattern depends on factors such as coverage area, tag orientation, distance, and interference considerations [8].

One prevalent antenna type in RFID systems is the directional antenna, known for focusing its energy in a specific direction, resulting in a high gain in that direction and a narrow beam. This type of antenna is suitable for applications where energy needs to be concentrated in a specific direction [108].

Conversely, an omni-directional antenna radiates equal power in all directions perpendicular to an axis, creating a donut-shaped radiation pattern. Omni-directional antennas are commonly used in applications where coverage in all directions is essential, such as Wi-Fi access points or cell towers.

In the specific in-road RFID application discussed, the goal is to ensure ample tag illumination while effectively managing challenges posed by substantial energy reflections from the underside of vehicles. To achieve this, a directional radiation pattern is chosen to optimise energy emission

toward the license plate tag's travel path. Figure 24 illustrates this desired radiation pattern and provides a comparison to the conventional patch antenna often used in RFID reader systems.

Furthermore, observations indicate that there are typically a maximum of five legally positioned license plate tags within the reading zone at any given moment, and it takes approximately 80 milliseconds to read all five UHF RFID tags. This suggests that tags traveling at 180 km/h require approximately 4 meters of reader illumination to complete the reading cycle. Practical measurements will be conducted to verify the feasibility of achieving this in real-world conditions according to [8].




Figure 24: Desired radiation pattern [8]

2.3.3.5 Reader configuration

Configuring the RFID reader ensures that it operates at its best performance for the specific application. This includes setting the appropriate power levels, sensitivity, and frequency settings to maximise tag read rates and read distances. Properly configuring these settings will ensure minimal interference from other devices or environmental factors and will minimise false tag reads.

The two main reader configuration modes as defined in [107] are:

High throughput mode – High throughput mode refers to a configuration or operating mode designed to maximise the speed of tag data collection and processing. This mode is particularly relevant in applications where a small population of RFID tags need to be read and processed quickly.

Dense reader mode – Dense reader mode ensures minimal interference between two or more readers when deployed in close proximity. These readers seamlessly transition between channels within a specific frequency range and may be obligated to wait for a signal before utilizing a channel. However, it's important to note that while dense reader mode minimises interference, it may impact the overall read rate to achieve this interference-free environment.

In the specific scenario, we anticipate a small population of fast-moving tags passing through the reading zone. While high throughput mode might seem suitable due to the tags' speed, it's essential to note that each lane requires its own reader. As a result, the readers will be in close proximity to one another, potentially causing interference. Through this investigation, we will determine which reader mode best ensures optimal performance in the high-speed traffic environment.

2.3.4 Factors affecting the EVI RFID system

The acceptance of RFID systems hinges on their ability to establish reliable communication with RFID tags. While commercially available RFID systems, including tags, readers, and antennas, undergo calibration in controlled lab settings, real-world scenarios present a host of challenges. These challenges encompass cross-interference and collisions, attenuation issues, and the presence of RF-reflecting objects in the environment. Additionally, factors like the position of the reader in space, as well as the position and orientation of the tag relative to the reader, play crucial roles in system performance [12]. This section will delve into the natural, operational, and driver factors that influence sensor performance and impact the system's ability to detect suspect vehicles.

2.3.4.1 Natural

Natural or environmental factors can impact RFID systems from the outside, and operators have limited control over these influences. Some of these factors include an increase in tag population, the proliferation of other RFID readers, and interference from various RF sources.

Increase in Tag Population:

- UHF RFID tags are designed to be read at ranges exceeding one meter, and the technology is optimised for reading clusters of tags moving past the reader beam simultaneously.
- Challenges arise when RFID tags from different applications overlap, leading to difficulties in distinguishing between them. Issues may include:
 - Some tags not being read or tags from different applications being read.
 - Foreign tags being mistakenly identified as belonging to the system.

Addressing these challenges involves ensuring that each tag contains a unique code associating it with a specific application. Standards like the GS1 Electronic Product Code (EPC) [109] and ISO standards [107], such as the Unique Item Identifier (UII), help uniquely identify tags and manage unwanted tag interference.

Increase in RFID Readers and Other RF Sources:

- RFID readers are intentional emitters, transmitting radio frequency waves to locate corresponding RFID tags. However, this intentional transmission can interfere with other RFID readers in close proximity, affecting reliability [96].
- Various RF sources, such as speed cameras, active RFID tags, and transmitting radios, can also introduce interference, impacting RFID system performance [111].
- The interference from multiple readers and RF sources in close proximity can result in "dead" spots where tags passing through may go undetected [96].

Managing these challenges involves strategic planning, adherence to standards, and implementing measures to mitigate interference. Addressing issues related to increased tag population and RF interference is crucial for maintaining the reliability and effectiveness of RFID systems in diverse environments

To avoid the reader-to-reader interference one can:

- Have different readers operate at different frequencies. Each reader can operate in its own frequency band or a subset of a frequency band, which will eliminate cross interference [111].
- Shield the readers/RF sources from one another to minimise interference. However, it can be difficult to adequately shield the sources from one another and still maintain good performance. Shielding readers can also decrease the reading distance.
- Space the readers far enough from one another/RF sources so that the interference is negligible. This is an easy way to mitigate the effect of interference, but space can sometimes be limited. In a traffic scenario the space is limited to a section of the pavement.
- Control the duty cycle of the readers. Divide the active time between the different readers. When one reader is active, the rest may be inactive; move on to the next reader. However, in a high-speed environment where vehicles can travel up to 180km/h, a reader being inactive for a few milliseconds can cause a vehicle to go undetected [111].
- Adjust the power level of the reader to decrease the reading distance. A reader only transmitting in the required space where the tag will be, can ensure that there is less interference with other readers close by.

Surrounding materials

When applying RFID technology for Electric Vehicle Identification (EVI), the RFID systems encounter various challenges in the presence of metals and liquids. The characteristics of the antenna and the magnetic field distribution can become unpredictable in such environments. Metal surfaces, in particular, have a significant impact on the performance of RFID systems:

- **Effect of Metal Surfaces on Magnetic Field Intensity:**
 - Metals induce magnetic fields that influence the power and directivity of the RF field generated by the RFID antenna [14].
 - Figure 25 illustrates the effect of a metal surface on the magnetic field intensity generated by the antenna of an RFID tag. The metal surface absorbs the energy of the radio frequency field.
- **Directionality and Reader-Tag Distance:**
 - RFID systems are generally directional, meaning that the tag can only be read when the reader and tag are in each other's beam [95].
 - The presence of metals, especially with decreased aperture and altered radiation patterns, makes it less likely for the reader and tag beams to intersect as expected.
 - Closer reader-tag distances can help ensure that the reader and tag beams cross, making in-road antennas more effective than overhead readers in certain scenarios.

Addressing these challenges involves considering the impact of metal surfaces on the RFID system's magnetic field, radiation patterns, and directional characteristics. Understanding these effects is essential for optimizing the performance of RFID systems in environments where metals are prevalent, such as in Electric Vehicle Identification applications




Figure 25: Effects of metal on RFID performance [14]

Multipath effects

Multipath propagation in RFID refers to the phenomenon where radio waves transmitted by the reader reach the tag antenna through multiple paths, each characterised by different amplitudes, phases, and time delays. Various elements, such as the asphalt surface, the tagged vehicle, and nearby objects, contribute to this multipath propagation, leading to interference. The tag

receives power not only from the direct beam but also from reflected waves, introducing additional power. The interference from these reflected waves can affect both the transmitting and backscattered waves. The amount of reflected power is influenced by the relative positioning of the reader and tag. Placing RFID readers in open environments can be beneficial in minimizing interference from nearby objects [113].

2.3.4.2 Operational

Operational factors are factors that can be controlled to improve the performance of the system. One of the main operational considerations is the placement of the sensors. Placement plays a major role in the levels of RF interference. The visibility of other tags in and on the vehicle is also an operational consideration.

Reader placement

The placement of RFID readers on motorways plays an important role in the effectivity, reliability, and acceptability of the system. Placing readers naively across the road could lead to degraded system performance. Interference between readers and crosstalk also have an influence on the placement of readers. Other objects surrounding the readers and tag can also cause significant interference due to reflections. We want as little as possible interference, without compromising the reliability of the system.

Multi-lane reader placement

The naïve placement of readers evenly across the road causes significant interference and flaws in the detection of RFID tags. Placing the readers evenly produces an SINR (signal-to-interference power) of close to 0dB in the centre of the readers, producing 'dead' spots. Cars driving in these 'dead' areas can cause flaws in the detection system. Figure 26 shows the radiation zone of the readers when evenly spaced across the road.




Figure 26: Evenly spaced reader interference [1]

Two alternative approaches can be followed to combat interference and crosstalk:

Spacing the readers in an equilateral triangle shape or a diagonal line eliminates most of the 'dead' spots between the readers. The radiation covers all the critical areas, so drivers cannot avoid detection.

Reader placement along the travel path

Minimum distance

The minimum distance between detection points will determine how close RFID readers can be placed along the travel path. This distance is determined by looking at factors such as interference between different readers and the RFID detection rate.

Maintaining a safe following distance between vehicles is crucial for road safety, allowing drivers adequate time to react to sudden changes. International road safety campaigns often recommend following distances of 1 or 2 seconds. For a highway with a speed limit of 100 km/h (28 m/s), this translates to a minimum vehicle spacing of 28 meters.

Research conducted in [8] indicates that with a vehicle spacing of around 28 meters, it is possible to observe two tags or fewer within the desired read zone. In congested conditions, for travel speeds ranging from 60 km/h to 110 km/h, the worst-case scenario involves a vehicle distribution with the densest compaction of vehicle plate tags in a 10-meter stretch (for a three-lane road), resulting in up to 5 tags.

To verifiably identify a UHF RFID tag according to the ISO18000-63 standard [107], a minimum of 10.2 milliseconds is required. Reading 5 tags necessitates a minimum of 51 milliseconds, excluding the dead time between different tag reads. In practical setups, this read time can increase to approximately 80 milliseconds. Consequently, tags traveling at 180 km/h require at least 4 meters to complete the reading cycle [8]. For RFID detection points with read distances of 4m, the distance between two readers along the travel path should therefore not be less than 4m to avoid crosstalk and reader-to-reader interference.




Figure 27: Effective reading zone [69]

Figure 27 displays an example of the in-road reader setup. The white dots in the centre of the road represents the RFID readers, while the red indicates the ground plane surrounding the readers. The additional 1 meter red zone is proposed to accommodate for the blinding effect caused by the metal structure of the vehicle.

Additionally, the mode of an RFID reader directly affects how closely multiple RFID readers can be placed near each other without interference. This is primarily because these modes influence the frequency-hopping behaviour and the way readers handle interference mitigation [107]. The Dense reader mode can enable the readers to be placed closer to each other than with the High throughput mode, however at the cost of tag read rate. Later experiments in Chapter 6 will reveal what the optimal reader mode is.

Maximum distance

Currently SANRAL e-tolling incorporates 42 directional toll gantries based on ANPR systems at approximately 10 km spacing. The gantry locations are depicted in Figure 28 below.

To utilise the full potential of EVI for intelligent transportation, ANPR camera systems will be combined with in-road sensors. This means that in-road readers will have to be placed in the viewing angle of the ANPR cameras. So a spacing of at least 10 km is necessary to achieve reliable fusion with the ANPR systems currently installed on road networks. This maximum distance also influences the decision between the two reader mode. With a 10km spacing between stations, it does not introduce a requirement to use Dense reader mode.




Figure 28: SANRAL Gantry locations [80]

Visibility of other tags

In 2021 28 billion RAIN chips were produced. The tags are not only used on vehicle bodies, but on a lot of other objects like vehicle tyres, wooden pallets, containers, parcels, and clothing. These objects are often transported in vehicles and can be detected by the RAIN readers. An increase in the number of tags detected by the reader can increase the tag inventory/access time and may force the reader to only read part of the available data or may even cause a failed read.

Other than flooding the in-lane RFID readers, these tags can also cause interference in neighbouring tags. These stray tags (not part of the vehicle identification) can lead to performance issues for the application tags and can also cause failed reads.

An experiment conducted by [11] revealed fluctuations in the RSSI when a neighbouring tag is introduced. It was found that adjacent tags cause an increase or decrease in the backscattering communication budget. Finding the ideal spacing from tag-to-tag is essential to combat crosstalk and interference.

Tag posture

The performance of the system is notably affected by the orientation of the RFID tag concerning the reader. Ideally, a tag should possess orientation invariance, guaranteeing readability regardless of its orientation. Concentrating the tag's orientation in a particular direction can improve its performance, but in the worst-case scenario, the reader may not receive any response from the tag. To optimise the radiated power from the reader, tags should be oriented to align with the reader's radiation pattern. As suggested in [2], license plates with RFID tags should be mounted with no more than a 10-degree offset from the vertical.

2.3.5 Tag information

2.3.5.1 Tag data types

The GS1 RFID tag protocol typically implements two main commands, which are INVENTORY and ACCESS [109]. INVENTORY is used to identify individual tags, whereas ACCESS is used to communicate with an identified tag. There are two different tag data types, each used for a specific application and allowing different functionalities:

- Identification tag – It is 96-bit or larger and only implements tag INVENTORY. The tag uses a standard numbering system, which reduces its dependency on real-time cloud data [109].
- Information tag – It is a 96-bit or larger tag that implements tag INVENTORY and ACCESS. Simple sensor data and additional object identification data is recommended to be part of tag INVENTORY data. Additional object and integrity data is recommended to be accessed after INVENTORY. Tag data ACCESS is required to achieve privacy through on-chip cryptography [109].

For the purpose of EVI, additional information needs to be stored on a tag to enable offline identification and verification; however, storing additional information on the tag can lead to some technical hurdles.

2.3.5.2 Problem areas

The verification of a vehicle's authenticity is crucial for various reasons, and counterfeit or duplicate number plates have become a significant concern. Current methods for verifying vehicle identity often involve static checks, which can disrupt traffic flow. To address this, there is a need for a secure, in-motion vehicle identification (EVI) technique based on a unique identifier.

RFID technology offers a reliable solution for identifying vehicles in motion, but implementing such a system requires a reader to interrogate a tag within a minimal time interval. While reading only the Tag Identifier (TID) may be suitable for practical applications, additional security measures, such as a digital signature, are recommended for enhanced authentication [19].

A hurdle in employing RFID for Electric Vehicle Identification (EVI) lies in the restricted data capacity of tags and the imperative speed of reader interrogation. Tags adhering to the ISO18000-63 standard [107] generally store a Unique Item Identifier (UII) and 512 bits of data. To overcome this constraint, the adoption of Elliptic Curve Cryptography for the digital signature proves beneficial in substantially reducing the signature size without compromising cryptographic robustness. This adaptation facilitates more efficient tag reading at elevated speeds, particularly in typical traffic conditions where quick reading of number plates is essential [19].

2.3.5.3 Tag data structure

Tag singularity

To successfully distinguish between several tags, a unique code must be assigned to each tag. A unique factory code or UII is assigned to each tag and stored in non-programable memory allowing several tags to be differentiated from one another. This also allows tags from other applications to be filtered out.

Vehicle identification

Combining Computer Vision (CV) systems with Radio-Frequency Identification (RFID) technology can enhance the reliability and accuracy of vehicle identification. While CV systems can accurately detect details such as the model, make, colour, and registration number of a vehicle, RFID number plate tags can store this information when issued to a client. This integration allows for a comparison between imaging and RFID data, providing a more robust method of identification.

To counter cloning of tagged number plates, it is essential to store comprehensive vehicle information in the RFID tag, including the VIN (Vehicle Identification Number), engine number, and chassis number. This additional information helps prevent criminals from swapping number plates between identical-looking vehicles. Manual checks at roadblocks or ad hoc inspections can then easily detect cloned plates by verifying these additional numbers [19]

Security

To enhance the security of RFID-tagged number plates and make it more challenging for fraudsters to clone legal plates, digital signatures can be employed. The digital signatures incorporate encrypted versions of information that uniquely identifies the vehicle, including details such as the registration number, issuer, end user, and possibly the VIN (Vehicle Identification Number). The use of digital signatures adds an extra layer of security, making it nearly impossible for criminals to generate a matching signature for a cloned plate.

The authenticity of the number plate and its associated digital signature can be verified by comparing the human-readable contents with the decrypted signature. This security measure becomes more effective when plates are not cloned for the same make, model, and colour. Maintaining a record of digital signatures allows officials to trace and identify the origin of legally or illegally issued plates. Detection of illegal plate embossers can lead to legal action against them [19].

2.3.5.4 Tag memory capacity

Passive RFID tags have very limited memory capacity, ranging from 64 bits to 1 kilobyte [133]. Therefore special attention needs to be paid to the encoding of the RFID tag. From the report in [2], UHF RAIN enabled tags should use the ISO/IEC 20248 [110] in a 640 bit tag. It is recommended that the tag structure is as follows:

- 1) A unique factory identifier – 96 bits
- 2) Tag data – 224 bits
 - Jurisdiction and data structure identifier – 40 bits
 - Expiry information assuming a 26 year lifecycle – 8 bits
 - License plate number – 42 bits for 8 characters
 - Vehicle visual characteristics (colour, model, make, size) – 134 bits
- 3) Tag signature – 256 bits using ECC BLS with point reduction

2.3.6 Market acceptance

When weighing the advantages and disadvantages of the current commercial RFID plates, it appears that selling these plates are challenging. Cost is one of the large role players, as well as the time it takes to fit all the vehicle with new number plates and to deploy the RFID readers. The following may be some reasons why sales may be challenging:

Plate roll-out

The biggest problem with RFID for EVI is that it takes time to fit all the vehicles and to deploy the readers. In a regulatory environment, this time is typically 5 years or longer. The general issue plates have a lifetime equal to the vehicle lifetime, which is approximately 11 years. A changeover will cause double plate volumes for at least five years [1].

Cost impact

One of the main reasons why RFID plates have not been globally accepted is the additional costs that may be incurred when deploying such a system. RFID enabled plates require additional infrastructure which is expensive to deploy and difficult maintain.

In-road:

- In-road installation cost for the fleetVALID project on the University Campus, excluding maintenance and decommissioning: €15,000 (AU\$20,000) per lane. This was a first-time installation, and the cost should reduce when the process becomes standard [65].
- The maintenance and inspection cost for in-road readers are part of current road inspections and maintenance i.e. the same team that maintains the road infrastructure can also maintain the in-road readers. This means a yearly maintenance cost of less than 2% of the installation cost [65].
- Although the tag on a RFID enabled number plate costs only in the range of a few cents (US), the process to integrate these tags in a number plate will increase the cost. The Tönnjes IDePLATE tags are less than €10 [21].

Overhead Gantry:

- European four lane gantry deployment costs can be up to €40,000 per lane, while a similar structure in the US would cost up to €200,000 per lane [64].
- The yearly inspection and maintenance cost is 15% to 20% of the installation cost, which would equate to around €20,000 per year [65] [21].
- The deployment duration of gantry structures is three to six months with partial road closure [65].

Thus, the total cost of deploying an in-road system is significantly cheaper compared to an overhead gantry structure. The table below shows a simplified comparison between in-road and overhead setups.

Table IV: In-Road and Overhead Comparison

Item	Overhead	In-road
Direct costs	€40,000 to ≥€120,000 per lane	≤€15,000 per lane
Cost of ownership	€120,000 to ≥€200,000 per lane	≤€20,000 per lane
Maintenance	10% to 25% yearly of installation costs	≤2% yearly of installation costs
Lifespan	Typically 10 years	≥15 years
Location Limitations	High	Almost none
Decommissioning costs	≥€100,000	NULL
Construction work complexity	High	Low
Deployment duration	3 to 6 months	1 week
Traffic impact	3 to 6 months	1 week
Road closure	Full road 4 to 8 hours	Less than 2 hours per lane

2.3.7 Combining RFID and Computer vision for EVI

Computer vision stands as a potent tool employed by authorities for monitoring vehicle activities on roads, serving the dual purposes of traffic management and law enforcement. While computer vision systems exhibit high accuracy in ideal conditions, real-world environments pose challenges that can significantly impact their precision. The integration of RFID and computer vision technologies for two-factor authentication presents the potential for more reliable outcomes [18].

In a system proposed in [19], the approach relies on a combination of fixed RFID reader infrastructure and existing Automatic Number Plate Recognition (ANPR) or traffic camera systems. This integrated system aims to automatically identify vehicle number plates and other relevant features, leveraging the strengths of both RFID and computer vision technologies. The synergy between these technologies offers several benefits:

- Cloning of number plates become a formidable challenge for criminals, as generating a signature matching the factory ID becomes nearly impossible. Legally produced RFID tags are intricately matched with the legal vehicle's make, model, and colour, rendering cloning a costly endeavour. By utilizing RFID tags in conjunction with ANPR camera systems, the proposed system verifies tagged vehicles, detecting any attempt to use cloned number plates by cross-referencing the information stored on the tag with the captured ANPR image [19].

- Detection of tag removal is enhanced through the combined use of ANPR and RFID systems. This synergy enables the identification of vehicles with untagged number plates, allowing law enforcement officials to stop such vehicles for manual inspections [19].
- Addressing tag migration challenges, if a criminal attempts to remove the tag from a plate, the unique Tag Identifier (TID) complicates the exchange of RFID tags between two number plates. The distinctive factory tag ID does not align with the registration number on the illicit new plate. Digital signatures on tagged plates play a crucial role in detecting situations where a tag does not match the designated license plate [19].

2.4 Chapter conclusion

In summary, this chapter has undertaken a thorough exploration of the theoretical underpinnings and prior research studies pertaining to the envisaged in-traffic vehicle identification system. The examination revealed that both RFID and Computer Vision emerge as viable solutions for vehicle identification, with the prospect that their combined utilization could enhance overall system performance. By delving deeper into the theoretical landscape, the objective was to foster a nuanced comprehension of the problem's background and uncover potential pathways for its resolution. As we progress, the theoretical insights garnered from this literature review assume a pivotal role. They will function as a guiding framework, shaping the trajectory of this research in the subsequent design, implementation, and testing phases.

3 Preliminary System Design

Following the comprehensive exploration of relevant literature in the preceding chapter, the attention now shifts to the critical phase of designing the in-traffic vehicle identification system. This chapter lays the groundwork for the practical implementation and testing phases, providing the building blocks for our research efforts. In this chapter, the system architecture and the foundational framework for the detection system is explored. In the first section, we elaborate on the process of transforming theory into a conceptual design. Following that, we delve into the examination of vital elements, encompassing the selection of hardware and software during the preliminary design phase.

3.1 System requirements

The following section provides the functional, performance, physical, interface, environmental, design, legal, and ergonomic requirements of the proposed system.

3.1.1 Functional requirements

- 3.1.1.1 The system must process data for multiple vehicles simultaneously. RFID readers should identify multiple tags in the reading field, and ANPR cameras must extract information from recorded footage for all vehicles.
- 3.1.1.2 RFID shall store essential vehicle characteristics, including details verifiable by the computer vision system.
- 3.1.1.3 The vehicle identification information stored on the RFID tag shall not exceed a total of 640 bits in size
- 3.1.1.4 RFID tags shall include unique vehicle identifiers to detect discrepancies between tag data and ANPR data.
- 3.1.1.5 RFID readers shall reliably detect tags mounted at heights ranging from 0.2 m to 0.6 m from the road surface.
- 3.1.1.6 RFID tags must include a digital signature for data verification.
- 3.1.1.7 RFID and ANPR systems must be time-synchronised to correlate vehicle and tag data
- 3.1.1.8 A classification index for vehicle make, model, and colour should be used to encode details compactly within RFID tag memory.

3.1.2 Performance requirements

- 3.1.2.1 The RFID reader must support a 4-meter horizontal range to reliably detect vehicles traveling at speeds up to 180 km/h
- 3.1.2.2 The system shall process RFID tag data and ANPR video streams in less than 10 seconds
- 3.1.2.3 The system shall detect and record RFID tags from multiple vehicles simultaneously.

3.1.3 Physical requirements

- 3.1.3.1 The system shall fit within a 113 mm road cavity without exceeding the permissible height above the road.
- 3.1.3.2 Components shall be resistant to moisture, vibration, and road impact.
- 3.1.3.3 The system shall be housed in weatherproof enclosures with an IP67 rating or higher.
- 3.1.3.4 The system shall be designed to operate with low power consumption to maximise energy efficiency.
- 3.1.3.5 The system shall be easily installable and integrable with existing infrastructure.

3.1.4 Interface requirements

3.1.4.1 The RFID reader must interface with the external controller via a serial interface.

3.1.4.2 The RFID reader must connect to the antenna using an SMA connector.

3.1.5 Environmental requirements

3.1.5.1 Components shall operate between -40°C and 60°C for the RFID subsystem and -10°C to 50°C for the ANPR camera.

3.1.5.2 The system shall be resilient to electromagnetic interference (EMI).

3.1.6 Design requirements

3.1.6.1 The system must be easy to deploy and integrate into edge controller with minimal setup.

3.1.6.2 The classification system must allow for easy updates to accommodate new vehicle types or features

3.1.6.3 All components, including the RFID reader and antenna, must be easily replaceable

3.1.6.4 The system shall provide easy access for routine maintenance.

3.1.7 Legal requirements

3.1.7.1 RF output must comply with regulations limiting emissions to below 500 $\mu\text{V}/\text{m}$ at 30 m.

3.1.7.2 The antenna must comply with regulations on maximum permissible height for road-mounted objects.

3.2 Context diagram

A context diagram illustrates how the system integrates with its operational surroundings. As seen in Figure 29, the diagram provides a clear representation of the system's interactions with various external components. These components encompass a range of external elements, including vehicles, road infrastructure, and the physical environment, each facilitating the exchange of specific information with the vehicle identification system. Consequently, it's crucial to acknowledge that all these elements collectively influence the system's performance and must be considered during the system design.




Figure 29: System context diagram

3.3 Functional flow

A functional flow diagram depicts how the system shall function in its operating environment. A Functional flow breaks down the complex functions of a system into smaller manageable tasks, which can easily be analysed and allocated to specific resources. The flow diagrams typically begin at a broad, high-level overview and subsequently progress to more detailed and granular levels of analysis.

3.3.1 Level 0

Figure 30 shows the lifecycle of the vehicle identification system. The functional flow starts with the installation of the system, followed by the normal operating loop. In case of scheduled or unforeseen maintenance requirements, the operation temporarily halts to allow for necessary maintenance tasks. When the system approaches the end of its life cycle, it is uninstalled and replaced as part of its routine life cycle management.




Figure 30: System functional flow

3.3.2 Level 1

The Level 1 functional flow further elaborates on F/U 3, which pertains to the system's operation. This operational context centres on the core functions of the system within its environment. To begin, the system's startup sequence ensures that all components are operational, and all user configurations are correctly applied. Subsequently, the system enters into a continuous cycle where it captures traffic video footage and collects tagged vehicle identification data. Within this cycle, the system also extracts vehicle identification features from the received footage and processes the data from the read tags. The two streams of data are then compared to reveal the actual identities of the vehicles. In the event of a discrepancy in the data, a vehicle is marked as suspicious, enabling authorities to take further action against the potentially suspect vehicle. Figure 31 displays this in a diagram.




Figure 31: Operational context

3.3.3 Level 2

To gain a better understanding of the system's functionality, it is essential to direct attention towards specific functional units within the operational context. Figure 32 provides a detailed expansion of F/U 3.3. The initial stage in extracting vehicle features involves localizing the vehicles within the input video stream. Following this, the computer vision pipeline utilises the extracted vehicle images to execute ANPR and vehicle make and model recognition. These two feature extraction processes collectively yield a distinct visual identifier for each vehicle passing through the detection point.




Figure 32: Extract vehicle features context

Another critical functional unit that demands further investigation is F/U 3.4, as depicted in Figure 33. To begin, the reader emits a signal to scan for nearby tags, and only those with application-specific IDs respond. Subsequently, the reader gains access to specific tag memory banks, which securely store crucial vehicle identification information, including vehicle registration numbers and make and model details, using specialised encryption techniques. Once the vehicle identification data is successfully retrieved, the reader relays this data to a connected roadside unit for additional processing and cross-referencing with visually extracted data.




Figure 33: Capture tag data context

The functional flow section has provided a good understanding of the system's operation. These diagrams will serve as guiding blueprint in the transition from theory to practical implementation.

3.4 System architecture

The system architecture translates the functional components from the previous section into a concrete framework, selecting the necessary hardware, software, and interfaces to bring the system to life. Figure 34 depicts the high-level view of the system and its interfaces to external components.




Figure 34: High level architecture

To gain a more profound insight into the system's makeup, it's essential to further break down F/U 1.0. Figure 35 offers a comprehensive elaboration of the system, providing specifics on the individual components housed within the system boundary. The system comprises three distinct units: an edge controller unit, a traffic camera unit, and an in-road RFID unit. The traffic camera and RFID sub-systems are managed by the primary edge controller, tasked with collecting data streams from both these sub-systems. Subsequently, it performs data comparisons and communicates the verification results to an external operator.




Figure 35: Full system architecture

The system architecture diagram acts as a visual roadmap for the upcoming development phases, guiding us in the selection of diverse hardware components that align with the distinct needs of the sub-systems.

3.5 Component Selection

In this section, hardware components essential for fulfilling the unique requirements of each sub-system, are explored. In this thesis, specific focus will be placed on the hardware design of the in-road RFID reader. Please note that the hardware design of the main edge controller and traffic cameras falls beyond the scope of this study.

3.5.1 Edge controller

For the purpose of this study, a trade-off analysis of potential edge controllers will not be conducted. Instead, a PC will be used as the primary processing unit to simulate the functionality of the edge controller. This approach is intended to focus on validating the integration and data processing capabilities of the system without introducing the complexities of hardware-specific constraints. While the PC does not represent the final deployment hardware, it provides a platform to test the system's performance. Future studies may address the selection of an appropriate edge controller. Specifications of the PC can be found in section 5.2.1.1.

3.5.2 RFID reader

3.5.2.1 Description

At the core of the RFID system lies the RFID reader, a crucial component responsible for facilitating seamless information exchange between the tags, the reader itself, and the user interface. This reader possesses the capability to receive commands through the user interface and can provide the necessary power to RFID tags for efficient data exchange. RFID readers are typically categorised into two main types: integrated and stand-alone. In the integrated configuration, the RFID reader and antenna are combined into a single device, while stand-alone readers operate independently from the antennas. It's worth noting that most off-the-shelf integrated readers are bulky and equipped with standard RFID antennas, which do not align with the specific requirements of the application. Consequently, stand-alone readers, designed to be compact and embeddable in the pavement, are the prime focus to meet the unique demands of the project.

3.5.2.2 Selection criteria

In Table V below, a comparative analysis of various RFID reader types based on several critical parameters is done. The first criterion is the physical size of the RFID reader. This is of particular importance due to the constrained cavity within which the reader must be seamlessly integrated. The reader's size must strike a balance between being compact enough to fit inside the restricted cavity space while still delivering ample power to consistently read tags from considerable distances.

The second crucial parameter to consider is the read rate. In high-speed traffic environments, it is imperative that the reader can swiftly and accurately capture tag data [85]. Therefore, the reader's ability to read tags at a high rate is a significant factor.

Thirdly, the cost of the reader is a vital aspect. Cost plays a pivotal role in the nationwide deployment of this technology. The objective is to minimise costs while maintaining optimal functionality.

Lastly, the operating temperature is assessed. The reader must be capable of functioning effectively in high-temperature conditions, reaching up to 60 °C [116]. This is a critical factor in ensuring the reader's durability and reliability in diverse environmental settings.

Table V: Comparison between different RFID readers

Sensors	Selection Criteria				
	Size (L x B x H)(mm)	Read rate (tags/sec)	Cost (Rand)	Read distance (m)	Operating temperature (°C)
Lepton7x4 4-port RAIN RFID reader [117]	81 x 76 x 18	950	10131	5	-20 to +70
Sparkfun Simultaneous RFID reader [118]	79 x 54 x 12	200	6166	4.5	-40 to +60
ThingMagic M6E micro [114]	69 x 43 x 15	750	10669	9	-40 to +60
TSL 3417 [119]	68 x 42 x 6.7	1300	29587	5	-40 to +60

The table presented above offers a comprehensive comparison of various RFID readers tailored to the specific application. As we can note, the ThingMagic M6E micro emerges as the most suitable choice. The form factor of this RFID reader stands out as one of the most compact among the listed devices, making it well-suited for installation within the pavement. Its tag read rate is marginally lower compared to some of the other options; however, this shortfall is compensated for by its superior read range, which nearly doubles that of its counterparts. Its theoretical read range is 9 meters, though practical ranges may vary depending on the setup, antenna selection, and environmental conditions.

One more standout feature is the reader's wide operating temperature range, rendering it suitable for the demands of a traffic environment. Additionally, when considering cost, the ThingMagic M6E micro falls within the average price range, aligning well with its specifications when compared to the other readers.

3.5.2.3 Final selection

The RFID reader chosen for this application is the ThingMagic M6E Micro [114], selected for its compact form factor, which makes it well-suited for integration into space-constrained

environments. With a read range of up to 9 meters and the ability to process up to 750 tags per second under high-performance conditions, it meets the reliability and performance requirements of this system effectively.

Figure 36 present an image of the reader module mounted on an interfacing board.



Figure 36: ThingMagic M6E micro RFID reader module [115]

The table below provides some specification of the ThingMagic reader.

Table VI: ThingMagic M6E micro specifications

Specification	Value
Data interface	UART/ USB 2.0
Supply voltage	3.5V to 5.25V
Power consumption	5.5W @ 30 dBm, 2W @ 0 dBm
RFID protocol	ISO 18000-63
Power output	-10 dBm to +30 dBm

3.5.3 RFID tag

3.5.3.1 Description

For the specific context of RFID-enabled number plates, where the goal is to achieve a discreet, concealed placement within the number plate, passive RFID tags are the natural choice based on research in the previous chapter.

Moreover, when the objective is to accurately identify swiftly moving vehicles within a range of up to 10 meters, UHF (Ultra High Frequency) tags are the superior choice compared to HF (High Frequency) or LF (Low Frequency) tags. These UHF tags offer an extended reading range, well-suited to this application. In this section, the RFID tags deployed utilise RAIN RFID UHF technology to deliver optimal performance.

3.5.3.2 Selection criteria

Table VII below compares several commercially available license plate tags according to three characteristics.

The first characteristic is the read range of the tag. As we have seen from literature, the read range of a tag is vitally important as it allows the RFID system to cover a larger area, ensuring that tagged vehicles are reliably detected [8]. A longer read range can also increase the efficiency and speed of RFID systems, which is crucial in a high-speed traffic environment.

The second key consideration is cost. In the case of a nationwide deployment where every vehicle must be equipped with an RFID license plate, affordability is a critical factor. It's important to recognise that many general users may have budget constraints and may not be willing or able to invest in expensive license plates. Therefore, the lower the cost of these RFID license plates, the greater the likelihood of widespread public acceptance and adoption of the new system.

Finally, the packaging type is an important consideration. It's essential that the end user can effortlessly affix the RFID tag to the vehicle. Ideally, the tag should be discreet and concealed, leaving no opportunity for unauthorised removal. Additionally, such a design enhances protection against external damage, ensuring the tag's durability.

Table VII: Comparison between RFID tags

Sensors	Read range (m)	Cost (Rand)	Package
TransCore eGo [120]	9.6	425.62	Plate mountable tag
TÖNNJES IDePLATE [81]	18	160	License plate
Omni-ID Exo 750 [121]	7	158	Plate mountable tag
Confidex Steelwave Micro II [122]	5	561	Plate mountable tag

Based on the information provided in the table above, the TÖNNJES IDePLATE stands out as the superior choice. The IDePLATE is an integrated license plate that incorporates RFID technology, employing a proprietary design. This unique feature allows the license plate to fulfill regulatory vehicle marking requirements and be readily detected by an RFID reader, eliminating the need for additional devices to be mounted onto non-tagged license plates.

Furthermore, the TÖNNJES IDePLATE boasts the longest read range when compared to the other tags, making it well-suited for vehicle identification applications. Cost-wise, these tag integrated

license plates are competitively priced, and the company even offers free samples for those interested in conducting pilot studies.

3.5.3.3 Final selection

The chosen component is the TÖNNJES IDePLATE displayed in Figure 37. These numberplate tags employ RAIN RFID technology, enabling electronic vehicle identification. Each numberplate tag is secured by a distinctive embedded identification code, referred to as the UII (Unique Item Identifier), which can be cross-referenced with a database for authentication purposes. The proprietary nature of the IDePLATE warrants attention, as it carries implications for the practical adoption and scalability of the system. While the IDePLATE serves as an effective test artifact to validate the system concepts, its proprietary design introduces potential limitations that could affect wider implementation. These limitations include dependencies on specific manufacturers, potential restrictions related to intellectual property, and the associated costs of licensing or adopting proprietary technology.

3.5.3.4 Technical information



Figure 37: TÖNNJES IDePLATE [81]

The table below provides some specification of the IDePLATE.

Table VIII: Toenjes IDePLATE specifications

Specification	Value
Memory capacity	3520 bits
Security features	32-bit kill and access passwords
Data retention on chip	20 years
RFID protocol	ISO 18000-63
Operating frequency	860-960 MHz

3.5.4 RFID reader antenna

3.5.4.1 Description

RFID reader antennas play a crucial role in RFID systems by transforming the RFID reader signal into RF waves that RFID tags can detect. In the absence of an RFID antenna, communication between the RFID reader and RFID tags is not possible. Two main types of RFID antennas exist:

integrated and standalone. For the chosen application, which involves a standalone RFID reader, it is essential to use a standalone antenna to facilitate communication.

In addition, RFID antennas come in two primary polarization types: linear and circular. Literature has indicated that linear polarization is better suited for tags with a fixed orientation and offers an extended reader range compared to circular polarization. Given that number plates, when attached to a vehicle, typically maintain a constant orientation, opting for a linearly polarised antenna is the optimal choice for this application, ensuring the maximum achievable read range.

3.5.4.2 Selection criteria

Table IX below compares various types of available antennas according to several characteristics.

From the literature discussion in section 2.3.3.3, a specific focus will be placed on linearly polarised antennas. The initial criterion to consider is the antenna's maximum read distance. This parameter directly influences the dimensions of the dependable tag read zone. Despite the selected RFID reader boasting a 9-meter read range, the ultimate coverage distance is also contingent upon the antenna's read range.

The second factor of concern is cost. An elevated cost could impede the feasibility of nationwide deployment. Therefore, the objective is to maintain costs at an optimal level without compromising the antenna's performance. It is worth noting that custom-designed solutions often incur higher expenses than off-the-shelf alternatives. Achieving the desired radiation pattern may necessitate a custom antenna design.

The third parameter to evaluate is the IP rating, a vital indicator of the system's resilience to environmental conditions. A higher IP rating signifies enhanced resistance to environmental factors, bolstering the system's durability [1].

Lastly, the antenna's radiation pattern requires attention. As observed in Figure 24, a specialised radiation pattern is indispensable to facilitate effective number plate tag detection. Specifically, a radiation pattern close to the horizontal plane is required. Conventional RFID antennas typically focus their radiation pattern in a specific direction rather than distributing it broadly in proximity to the reader [8].

Table IX: Comparison between RFID antennas

Sensors	Cost (Rand)	IP rating	Radiation pattern (max gain relative to the pavement)
RCSensor [69]	Not published	IP 67	32° from the horizontal

Times-7 RFID antenna [123]	42160	IP 66	Not published
Impinj threshold [124]	9793	IP 54	60° from the horizontal

Referring to the data presented in the table above, it is evident that, among the three antennas (with the RCSensor being a custom design by Licensys), the RCSensor outperforms in key aspects such as IP rating, and radiation pattern.

Furthermore, it is worth noting that the radiation pattern requirements for the in-road application are not met by any of the other antennas. Most of these antennas exhibit radiation patterns that are closer to the vertical orientation, making them less dependable for the reliable detection of number plate tags.

Licensys developed the RCSensor antenna as a solution to address these specific needs, with a primary focus on being embedded in the road to read tagged number plates. This specialised antenna generates a dropped toroid radiation pattern and maintains a low profile, measuring under 22 millimeter in height. Furthermore, the antenna is durable and can withstand the weight of most vehicles. Consequently, the optimal choice for the RFID EVI application will be to procure the RCSensor, custom-built for this purpose.

3.5.4.3 Final selection

The chosen RFID antenna solution is the RCSensor, a product developed by Licensys. The RCSensor is a standalone antenna designed for use alongside a separate RFID reader. With its low profile, it can be embedded in the road surface without causing any disruption to traffic flow. The RCSensor has a radiation pattern displayed in Figure 24, which is well suited for detection of incoming tagged license plates. It also has a robust construction, with a thick aluminum plate that ensures structural integrity and serves as an efficient heat radiator into the pavement. With a designed lifespan of 10 years, the antenna is built to endure harsh weather conditions, making it a reliable choice for the application.

3.5.4.4 Technical information

Figure 38 shows an image of the RCSensor.




Figure 38: RCSensor developed by LicenSys [69]

The table below provides some specification of the RCSensor.

Table X: LicenSys RCSensor specifications

Specification	Value
Return Loss	≥ 23 dB
Dimensions	200 mm x 300 mm x 30 mm
Gain	4 dBiL
Polarization	Linear
Center frequency	917 MHz
Lifespan	10 years

3.5.5 CV Camera

3.5.5.1 Description

The CV camera is the hardware component used for capturing traffic video footage. The video data is then received and processed by a roadside controller, where certain vehicle features are extracted using computer vision algorithms.

3.5.5.2 Selection criteria

Table XI below provides a comparison of various types of computer vision (CV) cameras based on several key criteria.

The first criterion is camera resolution. According to [78], many unsuccessful ANPR attempts result from blurry images, which can largely be attributed to inadequate camera resolution. Higher resolution enables the camera to capture greater detail and clarity in an image, significantly improving ANPR performance.

The second criterion is frame rate, a crucial factor for capturing fast-moving objects. A higher frame rate allows the camera to record more individual images per second, resulting in a clearer and less distorted representation of motion. For vehicles traveling at high speeds, accurate detail capture requires processing more frames. According to [78], an effective frame rate of at least 12 FPS is necessary for vehicles moving at speeds exceeding 96 km/h.

The third criterion is the lens type. Cameras equipped with a varifocal lens offer the flexibility to adjust the focal length, enabling changes to the field of view. In contrast, a fixed lens has a set focal length, providing a single, predetermined viewing angle. Given that cameras may be installed in diverse traffic scenarios, a varifocal lens is advantageous as it allows greater adaptability for capturing footage effectively.

The final criterion is cost. High costs can hinder the feasibility of large-scale, nationwide deployment. Therefore, it is essential to strike a balance between maintaining optimal camera performance and ensuring affordability to support widespread adoption.

Table XI: Comparison between CV cameras

Cameras	Resolution (MP)	Frame rate (FPS)	Lens	Cost (Rand)
Hikvision 2MP 6MM Acusense IP Bullet Camera	2	Up to 30	Fixed focal lens	2000
HIKVISION 4MP IP Bullet Camera 6MM Acusense	4	Up to 30	Fixed focal lens	3586

HiLook 6MP Dual Light MD 2.0 IP Bullet Network Camera	6	20	Varifocal lens	3000
---	---	----	----------------	------

Referring to the data presented in Table XI, the HiLook 6MP camera emerges as the most suitable option, primarily due to its superior resolution. While its FPS is lower compared to the other two alternatives, this is less critical than achieving a high resolution for capturing clear and detailed images. Additionally, the HiLook camera features a varifocal lens, offering flexibility in adjusting the field of view to suit different scenarios. Furthermore, it is competitively priced, making it a cost-effective choice compared to the other options.

3.5.5.3 Final selection

The chosen camera is the HiLook 6MP Dual Light MD 2.0 IP Bullet Network Camera. The camera offers a good resolution and frame rate, suitable for capturing footage of moving vehicles at high speeds. The additional feature of a varifocal lens will assist greatly in setting up the camera with the correct field of view and optical zoom for testing scenario.

Figure 39 provides an image of the HiLook camera.



Figure 39: HiLook IP bullet camera [143]

3.6 Chapter conclusion

In conclusion, this chapter has played a pivotal role in bridging the theoretical exploration of relevant literature with the imminent practical phases of implementation and testing. As we transition from theory to practice, the groundwork laid in this chapter serves as the cornerstone for the research endeavours. The chapter delved into the core aspects of designing the in-traffic vehicle identification system, establishing the foundational framework essential for subsequent phases. The detailed examination of vital elements, including the judicious selection of hardware and software, marked the preliminary design phase's critical contributions.

4 Initial RFID testing and experimentation

Before the start of formal implementation and comprehensive testing of the vehicle identification system, it is crucial to lay the groundwork and develop a sound foundational understanding by first conducting initial testing and experimentation. In this chapter basic tests are performed to determine the fundamental capabilities and limitations of the RFID equipment. These initial experiments serve to guide us toward making informed decisions during the implementation phase. The data collected in this phase is crucial for the design and configuration of the RFID systems to be implemented later. Through this chapter, we aim to equip readers with an understanding of the initial testing and experimentation processes, preparing them for the subsequent chapters where we delve into the implementation and comprehensive testing of RFID systems.

4.1 Testing Conditions

4.1.1 Tags to be tested

Unless stated otherwise, testing procedures will be carried out on a set of 5 functional samples of license plate tags.

4.1.2 Tag initialization

All tested tags must be initialised with a dedicated UII. The UII memory should be unlocked, and both access and kill passwords must be set to zero.

4.1.3 Test environment

If not stated otherwise, the system is configured in an open area to minimise reflections from nearby surfaces.

4.2 Waveform

Tests involving forward link commands will utilise PR-ASK signalling, unless stated otherwise. PR-ASK is defined as a 25 μ s Tari, 64/3 DR, M=4 Miller subcarrier, 256 kHz BLF. The latter waveform parameters will be adjusted accordingly for tests involving fast-moving tags.

4.3 Protocol

For slow moving setups, the Reader shall inventory a tag using a Query command with the following arguments presented in Table XII.

Table XII: Reading protocol

Function	Register entry	Value
<u>DR</u>	1	64/3
M	10	4
TRext	1	Extended preamble

SEL	01	All
Session	01	S1
Target	0	A->B
Q	0000	0

The Reader shall acknowledge (ACK) all RN16 responses and verify a UII for bit errors. A successful read can only be declared when both a RN16 and correct UII are received.

4.4 Frequency

There are two suitable RFID frequency bands that can be utilised in South-Africa. These bands are, the ETSI band with frequencies 865-868 MHz and the FHSS band with frequencies 915.4-919 MHz. The centre frequency of the available RFID antennas is 917 MHz. For this reason, the tests described in this document will be performed in the FHSS band described above.

4.5 Physical Setup

The Reader and TUT (Tag Under Test) shall have the following general setup; the specific setup details are captured in each test case:

- The height of the tag shall be 0.5 m from the ground unless otherwise specified
- The orientation of the tag shall be such that the longest dimension is horizontal
- The reader and tag shall have direct line-of-sight
- The reader and tag are set up in a matched configuration
- The tag shall be oriented to face the reader antenna
- The tags shall be linearly polarised
- The default reader antenna shall be circularly polarised, unless otherwise specified

4.6 Transmit Power

Transmit power is determined by the RFID reader and can be adjusted in software for the specific practical need. The transmit EIRP is determined from the RFID reader output power, the cable loss and the transmit antenna gain. The EIRP is calculated as follows:

$$EIRP = P_T - L_c + G_A$$

Where:

$EIRP$ is the transmit EIRP in dBm

P_T is the output power of the transmitter in dBm

L_c is the cable loss in dB

G_A is the antenna gain in dBi

4.7 Experimental equipment

Table XIII: Experimental equipment

Equipment	Quantity	Description
Computer/Laptop	1	-RFID Software to record reads
RFID reader module	2	-Placed with test antenna on stand and in the ground -Bluetooth connectivity to roadside processor
Power supply	2	-60W
Antenna mounting bracket	2	-Mount fixed reader to adjustable stand
RFID cable	2	- 7m (including links and SMA male connector)
RFID cable 200mm with SMA male connectors	2	Short RF cable between RFID reader and inroad antenna
Power and data cable	2	Connects roadside unit to RFID module. Length = 7m
Test antenna	1	-Circular polarization -Known gain (8.8dBm) -902 to 928 MHz
In-road antenna	1	-200mm x 300mm x 30mm (b, l, h) -SMA connection -917 MHz centre frequency
RFID tags	10	-10 x license plate (IDePLATE)
Mounting stand	1	-Adjustable height: 200mm to 1.6m
Plastic angle adjusting spacer (3D print)	1	-Adjustment angle: 0 to 180 degrees
Ground plane	1	-16m ² -Steel mesh ~ 1mm mesh spacing
License plate holder	5	-Match size of license plate tag

Extension lead reel	1	-40m -220VAC -16A
---------------------	---	-------------------------

4.8 Test methods

4.8.1 Tag Read Range

4.8.1.1 Purpose

The objective of this test is to characterise the chosen tags in terms of their read range. The test will provide information on the sensitivity of the tag, which is directly associated with its forward link range. Tag read ranges are commonly constrained by the forward link range (with the reverse link range being higher).

4.8.1.2 Description

In this test, the tag is situated in the far-field of the antenna, and the transmit power is adjusted to identify the minimum threshold power required to activate the tag. A successful read is confirmed when the tag returns a valid RN16 and error-free UII while responding to the reader's Query and ACK commands. The read range is calculated based on the free space path loss to the tag, considering a transmit Effective Isotropic Radiated Power (EIRP) of 36dBm [66].

Numerous factors contribute to the nominal read range of a tag, including RF reflections, interference, and the presence of lossy materials. In this experiment, we assume that the tag and the reader are within an environment as described in 4.1.3. However, it is important to note that the actual range experienced in a deployed environment may deviate from the reported values due to various environmental conditions and factors.

4.8.1.3 Setup

Set up the system as specified in 4.5.

The tests will be repeated over 5 measuring points as depicted in Figure 40. The figure shows the top view of the setup, with the rectangle representing the RFID tag. These points are spaced 45 degrees apart at a distance of 1 meter from the numberplate tag centre. For each repetition, the reader will be placed on a reading point facing the numberplate tag.




Figure 40: Tag read range measuring points

The test will be done with the numberplate tag in free space (mounted on a mast as in Figure 44) and with the numberplate tag mounted on a vehicle. Figure 41 shows the test setup where the tag is mounted on the vehicle. Vehicle mounted tests are also repeated over the 5 measuring points as described previously.



Figure 41: Read range test setup (vehicle mounted tag)

4.8.1.4 Method

The method to conduct the test is as follows:

1. Mount the fixed reader and TUT on the masts in the default orientation specified in 4.5
2. Set the reader frequency as specified in 4.4. Set the reader EIRP below the minimum expected power to excite the TUT.
3. Repeatedly transmit the waveform defined in 4.2. Read the tags using the protocol defined in 4.3 and monitor the success rates.
4. Increase the reader EIRP in steps of 0.25 dBm. Record the EIRP where the tag starts to successfully respond.
5. Increase the EIRP by 1 dBm. From there, decrease the EIRP in 0.25 dBm steps and record the EIRP where the tag stops responding successfully.
6. Calculate the average of the two recorded EIRP values. Use the equation listed below to calculate the nominal TUT read range assuming a 36dBm EIRP [66].

$$\text{Read range (meters)} = r \times 10^{\left(\frac{36 - \overline{\text{EIRP}}}{20}\right)}$$

Where:

- r is the distance between the reader and TUT in meters
- $\overline{\text{EIRP}}$ is the average of the two measured EIRPs in dBm

7. Repeat the test for all the number plate tags
8. Repeat the test for 5 measuring points: 0, 45, 90, 270 and 315 degree angles, refer to Figure 42.
9. Mount the numberplate tag on the vehicle as displayed in Figure 43 and repeat steps 3 to 8.

4.8.1.5 Results

Table XIV and XV provide example calculations of the RFID number plate tag read range, where the tag is facing the reader at 0 degrees. The example results are based on the best performing tag in the population. The upper limit EIRP denotes the EIRP where the tag starts to respond, while the lower limit EIRP denotes the EIRP where the tag stops responding successfully. The gain of the test antenna is 8.8dBi. The test setup resulted in a TUT read range of approximately 9m for free space tags and 7m for vehicle mounted tags.

Table XIV: Read range calculation at 0 degrees (free space)

	Reader Power (dBm)	Cable loss (dB)	Antenna gain (dBi)	EIRP (dBm)
Upper limit EIRP (dBm)	8.6	0.5	8.8	16.9
Lower limit EIRP (dBm)	8.4	0.5	8.8	16.7
Average EIRP (dBm)	16.8			
TUT read range (m)	9.12			

Table XV: Read range calculation at 0 degrees (vehicle mounted)

	Reader Power (dBm)	Cable loss (dB)	Antenna gain (dBi)	EIRP (dBm)
Upper limit EIRP (dBm)	10.6	0.5	8.8	18.9
Lower limit EIRP (dBm)	10.4	0.5	8.8	18.7
Average EIRP (dBm)	18.8			
TUT read range (m)	7.244			

Figure 42 shows a circular graph of the read range results for the numberplate tag in free space and mounted on the vehicle.




Figure 42: RFID number plate read range

4.8.1.6 Discussion

The reliability of the RFID vehicle detection system is largely dependent on the sensitivity of the tag with regards to the distance and orientation relative to the reader. Although the tag characteristics are not the only factors influencing the detection performance of the system, it is necessary to evaluate the individual performance of the selected license plate tag in comparison with standard tags. Comparing the two graphs in Figure 41, we can see a significant drop in the read range when the numberplate tag is mounted in the vehicle. The theoretical required read range to detect speeding vehicles as determined in the theoretical and requirement sections of this research, is 4m. Thus, the measured read range for the vehicle mounted tag proves that the tag can reliably and accurately identify vehicles at the required distance.

From this discussion we can clearly see that the IDePLATE tag provides the necessary sensitivity to allow effective identification (when mounted on a vehicle) up to a distance of approximately 7m.

4.8.2 Tag orientation tolerance

4.8.2.1 Purpose

The performance of RFID systems in an EVI environment is significantly influenced by the orientation of the numberplate tag. In real-world scenarios, vehicle tags may assume various orientations in relation to the reader. This test aims to assess how the sensitivity of the tag is impacted as it undergoes rotation away from the orientation facing the reader antenna.

4.8.2.2 Description

The test involves placing the numberplate tag within the antenna's field and adjusting the transmit power until a successful response is received, following the procedure outlined in section 4.7.1.2. This process is repeated for various tag azimuth angles ranging from -90° to $+90^{\circ}$. The read range is then determined as detailed in section 4.7.1.4.

4.8.2.3 Setup

Set the system up as specified in 4.5. The reader and tag remain fixed in their positions, while the numberplate tag's azimuth angle in relation to the reader is adjusted. Figure 43 provides a schematic description of the tag orientation tolerance test setup.




Figure 43: Orientation tolerance test illustration

Figure 44 shows an example test setup where the tag is rotated around the axis of the vertical tag mounting pipe.



Figure 44: Tag orientation tolerance test setup

4.8.2.4 Method

The method to conduct the test is as follows:

1. Perform steps 1 through 6 in 4.7.1.4.
2. Perform the described procedure in step 1 for positive azimuth angles of 10°, 20°, 30°, 40°, 50°, 60°, 70°, 80°, and 90°. Note that a positive azimuth angle is defined as a counterclockwise rotation when viewing the tag from above.
3. Repeat step 2 for negative azimuth angles.
4. Plot the read range against the azimuth angles.

4.8.2.5 Results

Table XVI: Tag orientation test result

		Azimuth Angles									
Angle (degrees)		0	10	20	30	40	50	60	70	80	90
Reader	Power	8.6	9.2	9.6	11.2	12.6	14.6	17.6	21.6	25.2	27.6
(Upper	limit)										
(dBm)											
Reader	Power	8.4	9	9.4	11	12.4	14.4	17.4	21.4	25	27.4
(Lower	limit)										
(dBm)											

Path loss	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Antenna gain	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8
Upper limit EIRP (dBm)	16.9	17.5	17.9	19.5	20.9	22.9	25.9	29.9	33.5	16.9
Lower limit EIRP (dBm)	16.7	17.3	17.7	19.3	20.7	22.7	25.7	29.7	33.3	16.7
Average EIRP (dBm)	16.8	17.4	17.8	19.4	20.8	22.8	25.8	29.8	33.4	16.8
TUT read range	9.12	8.51	8.12	6.76	5.75	4.57	3.23	2.04	1.34	9.12

Table XVII: Tag orientation test result

	Azimuth Angles									
Angle (degrees)	270	280	290	300	310	320	330	340	350	360
Reader Power (Upper limit) (dBm)	27.2	27.2	23.6	19.2	15.6	12.2	10.2	9.2	8.6	8.6
Reader Power (Lower limit) (dBm)	27	27	23.4	19	15.4	12	10	9	8.4	8.4
Path loss	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Antenna gain	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8	8.8
Upper limit EIRP (dBm)	35.5	35.5	31.9	27.5	23.9	20.5	18.5	17.5	16.9	16.9
Lower limit EIRP (dBm)	35.3	35.3	31.7	27.3	23.7	20.3	18.3	17.3	16.7	16.7
Average EIRP (dBm)	35.4	35.4	31.8	27.4	23.8	20.4	18.4	17.4	16.8	16.8
TUT read range	1.07	1.07	1.62	2.69	4.07	6.02	7.58	8.51	9.12	9.12

Table XVI and XVII presents the results for the tag orientation tolerance test, while Figure 45 displays these results on a graph. At an orientation angle of 0 degrees the tag can be read at a distance of almost 10 meters away. As the tag orientation is increased relative to the angle facing the reader, its reading range decreases. When it reaches an orientation which is parallel with the radiation direction of the reader, the read range is less than 1 meter.




Figure 45: Tag orientation tolerance

4.8.2.6 Discussion

In a real world scenario drivers may change lanes upon entering the detection zone, which will affect the incidence angle of the tag relative to the reader. When observing Figure 40, we can form a picture of how the tag will behave when its orientation relative to the reader is changed. The ideal orientation of a tag is when it is directly facing the reader, i.e. perpendicular to the reader's path of maximum gain. When the vehicle enters the detection zone with an ideal orientation, i.e. 0 degrees, the tag can be detected from approximately 10m away. As the tag moves away from the ideal orientation facing the reader, its sensitivity is drastically affected. We can note that even at an incidence angle of 50 degrees relative to the direction of movement, the tag still provides a read range of almost 5 meters, which is more than the required 4m reading capability. Therefore, we can see that the tag read range is more than sufficient for the vehicle detection application for different angles relative to the reader.

This information helps in optimizing the placement and configuration of RFID readers to achieve the desired coverage and minimise blind spots.

4.8.3 Tag spacing from metal background

4.8.3.1 Purpose

Metal surfaces reflect and shield electromagnetic waves. In the presence of metals, the performance of a RFID system can become unpredictable due to a change in the characteristics of the antenna and magnetic field distribution.

The purpose of this test is to evaluate the effect of different metal surfaces placed behind the license plate tag. The test will predict the probable influence of a metallic vehicle body part on the tag.

4.8.3.2 Description

The tag will be positioned in the antenna field, and the transmit power will be adjusted until a successful response from the tag is achieved, following the procedure outlined in 4.7.1.2. Two distinct measurement setups will undergo testing:

1. The tag distance from a flat metal surface oriented parallel to the tag is varied.

The tests will be repeated over 5 measuring points as depicted in Figure 46. The figure shows the top view of the setup, with the white rectangle representing the metal plate. These points are spaced 45 degrees apart at a distance of 1 meter from the numberplate tag centre. In effect, we will not only determine the impact on the tag sensitivity, but also the impact on the tag radiation pattern.




Figure 46: Tag measuring points

4.8.3.3 Setup

Set the system up as specified in 4.5, including the following adjustments:

1. The tag is placed parallel to a flat metal surface.
2. The spacing between the tag and metal background is varied from 0mm to 30mm.

Figure 47 provides an illustration of the test setup.



Figure 47: Metal background test setup

4.8.3.4 Method

The method to conduct the test is as follows:

1. Start with a measurement position perpendicular to the tag.
2. Perform steps 1 through 6 in 4.7.1.4 with a spacing of 0mm.
3. Repeat step 1 for 5mm, 10mm, 15mm, and 20mm, 25mm, 30mm spacings.
4. Repeat the test for the other measurement positions.

4.8.3.5 Results

Figure 48 depicts the results for the setup. The circular axis on the outside is the angle in degrees, while the axis in the centre is the distance in meters. The red rectangle in the centre of Figure 48 indicates the top view of the number plate tag, while the orange rectangle represents the top view of the metal plate positioned parallel to the tag. The tag is placed parallel to the metal surface while the spacing is adjusted. From the figure, when comparing 'no metal' to the rest of the graphs, we can see that the metal background improves the tag read range in the direction perpendicular to the metal surface.




Figure 48: Parallel setup

4.8.3.6 Discussion

A substantial metal surface positioned behind an RFID tag, such as a vehicle body, can induce alterations in the tag's radiation pattern. Typically, this results in increased directionality (decreasing aperture) perpendicular to the metal surface. However, the presence of curved metal surfaces introduces a more intricate deformation in the radiation pattern, especially given the diverse forms and materials of vehicles. Despite conducting tests with known metal surface sizes, forms, and orientations, these experiments provide some insight into the tag's behaviour when affixed to the front or back of a vehicle.

Examining Figure 48, where the metal sheet aligns parallel to the tag, it becomes evident that the effective read range of the tag actually expands in the direction perpendicular to the metal surface. This implies that the tag sensitivity might even intensify when affixed to a vehicle with a metallic bumper or grille. Notably, Figure 48 also reveals that the slotted metal antenna of the license plate tag experiences a slight reduction in read range when in contact with the metal surface. Therefore, cautious consideration is necessary when affixing license plates to vehicles, as the tag ideally requires a minimum spacing of 5mm from a metal surface.

Given that most RFID systems are directional, relying on the tag and reader being within each other's beam for successful reading, the reduction in aperture becomes crucial. The measurement results in Figure 48 underscore the impact of a metal surface on tag directionality, particularly revealing a decrease in tag orientation tolerance for spacings less than 15mm. Conversely, spacings exceeding 15mm result in less directional radiation patterns, enhancing the likelihood of successful communication, particularly in dynamic scenarios like vehicle identification, where tags may assume various orientations.

In summary, the tests indicate that optimizing the tag's read range and orientation tolerance is possible by appropriately spacing it from a metal surface. However, this must be approached with caution due to the unpredictable effects of a metallic vehicle body on the tag's radiation pattern.

4.8.4 In-road reader antenna radiation pattern (tag facing reader)

4.8.4.1 Purpose

The goal of this test is to characterise the in-road reader antennas based on their radiation patterns, which illustrate the directional dependence of radio wave strength emitted by the antenna. The radiation pattern serves as a guide to predict the optimal orientation and effectiveness of the antenna design [67]. Prior research has suggested that incorporating a conductive layer in the in-road setup can enhance the readability of license plate tags. This layer works by altering the radiation pattern of the antenna, directing it away from the ground to achieve maximum gain in proximity to the vehicle license plate. Additionally, this test aims to identify the most effective size for the conductive layer.

4.8.4.2 Description

The test involves placing the tag in the antenna's field while adjusting the transmit power until a successful tag response is achieved, as outlined in section 4.7.1.2. This procedure is repeated for reader azimuth angles ranging from -90° to $+90^{\circ}$. Figure 49 displays the test setup, where the reader is displayed as the square in the centre while the tag is displayed as a short line at each angle. The distance between the reader and tag centres is maintained at 1 meter. The read range is determined following the guidelines in section 4.7.1.4. It's worth noting that the conductive layer utilised in these experiments is an aluminium mesh sheet.




Figure 49: In-road reader antenna test setup

4.8.4.3 Setup

An example system setup is provided in Figure 50:

- The RFID reader and antenna shall be placed inside a cavity in the ground
- The conductive layer (mesh sheet) shall be placed under the transmit antenna
- The tag antenna shall have a height of 400mm above the ground



Figure 50: In-road antenna test setup

4.8.4.4 Method

The method to conduct the test is as follows:

1. With the reader fixed in one location in the ground, perform steps 1 through 6 in 4.7.1.4 with the tag positioned zero degrees with respect to the reader, and with the tag facing in the direction of the reader.
2. Repeat step 1 from +90 to -90 positive azimuth angles (angle formed by the position of the tag relative to the reader) in steps of 15°. The tag faces the reader at each position.

4.8.4.5 Results

Figure 51 displays the results of the in-road antenna test. From the graph we can clearly see the antenna radiates more to the sides, accommodating for the fact that vehicles may enter the detection zone from the sides. When observing the effect of the added conductive plane, it is clear that there is a significant improvement in read range when compared to the graph where no ground plane is present. We can also see that there is only a slight change in the tag read range for different ground plane sizes.




Figure 51: In-road antenna radiation pattern

4.8.4.6 Discussion

Upon examining Figure 51, we can observe one half of the symmetric pattern of the in-road RFID reader. In a typical traffic scenario, vehicles approach the reader from the zero-degree direction and traverse over the reader to exit the detection zone in the 180-degree direction. It is evident that the reader, coupled with its conductive layer, delivers a read range of up to 10 meters, aligning with the desired requirement of 4 meters outlined in the literature study.

A comparison of the radiation pattern graph without a ground plane to those with a ground plane reveals a notable expansion in the read range across all directions. In the 0-degree direction, there's a discernible improvement from 7 meters to 10 meters. This additional 3-meter read range allows the system more time for tag interrogation.

It's worth noting that the introduced conductive layer significantly amplifies the read range to the sides of the reader. This becomes crucial in scenarios where the tag orientation remains fixed along the vehicle travel path. As the tag approaches the 90 and 270-degree angles, it aligns

perpendicular to the direction of radiation from the reader, resulting in diminished RF power reception and complicating communication with the reader. Maximizing radiation power towards the sides enhances the reader's ability to interrogate the tag during stages where its orientation is suboptimal.

This experiment provides valuable insights into the radiation pattern of the in-road antenna and confirms the appropriateness of the antenna design along with the required conductive plane. For installation on a public road, it is recommended to incorporate a conductive plane with a radius no less than 1 meter around the in-road reader to enhance its reading capabilities. This conductive plane can be implemented using materials such as a steel mesh or a carbon fabric, applied to the surface with epoxy road marking paint.

4.8.5 RFID signal Illumination area

4.8.5.1 Purpose

The purpose of this test is to capture the signal illumination area when the tag is oriented in a fashion which is encountered in a real-world setup. This test differs from the previous test in that the tag orientation stays fixed in one direction as it moves around the in-road reader.

4.8.5.2 Description

In this test, the tag is placed within the antenna's field, and the transmit power is adjusted until the tag responds successfully, following the procedure outlined in section 4.7.1.2. The test is then iterated for tag azimuth angles ranging from -90° to $+90^{\circ}$. Figure 52 displays the test setup, where the reader is displayed as the square in the centre while the tag is displayed as a short line at each angle. The distance between the reader and tag centres is maintained at 1 meter. The read range is subsequently determined in accordance with the instructions provided in section 4.7.1.4.




Figure 52: Tag illumination area test setup

4.8.5.3 Setup

The system setup is identical to that described in 4.7.4.3 and is displayed in Figure 53.



Figure 53: Illumination area test setup

4.8.5.4 Method

The method to conduct the test is as follows:

1. With the reader fixed in one location in the ground, perform steps 1 through 6 in 4.7.1.4 with the tag angle zero degrees relative to the reader (the tag facing the shortest dimension of the reader).
2. Repeat step 1, incrementing the positive azimuth angles from +90 to -90 in 15° increments. Maintain a fixed tag orientation in the lateral direction as it rotates around the reader.

4.8.5.5 Results

By referring to the plot depicted in Figure 54, the distance at which effective readability can be determined varies based on the travel path's lateral offset from the reader position. From the figure it can be seen that the detection system meets the specified 4m reading zone requirement as defined in [8].




Figure 54: RFID signal illumination area




Figure 55: Max tag read range vs vehicle lateral position

4.8.5.6 Discussion

The findings from the preceding in-road reader test underscored its radiation pattern, emphasizing an optimal tag orientation facing the reader consistently. This highlighted that the reader supplied ample power, enabling successful tag detection at distances approaching 10 meters.

It is imperative to note that the illumination area was measured in a controlled setup with favorable conditions. In less ideal real-world conditions, including reflections from nearby objects and vehicles, the effective coverage significantly diminishes. Introducing readers in adjacent lanes poses its challenges, with potential signal area overlapping. While this overlap fortifies the detection zone and minimises the risk of missed detections, excessive overlap can lead to interference, creating "dead" spots where vehicles may pass undetected. To counter reader-to-reader interference, meticulous attention to reader spacing is crucial. Adopting a zigzag configuration, instead of collinear spacing, can eliminate the creation of these problematic "dead" spots. Subsequent in-road tests will unveil the system's coverage in non-ideal environments, factoring in interference from vehicles and additional in-road readers.

Figure 54 lacks clarity on the effective readable distance of a tag when traveling along different lateral paths relative to the reader. Figure 55 addresses this by illustrating readable distances as a function of lateral position relative to the in-road reader. For instance, a vehicle with a 1.75m offset from the reader's center line begins detection at approximately 9.5m from the reader. Even if a vehicle attempts to avoid detection by traveling directly between two lanes, it will still be successfully detected, considering the minimum 9.5m illumination available on either side of the RFID reader. Consequently, the system consistently meets the specified 4m reading zone requirement outlined in [8].

4.9 Chapter conclusion

In conclusion, this chapter has functioned as an essential prelude to the official deployment and comprehensive testing of the vehicle identification system. We began this chapter with some initial testing and experimentation because we realised how important it was to lay a solid foundation. The read range, orientation tolerance, metal background, and lighting area tests were among the studies carried out in this chapter. The main goal was to obtain understanding of the basic capabilities and constraints of the RFID equipment in order to prepare for making well-informed decisions throughout the next implementation phase.

5 Detail Design

The process of component selection has furnished us with the essential hardware elements tailored to the specific needs of our in-road RFID reader. Moving beyond this phase, we now delve into the intricacies of combining these components with the Computer Vision system to form a cohesive, functional, and highly efficient system. This chapter will entail a detailed discussion of the hardware and software components that constitute the system. By the end of this chapter, we aim to provide a comprehensive blueprint for the construction and assembly of the system, ensuring it aligns with the project's specific objectives and requirements.

5.1 System architecture

Figure 56 presents the new system architecture which is updated with the selected RFID technology. For the purpose of this study, we have also selected a PC as edge controller to manage incoming data from the RFID and camera systems. I/F (Interface) 4.0 and 5.0 establish connections from the traffic cameras and RFID system to the edge controller (PC) through serial protocols. Interface 6.0 denotes the connection between the RFID reader and the antenna, employing a conventional SMA connector and coaxial cable. Finally, interface 3.0 serves not only as an RF connection between the RFID tag and the RFID reader, but also as an interface between the hardware components and the physical environment.




Figure 56: Detailed system architecture

5.2 Hardware

5.2.1 RFID system hardware

5.2.1.1 Roadside controller (PC)

The roadside controller used for the recognition system is a PC with the following specifications: Intel Core i7, 16GB DDR4 RAM, NVIDIA Geforce GTX 1060 6GB graphics. The PC receives data from both the RFID reader and the camera.

5.2.1.2 In-road RFID reader and antenna

Figure 57, displayed below, provides an overview of the fully assembled in-road RFID reader. Within this assembly, the ThingMagic M6E Micro reader (housed within a sealed cylindrical tube labelled as (1)) is securely affixed to the baseplate of the antenna (labelled as (2)). The tube has a diameter of 150mm and a height of 200mm, ensuring a compact and protective enclosure. To ensure waterproofing, the cables linking the reader to the edge controller pass through a cable gland. These cables are then enclosed within a conduit, strategically positioned beneath the road's pavement, safeguarding them from potential damage and external interference.

The baseplate of the antenna is engineered to rest in the top layer of the pavement and can be firmly secured in place using four bolts. Notably, the only component protruding above the road's surface is the radome of the antenna, which complies with regulatory limits, standing at a height of 20mm. This design ensures the system's seamless integration into the road infrastructure.




Figure 57: Assembled in-road RFID reader

5.2.1.3 RFID conductive layer

As outlined in [69], the in-road RFID antenna necessitates the presence of a conductive layer as an integral component, to be installed on the road's surface. This conductive layer plays a pivotal role in sculpting the radiation pattern, ensuring it aligns with the desired configuration depicted in Figure 24. To meet the antenna's operational requirements, the conductive layer must have a minimum radius of 1 meter.

From the investigations in the initial testing and evaluation chapter, it was observed that varying the size of the conductive layer, even exceeding 1 square meter, has a minimal impact on the antenna's performance. These tests indicated that the 2.4m x 2.4m ground plane produced the largest read distance to the front of the reader. Therefore, a ground plane size of 2.4m x 2.4m is selected, which amply covers the test lane described in the upcoming chapter. In Figure 58, we

provide an illustrative example of the conductive layer that must be employed in conjunction with the in-road antenna to achieve optimal functionality. In a practical installation, the conductive layer can be painted onto the surface of the pavement using an epoxy road marking paint. This application serves as a protective measure against wear and tear. Additional installation methods are described in [69].



Figure 58: Required RFID antenna ground plane

5.2.1.4 RFID tag

The selected RFID tag, IDePLATE manufactured by Tönnjes, is fixed to a vehicle by slotting it into the standard license plate holder.

5.2.1.4.1 RFID TAG ENCODING

To meet the system requirements, the RFID tags must store essential vehicle information, including the vehicle's license plate characters, make, model, and colour. Given the limited memory available on RFID tags, strategies need to be devised to optimise storage.

The IDePLATE includes a passive RFID tag with multiple memory banks. The following memory considerations will allow for an efficient data storage design:

Make and Model:

Vehicle make and model is stored as 32-bit values and is referenced via a lookup table, where each entry is assigned a unique 32-bit identifier. This accommodates the extensive variety of vehicle types, including trucks and buses.

Vehicle Colour:

Colours are similarly encoded using an 8-bit lookup table, sufficient to represent 256 vehicle colours detectable by computer vision cameras.

License Plate Characters:

South African license plates consist of 7 to 9 characters, depending on the region and personalization. Representing these plates using standard ASCII encoding (8 bits per character) would require up to 72 bits for a 9-character plate. To optimise storage, a 6-bit encoding scheme was implemented, utilizing only uppercase letters (A–Z) and digits (0–9), for a total of 36 characters. This encoding reduces the storage requirement to 54 bits for a 9-character plate. To meet the alignment requirements of RFID systems, 2 additional padding bits are added, resulting in a 56-bit representation for the license plate.

Total Data Size:

Combining the 56 bits for the license plate, 32 bits for the make and model, and 8 bits for the colour results in 96 bits of total vehicle data, which fits in the EPC memory bank of the RFID tag.

Security Considerations:

To prevent cloning, digital signatures are used. These signatures, generated using a private key unique to the issuer, can be verified using a corresponding public key. While effective against cloning, digital signatures introduce significant storage challenges. A typical signature size is 512 bits. Including the 96 bits of vehicle data, 608 bits would need to be stored, exceeding the capacity of the EPC memory bank. To address this storage issues, an option is to store the vehicle data in the EPC memory bank and the digital signature in the user memory bank. However, accessing both memory banks increases the overall tag interrogation time.

5.2.2 Camera system hardware

Figure 59 displays the hardware of the roadside camera system. This system is composed of the HiLook Bullet network camera mounted on a tripod. The camera communicates with the Edge controller (PC) through an ethernet port. The main video stream of the camera is at 60Hz, 20 fps (frames per second), with a resolution of 1280 x 720 [125].



Figure 59: Camera system hardware

5.3 Software

Figure 60 shows a software flow diagram, depicting the execution of the program in a real-world setting. The software of the system consists of two separate components running synchronously to identify vehicles. The first component is the software used to communicate with the RFID reader and manage the incoming tag data. The second component is the Computer Vision software used to extract vehicle identification features from traffic footage. From Figure 60, we can see that both systems first initiate an identification process where it searches for an RFID tag or a vehicle. Upon detecting a tag or a vehicle, the system timestamps the data, whereafter it extracts the relevant information from the raw data. The steps related to synchronisation and comparison are discussed in section 5.3.4.




Figure 60: System software flow diagram

5.3.1 RFID system software

The software for the RFID system is written in the Python language and handles the tasks such as reading the RFID tag data, configuring the reader settings, managing the RFID tag data, and handling events triggered by the reader. The software uses the ThingMagic Mercury API (Application Programming Interface). The Mercury API is a (SDK) software development kit provided by ThingMagic to enable developers to interact with and control ThingMagic RFID readers and devices programmatically. For in-road tests, the RFID software is hosted on a PC, functioning as the edge controller, and interfaced with the RFID reader via a serial connection. As depicted in Figure 55, the software initiates its operation by verifying a successful connection with the RFID reader. Subsequently, the system commences the process of scanning for RFID tags. Upon detecting a tag, the program undertakes an inventory process, filtering tags according to their unique EPC numbers. The application then accesses the data fields within the tag, extracting relevant information. This information is also linked to the specific access time to allow subsequent comparison with visual information obtained from traffic camera footage.

5.3.2 Digital signature

The Ed25519 algorithm, a variant of the Edwards-curve Digital Signature Algorithm (EdDSA), is used to generate the public and private keys for the digital signature. The choice of using the Ed25519 algorithm follows its high verification speed and robust security features.

Using this library, the private key is generated first, followed by deriving the public key through the algorithm associated with the private key. The private key is securely stored on the host PC, while the public key is integrated into the system's codebase. This allows the public key to verify any data transmitted from the RFID software, ensuring its authenticity and protecting against tampering or cloning attempts.

5.3.3 Computer vision system software

The software for the computer vision system is also written in the Python language and uses multiple open-source libraries to accomplish the identification tasks. As seen from the above figure, the first step is to capture the traffic footage, which can be effectively done using the OpenCV library. The OpenCV library accesses the connected traffic camera and extracts the individual frames from the video footage. Frame by frame, the video is then passed through the vehicle identification pipeline. The detection pipeline consists of the following:

- **Vehicle detector** –As informed by the literature section 2.2.2, one stage object detectors are more ideal for quick and efficient detection. At the time of writing, YOLOv7 was identified as one of the top-performing one stage object detectors with excellent real-time processing capability, high accuracy, ease of integration, scalability, and resource

efficiency. The YOLOv7 object detection algorithm is designed to detect vehicles efficiently and accurately within video frames. Input frames are pre-processed through resizing and normalization to align with the neural network's requirements. YOLOv7 employs a CNN architecture, which includes a detection head for making predictions about bounding boxes, class probabilities and vehicle colour. It subdivides the image into a grid of cells, with each cell predicting a set of bounding boxes for potential vehicles, each described by their coordinates, dimensions, and confidence score. Additionally, class probabilities are assigned to these boxes. Non-maximum suppression is applied to refine predictions by removing redundant and overlapping boxes, ultimately delivering a list of bounding boxes representing the detected vehicles in the input image, ensuring efficient and accurate vehicle detection. The YOLOv7 architecture is displayed in Table XVIII [128].

Table XVIII: YOLOv7 model architecture

Attribute	Specification
Total layers	80
Total parameters	37 million
Gradients	37 million
Compute complexity	104 GFLOPS
Input resolution	640 x 640

- VMMR** – In the Vehicle Make and Model Recognition (VMMR) process, following the extraction of vehicles from the input frame, the subsequent task is to identify the make and model of each vehicle. This entails taking the cropped image of each detected vehicle and passing it through the VMMR neural network, which leverages a BCNN for fine-grained vehicle recognition. BCNNs (Bilinear Convolutional Neural Networks) outperform traditional methods for vehicle make and model detection due to their ability to capture and model complex relationships between different image regions. Unlike standard CNNs, which treat features independently, BCNNs model the interaction between feature pairs, enabling them to identify subtle, fine-grained details in vehicle images that are essential for distinguishing between similar-looking models. This bilateral feature interaction helps BCNNs achieve higher accuracy in tasks with small intra-class variations. The BCNN consists of two distinct VGG16-based CNNs, each specializing in extracting various hierarchical features from the vehicle image, including details like shape, colour, and texture. Consequently, these convolutional layers generate two sets of feature maps, conventionally referred to as features A and features B, capturing diverse facets of the vehicle's appearance. The pivotal step involves bilinear pooling, wherein the outer product is computed between corresponding elements of features A and B for each location in the feature maps. This process effectively captures intricate feature

interactions, such as those among the grille, headlights, logos, and other distinctive vehicle attributes. These outer product matrices are subsequently summed, yielding a bilinear pooled feature representation. These bilinear pooled features serve as input to fully connected layers responsible for mapping these features to the respective make and model classes. The model training process commences with transfer learning, utilizing pre-trained VGG16 models from the ImageNet dataset, followed by fine-tuning on the Stanford Cars dataset, housing 16,185 images spanning 196 classes. Additional insights into the training and evaluation of the VMMR model are expounded upon in the following chapter.

The VGG16 architecture is displayed in Table XIX [129].

Table XIX: VGG16 model architecture

Attribute	Specification
Total layers	16
Total parameters	138 million
Gradients	138 million
Input resolution	224 x 224

- **ANPR** – The ANPR pipeline outlined in this thesis leverages state-of-the-art open-source methodologies to construct a robust system comprising multiple stages. In the initial phase, the focus is on number plate extraction from the input image, and for this task, an identical YOLOv7 model is employed, as previously described, ensuring efficient and fast license plate detection. Input vehicle images are fed into the YOLOv7 model, which returns bounding boxes indicating potential number plates. Each bounding box is accompanied by a confidence score, aiding in the identification of the most likely bounding box encapsulating the number plate. Subsequently, the highest confidence score-bearing bounding box is cropped from the input image and funneled into the next stage, dedicated to character recognition. Before character recognition commences, the image undergoes preprocessing to improve its quality, involving resizing and noise reduction. The optical character recognition (OCR) phase uses PaddleOCR for dependable character recognition. PaddleOCR takes the cropped license plate images as input and employs CNNs to detect the positions and contours of text regions or text lines. The result is a collection of bounding boxes delineating the identified text. Once these text regions are identified, PaddleOCR proceeds with text recognition, typically relying on RNNs to decipher the characters within the bounding boxes. The result is the extracted text content within each of these boxes. The PaddleOCR model is described in Table XX.

Table XX: PaddleOCR model architecture

Module	Purpose	Architecture
Text Detection	Detects text regions in images	DBNet (Differentiable Binarization)
Text Recognition	Recognises text content within detected regions	CRNN (with CNN backbone, Bi-LSTMs, and CTC loss)
Text Direction Classification	Identifies and corrects text orientation	Lightweight CNN (e.g., MobileNet)
Multilingual Support	Handles text in different languages	Pre-trained multilingual dictionaries and models

The vehicle identification pipeline detailed above yields four distinctive visual vehicle attributes as its output: the make, model, colour and license plate number of the vehicle in question. Each piece of information extracted from every frame is timestamped to correspond with a specific time slot. Figure 61 schematically displays the detection pipeline.




Figure 61: Vehicle detection pipeline

5.3.4 Data synchronisation

The synchronization process between the camera system and the RFID system ensures that data from both sources is accurately matched and discrepancies are identified. The process involves iterating through all stored data streams from both systems and comparing them to find matches.

Initially, the system checks whether the timestamp of the data extracted from both the camera and RFID systems falls within a defined time window. If the timestamps are aligned, the number plate data is compared. A match indicates successful synchronization. If the timestamps or number plate data do not align, the current data is compared with previously stored, unsynchronised data to attempt a match.

For each comparison, the system considers variations in timestamps due to processing delays or environmental factors, such as vehicle speed or signal interference. If no match is found within the allotted time frame, it suggests one of the following issues:

- The system may have failed to detect the vehicle due to tampering, environmental interference, or system errors.
- The vehicle might have evaded one system's detection or been deliberately obscured.

In such cases, the system flags the incident for further investigation. Alerts are sent to the relevant authorities, and the vehicle information is logged for review to determine whether it is associated with illegal activity or requires further follow-up.

5.3.5 Integration software

The integration software is used to combine the data streams from the RFID and computer vision code. The integration software was developed by a fellow student and detailed in [126].

To handle the continuous and asynchronous data streams from the RFID and computer vision systems, the integration software processes each stream independently. The RFID and computer vision data are captured and processed in separate processes to ensure no data is lost due to timing conflicts. The integration code then combines these streams, organizing the data into structured formats for further comparison and analysis according to Figure 55 [126].

Docker containers were used to isolate and run the RFID and computer vision software. The RFID software was hosted in an Ubuntu-based container with Python 3.6.0 and the Mercury API wrapper, while the computer vision system ran in a separate container with a CUDA-enabled environment, TensorRT, and Python 3.11.9. Communication between the containers and the integration code was achieved through a shared Docker network. This setup enabled parallel processing of both systems in isolated environments, ensuring good integration and compatibility [126].

5.4 Chapter conclusion

In conclusion, the meticulous process of component selection has equipped us with the fundamental hardware and software elements uniquely tailored to meet the specific requirements of our in-road vehicle identification system. Through comprehensive discussions, the chapter aimed to provide a clear and informative guide for the construction and assembly of the system. This blueprint is provided to ensure alignment with the specific objectives and requirements outlined in the project scope.

6 Testing and Evaluation

This chapter will provide an in-depth exploration of the phase that bridges theoretical research with empirical results. This chapter is dedicated to the implementation and testing of the in-road RFID and Computer Vision systems, which is a crucial precursor to uncovering valuable insights in the following chapter. Here we design and execute a series of tests in environments that simulate the challenges and complexities of an actual traffic scenario. The aim is to validate the system's performance, identify limitations, and generate empirical data that form the basis of the subsequent results and discussions.

6.1 RFID tests

6.1.1 Experimental parameters

The testing conditions, reader's transmit waveform, reader protocol, testing frequency, physical setup and transmit power are determined in an identical fashion as described in sections 4.1 to 4.6. The same equipment as shown in Table XIII is used for the in-road experiments.

Four tag data read sizes will be tested:

- 96-bit – Only includes the vehicle data (registration number, make, model and colour) in the EPC memory bank
- 192-bit – Includes the vehicle data and TID
- 480-bit – The 480-bit tag presents a comprehensive data structure, encompassing the vehicle data and allocating 384 bits for User memory.
- 640-bit tag – On the other hand, a 640-bit tag, inclusive of the vehicle data in EPC memory and 544 bits of User memory. Of the available User memory, 512 bits are used for the digital signature.

6.1.2 Test environment

Low-speed tests were executed on the North-West University Campus using a designated strip of road measuring approximately 200 meters in length. This controlled environment permits

vehicles to achieve speeds of up to 70 km/h. The visual representation in Figure 62 illustrates the specific conditions of the low-speed test setting.



Figure 62: Low speed test environment

Additionally, for higher-speed tests reaching up to 120 km/h, we conducted the experiments on a designated section of the N4 roadway adjacent to Engineering 4.0, a facility of the University of Pretoria. The stretch of road utilised for these tests was approximately 400 meters long. The high-speed test environment is visually depicted in Figure 63 for reference.



Figure 63: High speed test environment

6.1.3 Reader mode test

6.1.3.1 Purpose

The specific EVI application requires rapid reading of data from the tag. The ISO 18000-63 specification [107] provides different read modes that can optimise the readability of a tag for a certain application. The purpose of the Reader mode test is to characterise the performance of the system on a road and to determine the ideal read mode that will maximise the number of successful tag reads as well as the amount of data that can be read from a tag. Table XXI outlines the characteristics of the various reader modes utilised in this test.

6.1.3.2 Description

The tag is mounted on the test vehicle. The vehicle is positioned outside the field of the antenna at a distance that is adequate to reach the required speed. The transmit power of the reader stays fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives directly over the transmit antenna and stops when outside its field. The amount of successful read attempts is recorded as the data amount is increased. A successful tag detection attempt is when a tag responds with a valid RN16 and correct UUI. This test will be repeated for the following reader modes:

Table XXI: ThingMagic reader modes

Reader mode	Modulation	TARI	BLF	M	DR
High throughput	DSB-ASK	6.25	640	0	64/3
Dense mode M4	PR-ASK	25	250	4	64/3
Dense mode M8	PR-ASK	25	250	8	64/3

The TARI is the reference time interval for commands sent from the reader to the tag. The M value sets the number of modulations used to encode each bit. The reader will inventory tags in the Dual Target mode, i.e. tags in both A and B states.

6.1.3.3 Setup

Set the system up as specified below:

- The tag is mounted on a test vehicle according to regulatory requirements
- The height of the License Plate tags shall be 400 mm from the ground
- The RFID reader shall be mounted with the antenna in the road

- An aluminium ground plane of 2.4m x 2.4m shall be placed under the in-road reader as depicted in Figure 64.
- The in-road RFID antenna and reader shall be fixed to the surface of the road as shown below.
- RF cables shall connect the in-road reader and in-road antenna
- The vehicle shall move at a constant speed over the transmit antenna.




Figure 64: RFID in-road setup

6.1.3.4 Method

The method to conduct the test is as follows:

1. Position the TUT (Tag Under Test) on the vehicle is described in 6.1.3.3
2. Determine the tag read range by following the steps in 4.7.1.4
3. Set the reader frequency as specified in 4.4. Set the reader EIRP to the maximum.
4. Repeatedly transmit the High throughput mode waveform. Read the tags using the protocol defined in 4.3 and monitor the success rates.
5. Drive the vehicle at a constant speed of 20km/h over the antenna. Record the number of successful read attempts.
6. Repeat step 4 for 30km/h, 40km/h, 50km/h, 60km/h, 70km/h, 80km/h, 100km/h, and 120km/h speeds. Higher speeds up to 180km/h are extrapolated from the graphs.
7. Repeat steps 5 and 6 for increasing data amounts: 96 bits, 192 bits, 480 bits, and 640 bits.
8. Repeat the test for the remaining reader modes.

6.1.3.5 Results

Table XXII tabulates the results for the reader mode test for a 96-bit tag. The values in the table represents the number of tag reads for each mode at different speeds. Figure 65 graphs these results and presents the relationship between the number of reads and the vehicle speed for different reader modes when the tag is storing 96 bits of data. Note, the values shown in the table for speeds above 120 km/h are extrapolated.

Table XXII: 96-bit Reader mode test

		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
High throughput mode	# of successful reads	735	489	349	295	215	167	146	114	102	80	68	59
Dense M4 mode		124	72	50	42	33	24	21	16	10	9	8	6
Dense M8 mode		73	47	39	28	21	16	14	12	8	7	6	5




Figure 65: Reader mode test results for 96 bits

Table XXIII and Figure 66 presents the results for a tag storing 192 bits of data.

Table XXIII: 192-bit Reader mode test results

		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
High through-put mode	# of successful reads	339	236	201	151	115	85	53	43	39	31	26	22
Dense M4 mode		76	47	37	31	26	15	11	9	6	5	4	4
Dense M8 mode		51	34	27	20	16	12	10	8	5	4	4	3




Figure 66: Reader mode test results for 192 bits

Table XXIV and Figure 67 presents the results for a tag storing 480 bits of data.

Table XXIV: 480-bit Reader mode test results

		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
High through-put mode	# of successful reads	247	150	123	113	86	62	45	37	33	27	23	20
Dense M4 mode		49	28	25	20	17	11	8	7	4	4	3	3
Dense M8 mode		36	24	21	14	13	9	6	5	2	2	2	1




Figure 67: Reader mode test results for 480 bits

Table XXV and Figure 68 presents the results for a tag storing 640 bits of data.

Table XXV: 640-bit Reader mode test results

		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
High throughput mode	# of successful reads	213	141	87	79	61	46	39	31	29	21	18	16
Dense M4 mode		37	19	16	15	14	10	8	7	3	3	3	2
Dense M8 mode		25	16	13	12	10	8	6	4	2	2	2	1




Figure 68: Reader mode test results for 640 bits

6.1.3.6 Discussion

Reader modes play a crucial role in defining the over-the-air data transmission between the reader antenna and the tag or endpoint in RFID systems. These modes provide a means for users to tailor the balance between read rate and resistance to interference, allowing customization to best fit the requirements of their specific use cases. This customization is achieved by altering the modulation scheme used to encode the data.

The number of modulations employed to encode each bit directly influences the resistance of data transmission to interference from other readers or devices in the vicinity. The High Throughput mode opts for FM0 encoding with the least modulation per bit, rendering it the fastest transmission mode available. In contrast, the Dense Mode M4 employs four modulations per bit, enhancing communication tolerance to interference. Lastly, the Dense Mode M8 takes this a step further by utilizing eight modulations to encode each bit, making it the most interference-tolerant mode among the options.

The graph presented in Figure 65 illustrates the tag read rates for three distinct reader modes focused on reading the stored EPC data bank. Reading solely the EPC memory bank entails the reader conducting an inventory of the tags without issuing an access command. This inventory process is notably faster than accessing other memory banks. A clear observation from a comparison between Figure 65 and the other figures reveals that the number of successful read attempts is nearly double across all modes, with High Throughput Mode enabling 59 successful reads at 180 km/h (extrapolated). Furthermore, Figure 65 shows that Dense Reader Modes M4 and M8 achieve lower read speeds as they are optimised for scenarios with a dense distribution of RFID readers in a specific area, prioritizing data integrity over read rate. At 180 km/h, Dense Modes M4 and M8 only achieve 6 and 5 reads, respectively (extrapolated). This may be

acceptable for single-lane, single-vehicle setups, but in multi-lane and multi-vehicle scenarios, the number of successful reads will likely decrease significantly.

Figure 66 displays the results of the reader inventorying the tag to extract its EPC number and accessing the tag to extract its TID. The TID, a factory-programmed identification number, remains unalterable. Accessing this additional memory bank enhances the reader's ability to distinguish between tags, increasing the accuracy and reliability of the system. However, the graph indicates that successful reads are almost half of the 96-bit tests due to the reader needing to issue an extra access command for the TID. High Throughput Mode still achieves 22 successful reads at 180 km/h (extrapolated), while Dense Modes M4 and M8 achieve 4 and 3 reads, respectively.

Figures 67 and 68 showcase the results for 480 and 640-bit read sizes. For 480 bits, the reader inventories the tag, reads its EPC, accesses the TID (an extra 96 bits), and accesses 288 bits from the user memory bank, containing user-programmable tag data like the vehicle registration number, make, and model. The 640-bit test increases data read from the user memory bank to 448 bits. Comparing these graphs, High Throughput Mode achieves a read rate of 20 and 16 for 480 and 640 bits at 180 km/h (extrapolated). Dense Mode M4 achieves two successful reads, and Dense Mode M8 achieves one successful read for both read sizes. Hence, Dense reader modes will struggle to reliably read tag user data at very high speeds, especially in high-density traffic environments with multiple vehicles and readers.

The presented results underscore the superiority of the High Throughput reader mode, which consistently produces favourable outcomes. This mode demonstrates the ability to successfully identify a tag carrying 640 bits of data, even when the tag is in motion at a high speed of 180 km/h. However, it's important to note that while this mode excels in speed optimization, it may exhibit less favourable results when deployed in close proximity to other readers, particularly in scenarios where multiple tags need to be accessed within a short timeframe. Subsequent tests outlined later in this chapter will delve into these specific factors to provide a more comprehensive understanding of the High Throughput mode's performance in different deployment scenarios.

These results also highlight a trade-off between system read rate and the amount of data that can be successfully read from the tag. A higher read rate (e.g., only reading the tag EPC) provides increased reliability in identifying a single tag but limits access to all data on the tag. Conversely, reading all memory banks may lead to missed detection. An algorithmic solution is proposed: first inventory all tag EPCs, then, if time permits, read TIDs, and if even more time is available, read the user memory bank.

Lastly, the collected data facilitates the determination of the RFID signal's illumination distance. In the High Throughput Mode, boasting a read rate of 750 tags/sec for stationary 96-bit tags, a tag in motion at 20 km/h, passing over the in-road reader, achieved successful reads 735 times. This implies an illumination duration of approximately 0.98 seconds, resulting in an effective tag illumination distance of 5.4 meters. Although this aligns with the specified 4m requirement in [8], it falls slightly short of the read range determined in the RFID lab test.

6.1.4 Multi Tag test

6.1.4.1 Purpose

This test aims to establish the maximum data read capacity when multiple tags are present in the read-zone, all traveling at identical speeds. Given the findings from [8], which identified a worst-case scenario of five tags in a 10m stretch under congested conditions for travel speeds between 60 km/h and 110 km/h, the objective is to assess the system's ability to successfully interrogate and retrieve data from these multiple tags within a specified time window.

6.1.4.2 Description

The tag is mounted on the test vehicle as described in 6.1.3.3. The vehicle is positioned outside the field of the antenna at a distance that is adequate to reach the required speed. The transmit power of the reader stays fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives directly over the transmit antenna and stops when outside its field. The amount of successful read attempts is recorded as the data amount is increased.

6.1.4.3 Setup

Set up the system as outlined in Section 6.1.3.3, with the tags mounted on the vehicle as shown in Figure 69. The TUT is positioned at the centre of the four surrounding tags, maintaining the same mounting configuration as in the previous test. Although this arrangement does not reflect a realistic real-world mounting scenario, it is intentionally designed to analyse the impact of multiple other tags on the TUT performance.



Figure 69: Multi-tag test setup

6.1.4.4 Method

The method to conduct the test is as follows:

1. Perform steps 1 through 7 in 6.1.3.4.
2. Mount four foreign tags on the vehicle and repeat step 1

6.1.4.5 Results

Table XXVI provides the results for the multi-tag test with the reader in high throughput mode. Values of speeds higher than 120km/h are extrapolated. Figure 70 displays the results of a tag interrogation with four other tags also moving through the reading zone. When looking at Figure 70 we can see that the reader was able to achieve just over 200 successful reads on 96 bit tags moving at 20km/h and 22 successful reads at 180km/h (extrapolated). For the largest read size of 640 bits, the reader is only able to interrogate the tags 4 times.

Table XXVI: Multi-tag test (High throughput)

Read size	# of suc	Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
96 bits		207	147	93	89	65	61	51	41	34	29	25	22
192 bits		139	112	46	52	46	39	31	24	20	17	15	13

480 bits		88	71	28	22	20	19	13	9	7	6	5	4
640 bits		69	54	24	18	16	14	10	7	6	4	4	3




Figure 70: Multi-tag test (High throughput mode)

Table XXVII provides the results for the multi-tag test with the reader in Dense M4 mode. Values of speeds higher than 120km/h are extrapolated.

Table XXVII: Multi-tag test (Dense M4)

		Vehicle Speed (km/h)											
Read size		20	30	40	50	60	70	80	100	120	140	160	180
96 bits	# of successful reads	25	17	12	10	8	6	6	4	4	3	3	2
192 bits		14	12	5	5	3	3	2	1	1	1	1	0
480 bits		12	10	5	4	3	2	2	1	1	1	0	0
640 bits		11	9	4	3	2	1	1	1	0	0	0	0




Figure 71: Multi-tag test (Dense M4 mode)

Table XXVIII provides the results for the multi-tag test with the reader in Dense M8 mode. Values of speeds higher than 120km/h are extrapolated.

Table XXVIII: Multi-tag test (Dense M8)

Read size		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
96 bits	# of successful reads	15	10	7	5	4	3	2	2	1	1	1	1
192 bits		12	8	6	4	3	2	2	1	1	1	1	0
480 bits		9	7	5	3	2	2	2	1	1	1	0	0
640 bits		9	7	4	2	1	1	1	0	0	0	0	0




Figure 72: Multi-tag test (Dense M8 mode)

6.1.4.6 Discussion

As evident from the preceding section, interrogating a single tag moving through the detection zone yields favourable results, enabling the reader to inventory and access the tag multiple times. However, introducing additional tags into the reading zone diminishes the reader's opportunity to extract the necessary data from all tags within a given time window.

In comparison to the previous test results, there is a notable reduction to approximately one-third in the number of successful reads for all data sizes. Notably, in the High Throughput mode, which prioritises speed over data integrity, the reader still achieves commendable results. Even at the extrapolated 180 km/h with a data size of 640 bits, this mode delivers four successful reads.

While this may seem suboptimal, comparing it with the results of other reader modes reveals their inability to produce successful reads at speeds exceeding 140 km/h (extrapolated).

Examining Figure 71, it becomes apparent that, for a data read size of 96 bits, the Dense M4 reader can interrogate tags twice at 180 km/h. However, as we move to larger read sizes, the reader encounters challenges in identifying tags. For 640-bit tags, the highest speed at which the reader can detect them is 120 km/h. Although the legal speed limit on highways in South Africa is 120 km/h, some illegal vehicles may surpass this limit to evade detection. Therefore, this reader mode can effectively be employed to detect speeding illegal vehicles, provided the amount of data stored on the tag is less than or equal to 192 bits.

Finally, Figure 72 presents the results for the Dense M8 reader mode, which performs the poorest in terms of the number of successful tag detections for various tag data sizes. The highest speed at which it can detect five 640-bit tags is 80 km/h. Vehicles traveling faster than this speed are likely to be missed. However, 96-bit tags can still be detected once at 180 km/h (extrapolated), making this reader mode suitable when there is a connection with a database to look up the EPC number and extract the corresponding vehicle data.

In summary, it is evident that the High Throughput mode stands out as the only mode capable of successfully detecting multiple tagged vehicles and extracting their stored vehicle characteristic data at extrapolated speeds up to 180 km/h. Nevertheless, this mode, which prioritises speed over interference resistance, may encounter challenges when all lanes in a multilane highway are equipped with readers. The subsequent Multi-reader test will delve into and assess the performance of this mode under such conditions.

6.1.5 Tag Height test

6.1.5.1 Purpose

The placement and height of license plates can vary widely between different vehicles. This means the performance of the in-road system may also vary between different vehicles. The tag height test will investigate the difference in tag performance for different vehicles at different license plate heights. This difference will be determined from the change in the amount of data that can be read from a tag traveling at varying speeds.

6.1.5.2 Description

The tag is mounted on the test vehicle as described in 6.1.3.3. The vehicle is positioned outside the field of the antenna at a distance that is adequate to reach the required speed. The transmit power of the reader stays fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives directly over the transmit antenna and stops when outside its field. The amount of successful read attempts is recorded as the data amount is increased. A successful read attempt is when a

tag responds with a valid RN16 and correct UII. Typical license plate heights range from 200mm - 1.3 meter; however, for this test, the height was limited to 0.6m.

6.1.5.3 Setup

Set the system up as specified in 6.1.3.3, including the following adjustments:

- The tag height relative to the ground will vary between 0.2 meter and 0.6m in 0.2 meter steps.
- The vehicle shall move at a constant speed over the transmit antenna.
- The reader operates in High throughput mode.

6.1.5.4 Method

The method to conduct the test is as follows:

1. Perform steps 1 through 7 in 6.1.3.4.
2. Repeat the test for different license plate heights.

6.1.5.5 Results

Depicting the correlation between the number of successful reads and vehicle speed, Figure 73 to 76 considers various tag heights and data read sizes.

Table XXIX provides the results for the tag height test for a read size of 96 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXIX: Tag height test (96 bits)

Height		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0.2m	# of successful reads	582	337	285	205	179	162	134	106	88	75	66	58
0.4m		735	489	349	309	215	188	166	130	107	90	78	69
0.6m		341	210	188	168	156	135	121	104	92	83	76	70




Figure 73: Tag height test results (96 bits)

Table XXX provides the results for the tag height test for a read size of 192 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXX: Tag height test (192 bits)

Height		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0.2m	# of successful reads	274	192	167	124	97	90	79	65	55	48	42	38
0.4m		339	236	201	151	115	107	94	76	64	55	49	44
0.6m		187	142	118	95	86	76	69	59	52	46	42	38




Figure 74: Tag height test results (192 bits)

Table XXXI provides the results for the tag height test for a read size of 480 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXI: Tag height test (480 bits)

Height		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0.2m	# of successful reads	198	146	92	71	64	52	44	35	28	24	20	18
0.4m		247	162	131	113	86	78	69	56	47	41	36	33
0.6m		131	98	69	45	37	32	27	20	16	13	11	10




Figure 75: Tag height test results (480 bits)

Table XXXII provides the results for the tag height test for a read size of 640 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXII: Tag height test (640 bits)

Height		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0.2m	# of successful reads	173	121	78	61	54	43	37	29	24	20	17	15
0.4m		213	141	87	79	61	51	43	33	27	23	19	17
0.6m		113	72	52	33	32	24	21	16	12	10	9	7




Figure 76: Tag height test results (640 bits)

6.1.5.6 Discussion

In a real-world environment, the RFID reader must possess the capability to detect tagged vehicles with varying license plate mounting heights. Notably, sports cars typically have license plates mounted close to the ground, while trucks may feature license plates positioned higher, possibly exceeding 0.5 meters. The in-road RFID antenna is designed with a specific radiation pattern, aiming to maximise power towards the travel path of the tag. Understanding the ideal and less ideal mounting heights is crucial.

The provided graphs clearly indicate that a height of 400mm above the ground yields the most favourable results. However, as speed increases, the graphs for different heights converge, suggesting that in high-speed traffic, the difference in tag height has a less substantial impact on system performance. Conversely, in lower speed traffic with higher vehicle density, tag height becomes a more critical factor. For a 640-bit tagged vehicle moving at 20 km/h (Figure 76), the difference between a 400mm and 600mm tag height is nearly 100 successful reads. When considering the increased traffic density and interference in a real-world setting, this difference can translate to the detection or non-detection of an offender.

Furthermore, these results lead to the conclusion that the RFID reader antenna generates maximum radiation power closer to the height of 400mm above the ground. To enhance detection capabilities for a diverse range of vehicles, including sports cars and trucks, a reinvestigation into the antenna design becomes necessary. Adjusting the antenna design to

accommodate various mounting heights can contribute to the overall effectiveness of the RFID system in real-world scenarios.

6.1.6 Vehicle lateral offset test

6.1.6.1 Purpose

From the initial experimentation and testing, we have determined the effective illumination area of the RFID tag and from the Orientation tolerance test we determined that the orientation angle of a tag relative to the reader can have a substantial effect on the performance. In this test we will investigate the performance of the system when a non-ideally oriented tag moves through the illumination area. We will relate the vehicle lateral offset relative to the reader, to the maximum amount of data that can be read from a tag traveling at a certain speed.

6.1.6.2 Description

The tag is mounted on the test vehicle as described in 6.1.3.3. The vehicle is positioned outside the field of the antenna at a distance that is adequate to reach the required speed. The transmit power of the reader stays fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives through the detection zone with different lateral offsets relative to the transmit antenna and stops when outside the detection zone. The amount of successful read attempts is recorded as the data amount is increased. A standard lane width of 3.5 meter means the furthest vehicle offset relative to a reader is 1.75 meter (assuming the placement of a reader in the centre of each lane). To achieve reliable and repeatable results, road markings are used to guide the vehicle on the correct travel path.

6.1.6.3 Setup

Set the system up as described in 6.1.5.3, with the following adjustments:

- The Tag height remains fixed at 0.4 meter
- The vehicle lateral offset varies between 0, 1 and 1.75 meters

6.1.6.4 Method

The method to conduct the test is as follows:

1. Perform steps 1 through 7 in 6.1.3.4 for the vehicle driving directly over the in-road reader.
2. Repeat step 1 for the rest of the vehicle offsets.

6.1.6.5 Results

Table XXXIII provides the results for the vehicle offset test for a read size of 96 bits. Values of speeds higher than 120km/h are extrapolated.

Presenting the number of successful reads for various lateral offsets at incremental vehicle speeds, Figure 77 to 80 exhibits the results graphically.

Table XXXIII: Vehicle offset test results (96 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0m	# of successful reads	735	489	349	309	215	198	171	134	111	94	81	72
1m		657	432	316	271	205	184	158	126	104	89	78	69
1.75m		561	385	226	204	192	175	139	112	94	81	71	63




Figure 77: Lateral offset test results (96 bits)

Table XXXIV provides the results for the vehicle offset test for a read size of 192 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXIV: Vehicle offset test results (192 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0m	# of successful reads	339	236	201	155	115	107	94	76	64	55	49	44
1m		319	201	188	133	94	89	78	62	51	44	38	34

1.75m		302	189	163	125	84	80	69	54	44	37	32	28
--------------	--	-----	-----	-----	-----	----	----	----	----	----	----	----	----




Figure 78: Lateral offset test results (192 bits)

Table XXXV provides the results for the vehicle offset test for a read size of 480 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXV: Vehicle offset test results (480 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0m	# of successful reads	247	162	131	113	86	78	69	56	47	41	36	33
1m		233	150	123	95	77	67	59	47	39	34	30	26
1.75m		218	141	107	90	72	62	54	44	36	31	26	24




Figure 79: Lateral offset test results (480 bits)

Table XXXVI provides the results for the vehicle offset test for a read size of 640 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXVI: Vehicle offset test results (640 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
0m	# of successful reads	213	141	87	79	61	51	43	33	27	23	19	17
1m		196	126	81	67	52	43	36	28	22	18	15	13
1.75m		184	112	68	65	47	39	33	25	20	16	14	12




Figure 80: Lateral offset test results (640 bits)

6.1.6.6 Discussion

From Figure 54, a comprehensive view of the RFID reader's illumination area on a road section is discernible. The illustration indicates that a solitary reader possesses the capability to cover multiple lanes. Despite this broad coverage, it is noteworthy that a tag situated farther from the reader experiences diminished power compared to a tag in closer proximity, leading to a decrease in successful transmissions. This correlation is substantiated by the accompanying graphs, revealing a decline in successful reads at varying speeds as the offset increases.

In scenarios where illicit vehicles attempt to evade detection, they often manoeuvre between two readers, necessitating an offset of 1.75 meter from each reader in the case of a 3.5 meter wide lane. For a vehicle equipped with a 96-bit tag traveling at 180km/h, detection occurs 63 times. Conversely, when the vehicle is equipped with a 640-bit tag and maintains the same speed, detection frequency diminishes to 12 times, in contrast to 17 times when directly over the reader. This underscores an elevated risk of missed detections.

Applying a calculation methodology similar to the Reader mode test, a one meter offset corresponds to an illumination distance of 4.9m, while a 1.75 meter offset results in a distance of 4.2 meter. Consequently, the detection zone spanning a single lane exhibits a minimum width of 4.2 meter and a maximum width of 5.4m at the lane centre on each side of the reader. As a result, the reader meets the theoretical requirement of a 4m wide reading zone on each side of the in-road reader [8].

6.1.7 Multi reader test

6.1.7.1 Purpose

In field, multiple readers will be deployed in close proximity at different sections of a roadway. Tag-to-Reader interference can occur as a result of multiple readers trying to access the same tag. This can also lead to additional 'dead' time between reads. It is critical to achieve as low as possible interference, while maintaining a high read rate.

The purpose of the multi-reader test is to assess how the readability of a tag is affected when an additional reader is introduced.

6.1.7.2 Description

The tag is mounted on the test vehicle as described in 6.1.3.3. The vehicle is positioned outside the field of the antennas at a distance that is adequate to reach the required speed. The transmit power of both readers stay fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives over the transmit antennas and stops when outside its field. The RUT (Reader Under Test) is the reader from which the readings are taken and saved to Excel. The amount of successful read attempts for both readers are recorded as the data amount is increased.

6.1.7.3 Setup

Set the system up as described in 6.1.5.3, with the following adjustments:

- An additional reader shall be introduced directly next to the original reader in the neighbouring lane. The RUT is positioned in the right lane, while the additional reader is positioned in the left lane.

6.1.7.4 Method

The method to conduct the test is as follows:

1. Introduce the second reader in the left lane
2. Perform steps 1 through 6 in 6.1.3.4, with the additional reader present. The test vehicle moves directly over the right side reader.
3. Repeat step 2 with the vehicle moving between the two readers.

6.1.7.5 Results

Table XXXVII provides the results for the multi reader test for a read size of 96 bits. Values of speeds higher than 120km/h are extrapolated.

Figures 81 to 82 illustrate the outcomes of the multi-reader test with the readers' modes set to High throughput. These graphs exclusively depict the results recorded from the RUT and not the additional reader. The orange graph represents the results when the vehicle traverses directly over the RUT, rendering the other reader incapable of detecting any tags. In contrast, the blue

graph displays the results when the vehicle moves between the two readers. During this movement, both readers register an equal number of detections, though only the results from one reader are presented.

Table XXXVII: Multi reader test (High throughput mode - 96 bits)

Vehicle Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	137	106	90	89	77	71	67	60	54	50	47	44
over RUT		553	396	267	243	196	170	150	121	102	88	77	69




Figure 81: Multi reader test results (96 bits)

Table XXXVIII provides the results for the multi reader test for a read size of 640 bits. Values of speeds higher than 120km/h are extrapolated.

Table XXXVIII: Multi reader test (High throughput mode - 640 bits)

Offset		Vehicle Speed (km/h)											
--------	--	----------------------	--	--	--	--	--	--	--	--	--	--	--

		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	47	34	25	17	15	12	11	8	7	6	5	4
over RUT		153	95	44	25	19	14	10	6	4	3	2	2




Figure 82: Multi reader test results (640 bits)

Table XXXIX provides the results for the multi reader test for a read size of 96 bits in Dense M4 mode. Values of speeds higher than 120km/h are extrapolated.

Figures 83 and 84 show the results of the Dense M4 mode for 96 and 640 bit tags.

Table XXXIX: Multi reader test (Dense mode M4 - 96 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	6	6	4	2	2	1	1	1	0	0	0	0
over RUT		85	79	53	22	23	19	15	11	9	7	6	5




Figure 83: Multi reader test results (96 bits)

Table XL provides the results for the multi reader test for a read size of 640 bits in Dense M4 mode. Values of speeds higher than 120km/h are extrapolated.

Table XL: Multi reader test (Dense mode M4 - 640 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	3	1	2	1	0	0	0	0	0	0	0	0
over RUT		28	21	16	5	5	4	3	2	1	1	1	0




Figure 84: Multi reader test results (640 bits)

Table XLI provides the results for the multi reader test for a read size of 96 bits in Dense M8 mode. Values of speeds higher than 120km/h are extrapolated.

Figure 85 and 86 shows the results for the Dense M8 mode for 96 and 640 bit tags.

Table XLI: Multi reader test (Dense mode M8 - 96 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	6	3	3	2	1	1	1	0	0	0	0	0
over RUT		17	14	8	6	5	4	3	2	2	1	1	1




Figure 85: Multi reader test results (96 bits)

Table XLII provides the results for the multi reader test for a read size of 640 bits in Dense M8 mode. Values of speeds higher than 120km/h are extrapolated.

Table XLII: Multi reader test (Dense mode M8 - 640 bits)

Offset		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
centre	# of successful reads	3	3	2	1	1	0	0	0	0	0	0	0
over RUT		13	12	8	4	4	3	2	2	1	1	1	1




Figure 86: Multi reader test results (640 bits)

6.1.7.6 Discussion

In scenarios where multiple RFID readers are deployed in close proximity, the potential for interference between these devices exists. Such interference can manifest as signal collisions, where signals from different readers overlap, impacting the communication capabilities with RFID tags. Rigorous testing in these scenarios is essential to identify and mitigate interference issues, ensuring effective communication between each reader and the tags within its coverage area. Strategic testing of multiple readers in close proximity allows for the optimization of reader placement to minimise interference and maximise coverage.

Examining the orange graphs in Figure 81 and 82, representing a standard test with the vehicle directly over the RUT, reveals a noticeable reduction in reading rates when an additional reader is introduced in a neighbouring lane. Although the neighbouring reader is unable to read any tags, its mere presence affects the RUT's read rate. To counter this interference, adjusting reader placement from a collinear to a zigzag fashion may minimise radiation pattern overlap, expanding the detection zone's coverage. The zigzag spacing test forms part of future work.

The blue graphs illustrate read results when the vehicle travels between two readers. A significant reduction in read rate compared to a vehicle directly over the reader is observed. This reduction is attributed to two main factors. First, as the vehicle deviates from the lane centre, the tag spends less time in the reader's illumination area, potentially leading to insufficient signal strength for successful communication. However, this is less likely the primary cause, as Lateral Offset tests with a 1.75 meter deviation did not yield such drastic effects. The more probable cause is the second reason.

The second reason is linked to the highest signal overlap between the two readers, where both attempt to access the same tag simultaneously. This means the total number of tag reads is the sum of the total reads from each reader. If we multiply the values of the blue graph with two, it still does not result in as many reads as when the vehicle drives directly over the reader. This discrepancy may be a result of collisions. Given that the High Throughput mode is not optimised for interference tolerance, the frequency of collisions may be significantly increased. Collisions trigger anti-collision protocols, introducing delays as readers implement backoff and retry mechanisms. This increased latency can lead to missed tag reads, especially in high-density, high-speed environments. The High Throughput mode demonstrates significant struggle in the presence of an additional reader. Nevertheless, it still achieves successful reads at speeds up to 180 km/h for a 640-bit tag. Despite the added complexity, this mode remains viable for the vehicle detection application, delivering sufficient performance under challenging conditions.

The orange graphs in Figures 83 and 84 indicate a noteworthy decrease in the read rate of the Dense M4 mode when compared to results without an additional reader. Specifically, a 96-bit tag can only be detected 85 times at a speed of 20 km/h, whereas it registers 124 times without another reader. When the vehicle maneuvers between the two readers, the tag detection is limited to a speed of 100 km/h for a 96-bit tag and 50 km/h for a 640-bit tag.

Similar observations can be made regarding the Dense M8 mode when examining Figures 85 and 86. The read rate experiences a significant reduction with the introduction of an additional reader. In the central position between the readers, a 96-bit tag can only be detected at speeds up to 90 km/h, while a 640-bit tag reaches a limit of 60 km/h.

These findings underscore that the dense modes do not exhibit a substantial improvement in interference tolerance compared to the High throughput mode. It is evident that despite being optimised for interference tolerance, the dense modes still face challenges when an additional reader is present. This could be attributed to both readers attempting to access the same tag, resulting in potential collisions. Although these collisions may be fewer than with the high throughput mode, they still contribute to an increase in the time required to access the tags. The Dense modes may demonstrate superior performance compared to the High throughput mode in environments where there is a higher density of readers situated closely together, such as in warehouse settings.

6.1.8 Worst case test

6.1.8.1 Purpose

The Worst-case test will reveal the performance of the RFID vehicle detection system when all of the worst-case conditions are combined for the different test scenarios.

6.1.8.2 Description

The tag is mounted on the test vehicle as described in 6.1.3.3. The vehicle is positioned outside the field of the antenna at a distance that is adequate to reach the required speed. The transmit power of the reader stays fixed at 30dBm, at the frequency specified in 4.4. The vehicle drives through the detection zone with a 1.75 meter offset relative to the transmit antenna and stops when outside the detection zone. The tag is also positioned, together with four other foreign tags, at a height of 0.6m above the road surface with 640 data bits.

6.1.8.3 Setup

Set the system up as described in 6.1.4.3, with the following adjustments:

- The Tag height remains fixed at 0.6m
- The vehicle lateral offset is fixed at 1.75 meter
- The data read size is 640 bits.
- Another in-road reader is placed in the neighbouring lane, reading tags with matched configured parameters.
- Four additional tags are fixed to the vehicle.

6.1.8.4 Method

The method to conduct the test is as follows:

1. Perform steps 1 through 7 in 6.1.3.4 for the vehicle driving along the 1.5 meter lateral offset line.

6.1.8.5 Results

Table XLIII provides the results for the worst case test. Values of speeds higher than 120km/h are extrapolated.

The worst-case test is able to reveal the system's performance under least favourable conditions. Figure 87 graphically represents the worst-case test results for the different reader modes. We can see that the High throughput mode is able to read a tagged vehicle twice at 180km/h. Dense M4 mode can still reads tags up to 100km/h, while Dense M8 mode only reads tags up to 80km/h.

Table XLIII: Worst case test

		Vehicle Speed (km/h)											
		20	30	40	50	60	70	80	100	120	140	160	180
High throughput mode	# of successful reads	47	34	18	14	13	9	7	5	4	3	3	2
Dense M4 mode		11	10	6	5	3	1	1	0	0	0	0	0
Dense M8 mode		8	7	7	4	3	1	0	0	0	0	0	0




Figure 87: Worst case test results

6.1.8.6 Discussion

The preceding assessments focused on scrutinizing the system's performance in meticulously chosen conditions and test scenarios. These tests were designed to unveil the system's behaviour under worst-case scenarios for specific parameters, including height and offset. While offering valuable insights into the impact of individual parameters on the system, these tests may not accurately reflect the system's performance in a real-world deployment.

In reality, multiple worst-case conditions can converge, exerting simultaneous influences on the system. The present evaluation addresses the system's performance in the most extreme conditions identified in prior tests — specifically, with five tags positioned at 0.6m above the road surface, traversing a path directly between two in-road readers.

Figure 87 illustrates the outcomes of the three reader modes under these limiting conditions. The data indicates a notable reduction in the number of successful reads across all reader modes compared to the base case displayed in Figure 68. For Dense M4 and M8 modes, their efficacy appears limited to low-speed traffic conditions below 100km/h, typical of urban areas or city centres. Conversely, the High Throughput mode graph reveals that even in the face of extreme conditions encountered in a traffic scenario, the system remains capable of successfully identifying vehicles traveling at speeds up to 180km/h. This underscores the system's suitability for installation on highways characterised by dense and high-speed traffic, with the ability to detect vehicles in a wide range of real-world circumstances.

6.2 Computer vision testing

6.2.1 ANPR

This section provides an in-depth exploration of the testing and experimentation carried out on the developed Automatic Number Plate Recognition (ANPR) system. The ANPR system comprises a YOLOv7 license plate detection model and a PaddleOCR license plate recognition model. These models underwent thorough training on a custom dataset consisting of thirty thousand vehicle license plate images, enabling a comprehensive evaluation of the system's performance.

The YOLOv7 model, initially pretrained on ImageNet, underwent a fine-tuning process using the specialised license plate dataset. The experimentation setup was implemented on Google Colab, utilizing the computational capabilities of a T4 GPU and 12GB of RAM. Further insights into the model architecture are provided in Section 5.3.2, offering a detailed understanding of the system's design and configuration.

6.2.1.1 Custom ANPR Dataset overview

The custom ANPR dataset encompasses a diverse collection of vehicle images, each featuring a clearly visible number plate. Comprising three thousand vehicle images captured from various angles, it is important to note that the dataset is specifically curated to include images primarily

focused on the frontal view of vehicles. Moreover, the dataset encompasses a variety of light and weather conditions, ensuring the robustness of the developed model.

An illustrative example of the ANPR dataset employed for training and model evaluation is depicted in Figure 88. To facilitate effective model training and assessment, the dataset was judiciously divided into 80% for training and 20% for test sets, striking a balance between training efficiency and model generalization. Furthermore, 20% of the training set is used for model validation.




Figure 88: Custom ANPR dataset example

6.2.1.2 Model training and testing

Solely the YOLOv7 model underwent the training process, as the PaddleOCR library already encompasses a pre-trained OCR model. The YOLOv7 model, initially pre-trained on the ImageNet dataset, was further refined through fine-tuning using the custom ANPR dataset. Notably, the ImageNet dataset comprises a vast array of images across diverse classes, including vehicles, animals, and plants. Leveraging this pre-trained model significantly expedites the training and learning phases, particularly in scenarios where resources and time are constrained.

The fine-tuning of the YOLOv7 model involved exposing it to the training set for 50 epochs with a batch size of 8 samples. Post fine-tuning, the model underwent evaluation on the separate test set. It's crucial to emphasise that the model had not encountered any images from the test set during its training phase. Consequently, the detection results obtained serve as a reliable indicator of the model's potential performance in a real-world environment, providing valuable insights into its generalization capabilities.

6.2.1.3 Results

The presented table highlights notable achievements in detection and recognition accuracies, showcasing an impressive 98.80% accuracy in plate detection on the test set and 88.33% accuracy

in recognizing plates in the test set. The proposed model demonstrates good performance, achieving an accuracy slightly below the 92% reported for the model in [127]. This remains impressive, particularly given the constraints of the training dataset size.

Table XLIV: ANPR model accuracies

Model	Accuracy Training set (%)	Accuracy Test set (%)
License Plate detection (YOLOv7)	99.80	98.80
License Plate recognition (PaddleOCR)	96.40	88.33

6.2.1.4 Discussion

While the license plate recognition achieved commendable results, it has nearly 10% lower accuracy than license plate detection. This deviation is in line with expectations, given the intricacies of the Optical Character Recognition (OCR) process. Variables like the positioning of the number plate, potential obstructions, and variations in fixing methods can exert significant influence on OCR accuracy, potentially leading to misidentifications of certain letters. It's crucial to acknowledge that attaining high accuracy in both detection and recognition entails addressing these real-world challenges, thereby bolstering the system's robustness and overall performance.

Table XLIV provides a comprehensive overview of the ANPR model's performance, showcasing impressive results that closely align with those presented in literature. In [40], the model comprised a YOLOv2 license plate detector and a sequence labelling method recognizing the entire license plate without character segmentation. Although the method in [40] achieves a higher accuracy of 91.46% for plate recognition, it was trained on a larger dataset consisting of 6648 vehicle images. In comparison, the proposed model achieves a lower accuracy than [48], employing a nine layer CNN, with 97.62% for plate recognition. Furthermore, when compared to the architecture presented in this thesis, the system described in [46] reported a lower overall accuracy (86%) using traditional machine learning techniques and handcrafted features. Thus, it is evident that state-of-the-art deep learning methods yield very high recognition accuracies.

6.2.2 ANPR road testing

6.2.2.1 Purpose

The purpose of this test is to evaluate the performance and reliability of the ANPR detection system in a real-world environment, assessing its accuracy and robustness under dynamic conditions such as varying lighting, vehicle speeds, which are not fully represented in the predefined test set.

6.2.2.2 Description

The test will examine how well the system identifies and processes license plates while accounting for environmental factors that may influence detection accuracy. By using different vehicles and real-world setups, the test will provide a realistic measure of the system's operational effectiveness. Vehicles will drive past a roadside camera at varying speeds, with the camera mounted 2 meters above the ground. The system will capture video footage of the passing vehicles and apply the recognition model to analyse the data. The detection process will start immediately after each vehicle exits the camera's field of view. It takes several seconds for the software to output the detection result(s). Tests are conducted during midday.

The two vehicles used in this setup are:

1. Vehicle 1
 - a. Make and model: Suzuki Swift 2018
 - b. Colour: white
 - c. Registration number: KRT203NW
2. Vehicle 2
 - a. Make and model: Suzuki Vitara Brezza 2022
 - b. Colour: Orange
 - c. Registration number: LS55ZRGP

6.2.2.3 Setup

Set the system up as specified below:

- The camera is mounted 2 meters above the ground at a 30° angle from horizontal, facing the testing lane.
- The camera is powered via an inverter providing 12V at 0.9A.
- The camera connects to a laptop through an Ethernet connection, using the RTSP protocol to interact with the code.
- The camera lens is initialised and zoomed in to ensure a clear view of vehicles in the video feed.

Figure 89 displays an example of the test setup.



Figure 89: Computer vision vehicle setup

6.2.2.4 Method

The method to conduct the test is as follows:

1. Position the test vehicle at a sufficient distance from the camera's frame to reach the target speed.
2. Drive the vehicle at a constant speed (e.g., 10 km/h) through the camera's frame until it exits.
3. Allow the system to analyse frames and compile detection results.
4. Perform the test for speeds of 10 km/h, 20 km/h, 30 km/h, 40 km/h, 60 km/h, and 70 km/h.
5. Perform 5 iterations for each vehicle speed.
6. Repeat the test for the second vehicle.

6.2.2.5 Results

Table XLV and XLVI displays the results for the ANPR road testing.

Table XLV: ANPR test results (vehicle 1)

Vehicle speed (km/h)	Number of successful detections
10	5
20	5
30	5
40	5
50	5
60	4

70	5
----	---

Table XLVI: ANPR test results (vehicle 2)

Vehicle speed (km/h)	Number of successful detections
10	5
20	5
30	5
40	4
50	5
60	5
70	4

6.2.2.6 Discussion

From Tables XLV and XLVI, it is evident that the ANPR detection system demonstrates high performance, with only three instances of unsuccessful license plate identification. This means a success rate of 91.4%. As illustrated in Figure 89, the presence of significant shadows in the vehicle lane could have contributed to the missed detections. Specifically, for vehicle 1, the system misinterpreted the letter 'K' as 'X,' while for vehicle 2, it misidentified the letter 'S' as the number '5.'

6.2.3 VMM

Concurrently with the assessment of the Automatic Number Plate Recognition (ANPR) system, this chapter delves into an additional dimension of Computer Vision—specifically, the recognition of vehicle make and model. A pioneering methodology was adopted, harnessing BCNN built upon the foundation of VGG16 models. These VGG16 models, initially pre-trained on the extensive ImageNet dataset, underwent detailed fine-tuning utilizing the Stanford Cars dataset. Further details into the model architecture are expounded upon in section 5.3.3.

6.2.3.1 VMMD dataset overview

The fine-tuning dataset utilised for Vehicle Model Recognition (VMMD) is the widely used Stanford Cars dataset, a benchmark in computer vision research for vehicle identification.

To facilitate a thorough evaluation of the model, the VMMD dataset is divided evenly into a 50-50 train/test split. In contrast, the ANPR model, discussed in the previous section, utilised an 80/20 split owing to a constrained availability of image samples. The ANPR dataset consisted of three thousand images, while the VMMD dataset boasts an expansive collection of approximately sixteen thousand images. This division facilitates a robust assessment of model performance. Additionally, within the training set, an additional 20% was allocated for validation purposes. The

accompanying figure visually represents the diverse scenarios encapsulated within the Stanford Cars dataset.



Figure 90: Stanford Cars dataset sample [142]

6.2.3.2 Model training and testing.

The initial phase of the training process involved passing the training data through a data augmentation generator which forms part of the training tools available for the VGG16 model. This step serves a dual purpose: augmenting the training set by introducing a variety of vehicle images and expanding the number of samples. Various transformations, including rotation, shifting, zooming, rescaling, and horizontal flipping, were applied to diversify the dataset. The augmented data is then ready to be fed into the model.

Subsequently, transfer learning is implemented, guiding the model to adapt to the new task of vehicle identification. The first stage involves fine-tuning the last layer of the model, a computationally lighter approach compared to fine-tuning the entire model. This fine-tuning of the last layer spans 20 epochs, utilizing a batch size of 32 samples and employing an SGD (Stochastic Gradient Descent) optimiser.

Following this, the entire model undergoes fine-tuning, adjusting the weights of all network layers for greater adaptability to the new task. This phase, lasting 20 epochs with a batch size of 32 samples, utilises the Adam optimiser. Once the model has been fine-tuned, it is subjected to evaluation on a previously unseen test set, ensuring a robust assessment of its performance in real-world scenarios.

6.2.3.3 Results

Figure 91 illustrates the model's accuracy when exposed to both the training and test sets.




Figure 91: VMMR training accuracy

Table XLVII: Summary of Computer Vision testing accuracies

Model	Accuracy Training set (%)	Accuracy Test set (%)
VMMR (VGG16 BCNN)	97.89	87.62

6.2.3.4 Discussion

The graph in Figure 91 indicates a discernible point of overfitting after 5 iterations, where the accuracy of the test set plateaus while the accuracy of the training set continues to rise. The accompanying table neatly summarises the attained accuracies. Notably, the model achieves an impressive VMMR accuracy of 87.62% on the test set. This accomplishment is particularly noteworthy given the challenge of distinguishing between different vehicle models from the same manufacturer. However, it is essential to recognise that this achievement implies that approximately 12 out of every 100 vehicles passing through the detection zone may not be accurately identified according to their make and model. While this result is commendable, it underscores the motivation behind this thesis—to develop a robust methodology employing multiple stages (ANPR, RFID, and VMMR) to ensure accurate vehicle identification. This multi-faceted approach aims to enhance the overall reliability of the system, addressing the complexities involved in accurately identifying vehicles in diverse real-world scenarios.

When benchmarking the achieved accuracy against results from the literature, [40] stands out with an impressive accuracy of 95.07%. This method employs a combination of global and parts-based techniques, enhancing its robustness. However, it is important to note that this model requires the annotation of both vehicle images and individual vehicle parts, rendering its architecture computationally expensive and time-consuming, especially during retraining whenever a new vehicle model is released. In contrast, approaches without deep convolution networks, such as Yang [60], achieve lower accuracy (76.7%) compared to our deep learning model. Boxcars [58], relying on deep networks with global feature representation, achieves slightly lower results than our model at 84.8%. The architecture presented in [59] utilises a pre-trained Inception-V3 convolutional neural network, achieving an overall accuracy of 82.36%. Notably, the architecture used in this thesis distinguishes itself by combining the feature representations of two separate models, contributing to its competitive performance.

6.2.4 VMRR road testing

6.2.4.1 Purpose

The purpose of this test is to evaluate the performance and reliability of the VMRR detection system in a real-world environment, assessing its accuracy and robustness under dynamic conditions such as varying lighting and vehicle speeds, which may not be fully represented in the predefined test set.

6.2.4.2 Description

The test will examine how well the system identifies and processes vehicle make and model characteristics. Vehicles will drive past a roadside camera at varying speeds, with the camera mounted 2 meters above the ground. The system will capture video footage of the passing vehicles and apply the recognition model to analyse the data. The detection process will start immediately after each vehicle exits the camera's field of view. It takes several seconds for the software to output the detection result(s). Tests are conducted during early afternoon (12pm to 2pm).

The two vehicles used in this setup are:

3. Vehicle 1
 - a. Make and model: Suzuki Swift 2018
 - b. Colour: white
 - c. Registration number: KRT203NW
4. Vehicle 2
 - a. Make and model: Suzuki Vitara Brezza 2022
 - b. Colour: orange

c. Registration number: LS55ZRGP

6.2.4.3 Setup

Set the system up as described in 6.2.2.3

6.2.4.4 Method

Perform steps 1 to 6 of 6.2.2.4

6.2.4.5 Results

Table XLVIII and XLIX displays the results for the VMMR road testing.

Table XLVIII: VMMR test results (vehicle 1)

Vehicle speed (km/h)	Number of successful detections
10	5
20	5
30	4
40	5
50	5
60	4
70	3
Total	31

Table XLIX: VMMR test results (vehicle 2)

Vehicle speed (km/h)	Number of successful detections
10	5
20	3
30	5
40	5
50	5
60	4
70	4
Total	31

6.2.4.6 Discussion

From Tables XLVII and XLIX, the VMMR detection system demonstrates strong performance, achieving an 88.5% recognition success rate across all iterations. While this slightly exceeds the accuracy achieved with the predefined test set, it is worth noting that the vehicle population consisted of only two vehicles. The missed identifications for the first vehicle were primarily due

to the system identifying the vehicle colour as grey instead of white, likely caused by shadows in the test lane. Similarly, for the second vehicle, the orange colour was misidentified as brown, also attributed to shadow interference. Additionally, two missed identifications were related to make and model classification, where the Suzuki Vitara Brezza was incorrectly classified as an earlier model, the Suzuki Vitara.

6.2.5 Integration testing

6.2.5.1 RFID tag data packing

The memory of the RFID tag is divided into two main sections:

- EPC data (12 bytes): Vehicle data.
- User data (64 bytes): Digital signature.

Breakdown of Vehicle Data (12 Bytes):

- Bytes 1 to 7: Number plate
- Bytes 8 to 11: Vehicle make and model code.
- Byte 12: Colour code

Number Plate Data Handling:

- Converted to a binary string.
- Last 2 bits (padded bits) are removed.
- The result is converted back into a string using a bitmap.

Processing Flow:

1. After extracting and decoding the data, it is checked against an existing data structure
2. If the data already exists, it is ignored.
3. If new, the digital signature is verified using the public key and extracted vehicle data.
4. Successfully verified data: Make, model, and colour codes are decoded using lookup tables.
5. Combined with the number plate and timestamp, the data is stored in the data structure.
6. The processed data is sent to the comparison software for cross-referencing with camera data.

6.2.5.2 Single vehicle test under ideal conditions

6.2.5.2.1 PURPOSE

The purpose of this test is to assess the performance and reliability of the integrated detection system combining computer vision and RFID. It also aims to evaluate the effectiveness of the integration software in synchronising data streams from both systems.

6.2.5.2.2 DESCRIPTION

The test will consist of a single vehicle driving through the detection zone, consisting of the roadside camera and in-road RFID reader, at varying speeds. The system will capture video footage of the passing vehicle and applies the recognition model to analyse the data, while the RFID reader will detect the numberplate tag mounted on the vehicle. The RFID tag originally displays the word "SAMPLE," but to simulate a realistic scenario, a printed image of an actual number plate is affixed over the "SAMPLE" text. Tests are conducted during the early afternoon (12pm to 2pm) and early evening (3pm to 5pm). This is an ideal test setup, which means only a single vehicle and RFID numberplate will be used in the test.

The vehicle used in this setup is:

- Make and model: Suzuki Swift 2018
- Colour: white
- Registration number: KRT203NW

The above vehicle data is written to an RFID numberplate and stored in the EPC memory bank, while the generated digital signature is stored in the User memory bank.

6.2.5.2.3 SETUP

Set the system up as specified below:

- The in-road RFID system is set up as specified in 6.1.3.3
- The camera system is set up as specified in 6.2.2.3

The RFID reader is positioned approximately 4.0 m from the camera to ensure accurate synchronization between the computer vision and RFID systems. This placement ensures the vehicle is immediately within the RFID signal range as it leaves the camera's frame, eliminating the need to account for variations in vehicle speed. Placing the reader farther away, such as 8.0 m, could allow drivers to speed up or slow down within the additional distance, potentially causing the RFID tag to miss the timing window, resulting in failed data synchronization and comparison.

Figure 92 displays an example of the combined test setup. The camera mounted on the tripod can be seen on the left side, while the in-road RFID reader is displayed on the right.



Figure 92: Computer vision vehicle setup (early afternoon)

6.2.5.2.4 METHOD

The method to conduct the test is as follows:

1. Position the test vehicle at a sufficient distance from the camera's frame to reach the target speed.
2. Launch and initialise the integration software
3. Drive the vehicle through the camera frame and RFID reader zone at a steady speed.
4. Allow the system to analyse the data streams and compile detection results.
5. Perform the test for speeds of 10 km/h, 20 km/h, 30 km/h, 40 km/h, 60 km/h, and 70 km/h.
6. Perform 5 iterations for each vehicle speed.

6.2.5.2.5 RESULTS

Table L displays the results for the integration test under ideal conditions during morning hours.

Table L: Ideal integration test results (early afternoon)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	5	0	0	0	5	5
20	5	0	0	0	5	5
30	5	0	0	0	5	5
40	4	0	0	1	5	4
50	5	0	0	0	5	5
60	5	0	0	0	5	5
70	3	0	1	1	5	3
Total	32	0	1	2	35	32

Table LI displays the results for the integration test under ideal conditions during midday hours.

Table LI: Ideal integration test results (early evening)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	4	0	0	1	5	4
20	5	0	0	0	5	5
30	4	0	0	1	5	4
40	4	0	0	1	5	4
50	5	0	0	0	5	5
60	5	0	0	0	5	5
70	3	1	0	1	5	4
Total	30	1	0	4	35	30

6.2.5.2.6 DISCUSSION

Table L shows that 32 out of 35 iterations across all speeds resulted in successful identifications, yielding an overall success rate of 91.4% during the early afternoon hours. Similarly, Table LI indicates that 30 out of 35 iterations were successful, achieving an overall success rate of 88.5% during early evening. The performance of the integrated system was influenced by the computer vision component, as the RFID consistently delivered accurate detections.

During the early evening hours, the presence of significantly more shadows compared to the early afternoon was a key factor impacting computer vision accuracy. The reduced lighting conditions in the test lane led to a noticeable decline in colour classification accuracy. However, when colour was excluded as a determining factor, the recognition accuracy during the early evening improved to 97.1%.

6.2.5.3 Single vehicle multi-tag test

6.2.5.3.1 PURPOSE

The purpose of this test is to evaluate the performance and reliability of the integrated detection system, which combines computer vision and RFID, in a scenario involving a single vehicle. In this scenario, the RFID reader detects multiple tags, simulating a situation where the camera system fails to identify a second tagged vehicle, recognizing only one.

6.2.5.3.2 DESCRIPTION

The test will consist of a single vehicle driving through the detection zone, consisting of the roadside camera and in-road RFID reader, at varying speeds. The system will capture video footage of the passing vehicle and applies the recognition model to analyse the data, while the RFID reader will detect the numberplate tag mounted on the vehicle. A second RFID number plate is manually introduced into the RFID reader's detection zone as the vehicle passes. Tests are conducted during early afternoon.

The vehicle used in this setup is:

- Make and model: Suzuki Swift 2018
- Colour: white
- Registration number: KRT203NW

The above vehicle data is written to RFID numberplate and stored in the EPC memory bank, while the generated digital signature is stored in the User memory bank.

6.2.5.3.3 SETUP

Set the system up as specified in 6.2.5.1.3.

6.2.5.3.4 METHOD

The method to conduct the test is as follows:

1. Perform steps 1 through 6 in 6.2.5.1.4; when the vehicle enters the RFID detection zone, introduce the additional RFID numberplate so that it can be detected by the RFID reader.

6.2.5.3.5 RESULTS

Table LII displays the results of the integration test for a single vehicle, multi-tag scenario. These results only display the data for the vehicle. Table LIII shows that for the manually introduced numberplate tag, all the RFID identifications were successful, while all the camera identifications failed. This additional tag was marked illegal for every iteration.

Table LII: Single vehicle, multi-tag test results (vehicle 1)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	5	0	0	0	5	5
20	3	0	0	2	5	3
30	5	0	0	0	5	5
40	4	0	1	0	5	4
50	5	0	0	0	5	5
60	4	0	0	1	5	4
70	5	0	0	0	4	4
Total	31	0	1	3	34	30

Table LIII: Single vehicle, multi-tag test results (manually introduced tag)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches

10	0	-	-	-	5	0
20	0	-	-	-	5	0
30	0	-	-	-	5	0
40	0	-	-	-	5	0
50	0	-	-	-	5	0
60	0	-	-	-	5	0
70	0	-	-	-	5	0
Total	0	-	-	-	35	0

6.2.5.3.6 DISCUSSION

Table LII shows that 30 out of 35 iterations across all speeds resulted in successful identifications, achieving an overall success rate of 85.7%. Similar to the previous test, vehicle colour misidentification was the primary cause of detection failures, likely due to the lighting conditions in the test lane. Additionally, the results indicate that the presence of an additional tag had minimal impact on the overall success rate of the actual vehicle identification, with the RFID failing to detect only a single identification at 70 km/h. The system also successfully flagged the additional tag as illegal, due to no matching computer vision data. Finally, it is evident that the system can synchronise and match the corresponding computer vision and RFID numberplate data.

6.2.5.4 Multi-vehicle single-tag test

6.2.5.4.1 PURPOSE

The objective of this test is to evaluate the performance and reliability of the integrated detection system, combining computer vision and RFID, in a multi-vehicle scenario. In this setup, the camera system successfully identifies multiple vehicles, whereas the RFID system detects only a single numberplate tag. This test simulates a situation where the camera system operates as expected, but the RFID system detects just one vehicle.

6.2.5.4.2 DESCRIPTION

The test will consist of multiple vehicles driving through the detection zone directly behind each other. The system will capture video footage of the passing vehicles and apply the recognition model to analyse the data, while the RFID reader will detect the numberplate tag mounted on the second vehicle. Tests are conducted during early afternoon.

The two vehicles used in this setup are:

1. Vehicle 1
 - a. Make and model: Suzuki Swift 2018
 - b. Colour: white
 - c. Registration number: KRT203NW

2. Vehicle 2

- a. Make and model: Suzuki Vitara Brezza 2022
- b. Colour: orange
- c. Registration number: LS55ZRGP

The data of vehicle 2 is written to the RFID numberplate attached to it and stored in the EPC memory bank, while the generated digital signature is stored in the User memory bank.

6.2.5.4.3 SETUP

Set the system up as specified in 6.2.5.1.3.

Figure 93 shows the test setup. The orange vehicle is the Suzuki Vitara Brezza fitted with an RFID enabled numberplate, while the white Suzuki Swift does not have an RFID numberplate. Both vehicles were in the frame of the camera at the same time while driving down the road.



Figure 93: Multi vehicle, single tag test setup

6.2.5.4.4 METHOD

The method to conduct the test is as follows:

1. Perform steps 1 through 6 in 6.2.5.1.4 with both vehicles driving through the detection zone.

6.2.5.4.5 RESULTS

Tables LIV and LV display the results of the integration test for the multi-vehicles, single tag scenario.

Table LIV: Multi-vehicle, single tag test results (vehicle 1)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	5	0	0	0	0	0
20	5	0	0	0	0	0
30	5	0	0	0	0	0
40	5	0	0	0	0	0
50	5	0	0	0	0	0
60	3	0	0	2	0	0
70	4	0	1	0	0	0
Total	32	0	1	2	0	0

Table LV: Multi-vehicle, single tag test results (vehicle 2)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	5	0	0	0	5	5
20	5	0	0	0	5	5
30	2	0	1	2	5	2
40	5	0	0	0	5	5
50	5	0	0	0	5	5
60	5	0	0	0	5	5
70	4	0	0	1	5	4
Total	31	0	1	3	35	31

6.2.5.4.6 DISCUSSION

Table LIV shows that Vehicle 1, which was not equipped with an RFID-enabled number plate, had no successful RFID identifications as expected. However, the camera system successfully identified Vehicle 1 in 32 out of 35 iterations. Meanwhile, Table LV indicates that Vehicle 2, fitted with an RFID-enabled number plate, achieved 31 successful identifications out of 35.

These results demonstrate that the integrated system effectively matched the correct camera data with the RFID data, flagging Vehicle 1 as illegal since its RFID data (blank data) did not align with the computer vision data. Additionally, the test confirmed that the computer vision system was capable of detecting two vehicles passing through the frame simultaneously.

6.2.5.5 Multi vehicle multi-tag test

6.2.5.5.1 PURPOSE

The objective of this test is to evaluate the performance and reliability of the integrated detection system, combining computer vision and RFID, in a multi-vehicle, multi-tag scenario. In this setup, the camera system and RFID system can successfully identify multiple vehicles. This test simulates a practical scenario where multiple vehicles are fitted with RFID enabled number plates and drive through the detection zone.

6.2.5.5.2 DESCRIPTION

The test will consist of multiple vehicles driving through the detection zone directly behind each other. The system will capture video footage of the passing vehicles and apply the recognition model to analyse the data, while the RFID reader will detect the numberplate tags mounted on the vehicles. Tests are conducted during early afternoon.

The two vehicles used in this setup are:

3. Vehicle 1
 - a. Make and model: Suzuki Swift 2018
 - b. Colour: white
 - c. Registration number: KRT203NW
4. Vehicle 2
 - a. Make and model: Suzuki Vitara Brezza 2022
 - b. Colour: orange
 - c. Registration number: LS55ZRGP

The data of both vehicles are written to their respective RFID number plates. The data is stored in the EPC memory bank, while the generated digital signature is stored in the User memory bank.

6.2.5.5.3 SETUP

Set the system up as specified in 6.2.5.3.3.

6.2.5.5.4 METHOD

The method to conduct the test is as follows:

1. Perform steps 1 through 6 in 6.2.5.1.4 with both vehicles driving through the detection zone.

6.2.5.5.5 RESULTS

Tables LVI and LVII display the results of the integration test for the multi-vehicle, multi-tag scenario.

Table LVI: Multi-vehicle, multi tag test results (vehicle 1)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of successful matches
10	5	0	0	0	5	5
20	4	0	0	1	5	4
30	5	0	0	0	5	5
40	5	0	0	0	5	5
50	4	0	0	1	5	4
60	5	0	0	0	4	4
70	4	1	0	0	4	3
Total	32	1	0	2	33	30

Table LVII: Multi-vehicle, multi tag test results (vehicle 2)

	Camera identifications				RFID identifications	Result
Speed	# of successful identifications	# of wrong/no registration number identifications	# of wrong/no make and model identifications	# of wrong/no colour identifications	# of successful identifications	# of matches
10	5	0	0	0	5	5

20	5	0	0	0	5	5
30	3	0	1	1	5	3
40	5	0	0	0	5	5
50	4	0	0	1	5	4
60	4	1	0	0	5	4
70	5	0	0	0	5	5
Total	31	1	1	2	35	31

6.2.5.5.6 DISCUSSION

Table LVI reveals that Vehicle 1 was successfully identified in 30 out of 35 iterations, achieving an identification rate of 85.7%. Similarly, Table LVII shows that Vehicle 2 was identified successfully in 31 out of 35 iterations, resulting in an 88.5% success rate.

These findings demonstrate the integrated system's effectiveness in identifying multiple tagged vehicles passing through the detection zone. The system accurately matched each vehicle to its corresponding number plate tag, achieving an impressive average identification success rate of 87.1%.

6.2.6 Integration testing conclusion

The integration testing demonstrated that combining computer vision and RFID identification is feasible and effective, achieving a very good identification success rate of approximately 88%. However, the results also highlight that the computer vision component is the primary limiting factor, accounting for most missed or incorrect detections. Among these, missed colour identifications were the most frequent issue, likely caused by shadows in the test lane.

To address this, the colour recognition model could be improved by training it with additional data that includes shadowed or poorly lit conditions. Alternatively, removing the colour recognition component entirely increases the overall success rate to over 90%.

Additionally, the computer vision system currently lacks real-time detection capabilities, requiring approximately 15 seconds to compute the output. Future work could focus on developing an improved system capable of real-time processing to enhance the system's practicality and performance.

6.3 Chapter conclusion

In conclusion, this chapter marked a critical transition from the theoretical foundations of the research to practical application. The focus was on designing and conducting a series of tests that simulated the challenges of real-world traffic scenarios.

The individual tests for RFID and computer vision demonstrated that both systems are capable of reliably identifying vehicles in real-world conditions. Furthermore, the integrated testing confirmed that combining RFID and computer vision data enhances the overall system's robustness. When one system failed to detect a vehicle, the integration ensured that the vehicle was flagged for further processing or action, showcasing the effectiveness of this combined approach in addressing detection challenges.

7 Conclusion

The preceding chapters have presented conceptualisation, literature review, experimentation, detailed design, and comprehensive testing, collectively leading to the realisation of an efficient in-road vehicle identification system.

In this chapter, the key findings and insights derived from the extensive investigation are discussed by answering the stated research questions. Moreover, the limitations encountered during the research are addressed, acknowledging the challenges that shape the boundaries of the study. In this chapter future avenues of research are also explored, suggesting potential enhancements for the in-road identification technology.

7.1 Research questions answered

In Chapter 1, three fundamental research questions were identified. The subsequent discussion within this section will delve into the manner in which these questions were addressed and ultimately answered throughout the course of this research.

7.1.1 Can RFID and Computer Vision technology be combined to increase the effectiveness of illegal vehicles detection?

In this study, various techniques for effectively detecting illegal vehicles were explored. The investigation revealed that each individual technique possesses the ability to identify vehicles to a certain extent, yet lacks a specific characteristic, rendering the system unreliable for comprehensive vehicle identification. Notably, RFID technology demonstrated the capacity to effectively identify a vehicle, but it fell short in verifying the authenticity of this identity. On the other hand, Computer Vision, while adept at accurately identifying vehicles by extracting visual characteristics, faced limitations in distinguishing between legal and illegal license plates. Motivated by these observations, the research aimed to assess whether the combination of these two techniques could yield greater reliability compared to their individual use.

From individual testing, the findings indicate that the in-road RFID system is well suited for detecting vehicles in a real traffic scenario. The in-road concept performed very well and was able to detect vehicles up to speeds of 180km/h. On the other hand, the computer vision models developed in this research demonstrated commendable accuracy in extracting visual characteristics from moving vehicles, however, additional improvements in the overall model data exposure and training can significantly improve the computer vision system.

Additionally, we ascertained that the data streams from the RFID and connected camera systems can be effectively combined by a connected roadside controller (PC). This integration of two data streams, each possessing its unique degree of accuracy, enhances the overall system's reliability and robustness. The integration test results showed that the system correctly identified vehicles approximately 88% of the time. Analysis indicated that the primary source of misidentification was the

computer vision system, which was significantly affected by the lighting conditions in the test environment. Enhancing the computer vision system, particularly its robustness to varying lighting conditions, has the potential to increase the recognition rate to over 90%.

The system, however, lacks real-time detection capabilities, as the roadside controller requires several seconds to process vehicle data and generate results. Potential solutions to address this limitation include optimising the integration software by minimizing reliance on multiple software containers and upgrading the hardware of the roadside controller (PC) to enhance the speed of visual vehicle classification.

In conclusion, the study emphasises the potential synergy between RFID and Computer Vision technologies, showing that their combined use can lead to a more reliable system for identifying and distinguishing between legal and illegal vehicles.

7.1.2 What are the performance criteria of the proposed in-road RFID system?

This research inquiry stems from the overarching main research question and aims to assess the viability of utilizing an in-road RFID system as a means of vehicle identification. Extensive literature review identified three primary RFID-based vehicle identification configurations: gantry setup, roadside setup, and in-road setup. Theoretical analysis suggested that the in-road setup offered an optimal solution in terms of cost and maintenance considerations. However, given its novelty in RFID applications, empirical verification was essential to substantiate the theoretical propositions.

Consequently, a comprehensive investigation was conducted to determine the performance criteria of the in-road RFID subsystem. Rigorous testing and experimentation yielded conclusive results indicating that the RFID system successfully met the requisite performance standards for identifying vehicles in high-speed, high-density traffic environments. The following key performance criteria were thoroughly explored in this research:

- **Reader Illumination Distance:** The in-road reader's illumination area demonstrated the capability to cover up to three lanes, each with a width of 3.5 meter. The illumination area extended approximately 5 meter on either side of the reader for moving vehicles.
- **Reader Mode:** This study established that only the High Throughput mode enables the reader to detect tags at speeds of up to 180km/h. It was observed that this mode prioritises speed over data reliability, making it more susceptible to interference from other readers. Dense reader modes, on the other hand, failed to meet the system requirements.
- **Multiple RFID Tags:** Through experimentation, it was determined that a maximum of five tags can be present simultaneously in the detection zone. The High Throughput mode

proved capable of reading five tags at high speeds, while Dense modes experienced challenges due to their emphasis on data integrity.

- **RFID Tag Height:** Ideal tag height was identified to be approximately 0.4m above the ground. Deviating from this range, either higher or lower, significantly diminished system performance.
- **Vehicle Offset:** In a multi-lane roadway scenario, in-road RFID readers successfully detected a tagged vehicle even when it travelled directly between two readers with a 1.75 meter offset from each reader for a 3.5 meter lane width.
- **Multiple Readers:** The study revealed that the performance of a single reader is notably affected when another reader is positioned in an adjacent lane. However, both readers remained capable of detecting the tagged vehicle at speeds of up to 180km/h.
- **Data read capacity:** The in-road reader is able to accurately interrogate tags carrying 640-bits of data up to speeds of 180km/h. Higher data read sizes will be possible at lower speeds.

In conclusion, this research not only validated the theoretical advantages of the in-road RFID setup but also provided valuable insights into the specific performance parameters crucial for its effective deployment in high-speed, high-density traffic environments.

7.1.3 Does the system meet the specific theoretical requirements as stated in literature?

The literature examined in [1], [2], [8], and [35] delineated specific prerequisites essential for the successful implementation of an in-road RFID reader system in a traffic environment. Initially, it was stipulated that a minimum 4m reading zone on each side of the reader is imperative to facilitate the effective interrogation of five high-speed tags within a designated timeframe. The experiments conducted in this study unveiled that the reader generated an illumination zone exceeding five meters, surpassing the initially stated theoretical requirement by an additional meter.

Moreover, [2] emphasised the necessity for a tag to store a minimum of 640 bits of data while remaining readable at elevated speeds. Given that vehicles often traverse traffic scenarios at speeds reaching up to 180km/h, a proficient vehicle identification system must detect vehicles at such velocities. The extensive tests demonstrated that a 640-bit tag, storing a unique identifier, visual vehicle characteristics, and a digital signature can be reliably read at least once at 180km/h, even when subjected to adverse conditions.

Addressing the computer vision system requirement, it was mandated that the system accurately differentiate between various vehicles based on their visual features. Model testing

substantiated that the vision model adeptly distinguishes between vehicles. The ANPR model achieved very good accuracy, with an overall accuracy of 91.4% for on-road testing. Conversely, the VMMR system encountered challenges in discerning subtle variations among vehicle models from the same manufacturer with an accuracy of 88.5% for on-road testing. Despite this, the overall performance of the vision model remains commendable.

Lastly, the investigation underscored the imperative to establish a system enabling offline verification of the identity of passing vehicles. The research revealed that a roadside controller plays a pivotal role in facilitating offline comparisons between RFID and visual data, thereby confirming the authenticity of the vehicle identity. This innovative approach enhances the robustness of the system, ensuring reliable verification even in scenarios where real-time connectivity may be compromised.

7.2 Summary

This section gives an overview of the diverse chapters presented in this thesis. The initial chapter serves as an introduction, providing the reader with a background to the research topic, elucidating the research problem, proposed solution, project objectives, requirements, scope, and deliverables. Additionally, the first chapter introduces the three primary research questions.

Chapter two embarks on a comprehensive literature review, delving into key facets of RFID and computer vision systems by synthesizing insights from existing literature. This section places particular emphasis on the application of computer vision and RFID for precise electronic vehicle identification. It also explores real-world factors influencing these systems, unveiling theoretical requirements crucial for the vehicle identification system's formulation.

Chapter three transitions to the design phase of the vehicle identification system. Beginning with a conceptualization of the proposed system, this chapter explores the operational intricacies and physical architecture. Subsequently, it delves into the meticulous examination and selection of key hardware and software components aligning with the system's prerequisites, setting the stage for a tangible implementation.

Before advancing to the comprehensive implementation and testing of the in-road system in chapter six, it was imperative to establish a foundational understanding of the hardware components through initial testing and experimentation. This phase, discussed in chapter four, involved basic tests crucial for the detailed design and configuration of the system.

Chapter five marks the progression from preliminary design and initial testing to the detailed design phase. Here, the intricate fusion of RFID and computer vision for a reliable vehicle

identification system is dissected. This chapter features an in-depth discussion of the hardware and software components, ensuring alignment with specific requirements.

The final chapter, chapter six, is dedicated to the formal implementation and testing of the in-road RFID and computer vision system and their integration. It includes various tests illuminating the performance of the in-road RFID system and establishes criteria for successful adoption in practical traffic environments. Additionally, this chapter critically discusses the results, shedding light on limitations within the RFID and vision systems.

7.3 Future research

The work undertaken in this study builds upon the foundation laid by preceding research as outlined in the literature study. The exploration conducted suggests several avenues for potential future research:

- **Examination in High-Density Traffic Environments:** Further investigation in a high-density traffic environment open to the public is recommended. This approach would assess the system's performance in a real-world scenario, potentially featuring vehicle speeds surpassing 180 km/h. Such scrutiny would offer valuable insights into any limitations of the system, paving the way for a nationwide rollout and widespread adoption.
- **Expansion of Training Datasets for ANPR and VMMR:** Enhancing identification accuracy is proposed through the compilation of more extensive training datasets for both Automatic Number Plate Recognition (ANPR) and Vehicle Make and Model Recognition (VMMR). This dataset should encompass in-traffic footage reflecting worst-case conditions, including number plate obstruction, intentional detection avoidance instances, and other forms of unacceptable driver behaviour. The aim is to cultivate a more robust model applicable across diverse traffic environments.
- **Integration of State-of-the-Art ANPR and VMMR Models:** Continuous improvement in performance can be achieved by incorporating the most recent state-of-the-art ANPR and VMMR models for vehicle visual feature extraction. The dynamic landscape of open-source deep learning models, evolving continuously, provides an opportunity to elevate the accuracy and reliability of the system.
- **Development of a Custom Roadside Control Unit:** The creation of a bespoke roadside control unit capable of seamlessly combining RFID and visual data streams for vehicle identification and verification is proposed. This unit should be designed to integrate with existing roadside infrastructure. Additionally, the development of a cloud-based architecture, enabling the roadside unit to operate and communicate data with a cloud server, facilitates both online and offline verification. This holistic approach ensures adaptability to varying verification needs while leveraging existing infrastructure.

8 References

- [1] LicenSys, "fleetVALID Report – Sumner Pilot: 2018," fleetVALID, Queensland, Australia, 2018. Unpublished.
- [2] LicenSys, "fleetVALID Report – RAIN RFID Research," LicenSys, Queensland, 2022. Unpublished.
- [3] S. R. Meneses González and R. Linares y Miranda, "RFID radio channel performance analysis," **Journal of Electrical and Electronic Engineering**, vol. 1, no. 1, pp. 35–40, 2013, doi: 10.11648/j.jee.20130101.14.
- [4] M. Ling-fei, Q. Chun-fang, and T. Xiao-fen, "Velocity analysis for UHF RFID vehicle license plate," **Journal of Engineering**, vol. 1, no. 1, pp. 1–4, 2010.
- [5] Y.-P. Luh and Y.-C. Liu, "Measurement of effective reading distance of UHF RFID passive tags," **Scientific Research Publishing**, vol. 3, no. 3, 2013.
- [6] E. C. Jone, A. Garza, and S. Venkatesan, "RFID readability on different materials in motion testing," University of Nebraska-Lincoln Industrial and Management Systems Engineering.
- [7] P. Nikitin and K. Rao, "Performance limitations of passive UHF RFID systems," in **Proc. 2006 IEEE Antennas and Propagation Society International Symposium**, 2006, doi: 10.1109/APS.2006.1710704.
- [8] A. J. Pretorius, A. J. Hoffman, and Y. Wang, "Geometry-based analysis of an alternative RFID interrogator placement for electronic vehicle identification."
- [9] M. Yu, D. Zhang, Y. Cheng, and M. Wang, "An RFID electronic tag-based automatic vehicle identification system for traffic IoT applications," in **Proc. 2011 Chinese Control and Decision Conference (CCDC)**, 2011, doi: 10.1109/CCDC.2011.5968962.
- [10] H. Marais, M. Grobler, and J. Holm, "Modeling of an RFID-based electronic vehicle identification system," in **Proc. 2013 Africon**, 2013, doi: 10.1109/AFRCON.2013.6757662.
- [11] J. Choi, M. Kang, R. Elmasri, and D. Engels, "Investigation of impact factors for various performances of passive UHF RFID systems," in **Proc. 2011 IEEE Int. Conf. RFID-Technologies and Applications**, 2011, doi: 10.1109/RFID-TA.2011.6068631.
- [12] S. D'Mello, E. Mathews, L. McCauley, and J. Markham, "Impact of position and orientation of RFID tags on real-time asset tracking in a supply chain," **Journal of Theoretical and Applied Electronic Commerce Research**, vol. 3, no. 1, pp. 1–12, 2008, doi: 10.3390/jtaer3010003.

- [13] L. Zhenzhong, M. Nezih, X. George, O. Yuu, and L. Guocheng, "Effects of temperature and humidity on UHF RFID performance," **Smart Materials, Structures & NDT in Aerospace**, vol. 1, no. 1, pp. 1–10, 2011.
- [14] S. Wong and Y. Zheng, "The effect of metal noise factor to RFID location system," in **Proc. 2013 IEEE Int. Conf. Industrial Engineering and Engineering Management**, 2013, doi: 10.1109/IEEM.2013.6962424.
- [15] "Keeping Electronics Cool: Many Options, But Many Constraints As Well," **Electronics360**, 2022. [Online]. Available: <https://electronics360.globalspec.com/article/6115/keeping-electronics-cool-many-options-but-many-constraints-as-well>. [Accessed: Feb. 10, 2022].
- [16] B. Saoudi, "Analysis and study of the performance of coaxial cable passed on different dielectrics," **International Journal of Applied Engineering Research**, vol. 13, no. 3, pp. 1664–1669, 2018.
- [17] A. Almubarak, "The effects of heat on electronic components," **International Journal of Engineering Research and Applications**, vol. 7, no. 5, pp. 52–57, 2017, doi: 10.9790/9622-0705055257.
- [18] M. Rhead, R. Gurney, S. Ramalingam, and N. Cohen, "Accuracy of automatic number plate recognition (ANPR) and real-world UK number plate problems," in **Proc. 2012 IEEE Int. Carnahan Conf. Security Technology (ICCST)**, 2012, doi: 10.1109/CCST.2012.6393574.
- [19] A. Hoffman, D. Geldenhuys, and A. Pretorius, "Securing number plates based on digital signatures and RFID," in **Proc. 16th Int. IEEE Conf. Intelligent Transportation Systems (ITSC 2013)**, 2013, doi: 10.1109/ITSC.2013.6728531.
- [20] A. Rio-Alvarez, J. de Andres-Suarez, M. Gonzalez-Rodriguez, D. Fernandez-Lanvin, and B. López Pérez, "Effects of challenging weather and illumination on learning-based license plate detection in noncontrolled environments," **Scientific Programming**, vol. 2019, pp. 1–16, 2019, doi: 10.1155/2019/6897345.
- [21] LicenSys, "Business Analysis IQPLATE and RCsensor."
- [22] Arizona Department of Transportation, "Cost/Benefit Analysis of Electronic License Plates," Phoenix, 2008.
- [23] Alcatel Lucent, "Building future-proof networks for intelligent transportation systems," 2022.
- [24] J. León-Coca, D. Reina, S. Toral, F. Barrero, and N. Bessis, "Intelligent Transportation Systems and Wireless Access in Vehicular Environment Technology for Developing Smart

Cities," *Big Data and Internet of Things: A Roadmap for Smart Environments,* pp. 285–313, 2014. doi: 10.1007/978-3-319-05029-4_12.

[25] K. Kastell, "Network planning for intelligent transportation systems based on existing wireless networks," in *Proc. 2013 5th Int. Congr. Ultra Modern Telecommunications Control Syst. Workshops (ICUMT),* 2013. doi: 10.1109/icumt.2013.6798400.

[26] M. Matin and M. Islam, "Overview of Wireless Sensor Network," *Wireless Sensor Networks - Technology and Protocols,* 2012. doi: 10.5772/49376.

[27] V. Rajaravivarma, Y. Yang, and T. Yang, "An Overview of Wireless Sensor Network and Applications," in *Proc. 35th South-eastern Symp. Syst. Theory,* Morgantown, WV, USA, 2003, pp. 432–436.

[28] A. Nagpurkar and S. Jaiswal, "An overview of WSN and RFID network integration," in *Proc. 2015 2nd Int. Conf. Electronics Communication Syst. (ICECS),* 2015. doi: 10.1109/ecs.2015.7124956.

[29] L. Zhang and Z. Wang, "Integration of RFID into wireless sensor networks: Architectures Opportunities," in *Proc. 5th Int. Conf. Grid Cooperative Computing Workshops (GCCW),* Oct. 2006, pp. 463–469.

[30] G. Yang, M. Xiao, and C. Chen, "A simple energy-balancing method in RFID sensor networks," in *Proc. 2007 IEEE Int. Workshop Anti-Counterfeiting Security Identification,* Apr. 16–18, 2007, pp. 306–310.

[31] "RAIN flooding: ISO/IEC 20248."

[32] P. Anuradha and R. Sendhilkumar, "Design and implementation of Zigbee-RFID based vehicle tracking," in *Proc. Int. Conf. Sustainable Energy Intell. Syst. (SEISCON),* 2011. doi: 10.1049/cp.2011.0451.

[33] Y. Lai and J. Cheng, "A Cloud-Storage RFID Location Tracking System," *IEEE Trans. Magn.,* vol. 50, no. 7, pp. 1–4, 2014. doi: 10.1109/tmag.2014.2303810.

[34] D. Vinod and A. Mohan, "A Successful Approach to Bus Tracking Using RFID and Low Power Wireless Networks," in *Proc. TENCON 2018 - IEEE Region 10 Conf.,* 2018. doi: 10.1109/tencon.2018.8650342.

[35] A. Hoffman and A. Pretorius, "SmartRoad: A New Approach to Law Enforcement in Dense Traffic Environments," in *Proc. 2015 IEEE 18th Int. Conf. Intell. Transp. Syst.,* 2015. doi: 10.1109/itsc.2015.104.

- [36] M. A. Abdelwahab, "Accurate vehicle counting approach based on deep neural networks," in **Proc. 2019 Int. Conf. Innovative Trends Comput. Eng. (ITCE)**, 2019, pp. 1–5. doi: 10.1109/ITCE.2019.8646549.
- [37] E. Bas, A. Tekalp, and F. Salman, "Automatic vehicle counting from video for traffic flow analysis," in **Proc. 2007 IEEE Intell. Vehicles Symp.**, Jul. 2007, pp. 392–397. doi: 10.1109/IVS.2007.4290146.
- [38] L. Albini, M. Gutoski, and H. Lopes, "Car Make and Model Classification with Deep Learning Methods," in **Proc. 14th Brazilian Congr. Comput. Intell.**, 2020. doi: 10.21528/cbic2019-103.
- [39] F. Chollet, **Deep Learning with Python**, 1st ed. Shelter Island, NY: Manning Publications Co., 2017, pp. 50–300.
- [40] M. D. E. Besbes, "Vehicle identity recognition using deep learning," Ph.D. dissertation, CY Cergy Paris Université; Université de Sfax (Tunisie), 2021.
- [41] D. Greene, P. Cunningham, and R. Mayer, "Unsupervised learning and clustering," in **Springer Berlin Heidelberg**, Berlin, Heidelberg, 2008, pp. 51–90.
- [42] D. G. Lowe, "Object recognition from local scale-invariant features," in **Proc. 7th IEEE Int. Conf. Computer Vision**, vol. 2, pp. 1150–1157, Sep. 1999, doi: 10.1109/ICCV.1999.790410.
- [43] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C.-Y. Fu, and A. C. Berg, "SSD: single shot multibox detector," **CoRR**, abs/1512.02325, 2015.
- [44] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: unified, real-time object detection," in **2016 IEEE Conf. Computer Vision and Pattern Recognition (CVPR)**, pp. 779–788, Jun. 2016.
- [45] J. Redmon and A. Farhadi, "YOLO9000: better, faster, stronger," **CoRR**, abs/1612.08242, 2016.
- [46] C. N. E. Anagnostopoulos, I. E. Anagnostopoulos, V. Loumos, and E. Kayafas, "A license plate recognition algorithm for intelligent transportation system applications," **IEEE Trans. Intelligent Transportation Systems**, vol. 7, no. 3, pp. 377–392, Sep. 2006.
- [47] H. S. Ng, Y. H. Tay, K. M. Liang, H. Mokayed, and H. W. Hon, "Detection and recognition of Malaysian special license plates based on SIFT features," **CoRR**, abs/1504.06921, 2015.
- [48] H. Li, P. Wang, M. You, and C. Shen, "Reading car license plates using deep neural networks," **Image and Vision Computing**, vol. 72, pp. 14–23, 2018.

- [49] S. Z. Masood, G. Shu, A. Dehghan, and E. G. Ortiz, "License plate detection and recognition using deeply learned convolutional neural networks," **CoRR**, abs/1703.07330, 2017.
- [50] J. M. Guo and Y. F. Liu, "License plate localization and character segmentation with feedback self-learning and hybrid binarization techniques," **IEEE Trans. Vehicular Technology**, vol. 57, no. 3, pp. 1417–1424, May 2008.
- [51] N. A. Jalil, A. S. H. Basari, S. Salam, N. K. Ibrahim, and M. A. Norasikin, "The utilization of template matching method for license plate recognition: A case study in Malaysia," in **Adv. Computer and Communication Engineering Technology**, Hamzah A. Sulaiman et al., Eds., Cham: Springer, 2015, pp. 1081–1090.
- [52] S. A. Radzi and M. Khalil-Hani, "Character recognition of license plate numbers using convolutional neural network," in **Visual Informatics: Sustaining Research and Innovations**, H. B. Zaman et al., Eds., Berlin, Heidelberg: Springer, 2011, pp. 45–55.
- [53] T. K. Cheang, Y. S. Chong, and Y. H. Tay, "Segmentation-free vehicle license plate recognition using ConvNet-RNN," **CoRR**, abs/1701.06439, 2017.
- [54] S. M. Silva and C. R. Jung, "Real-time Brazilian license plate detection and recognition using deep convolutional neural networks," in **2017 30th SIBGRAPI Conf. Graphics, Patterns, and Images (SIBGRAPI)**, pp. 55–62, Oct. 2017.
- [55] H. He, Z. Shao, and J. Tan, "Recognition of car makes and models from a single traffic-camera image," **IEEE Trans. Intelligent Transportation Systems**, vol. 16, no. 6, pp. 3182–3192, Dec. 2015, doi: 10.1109/TITS.2015.2437998.
- [56] M. Biglari, A. Soleimani, and H. Hassanpour, "A cascaded part-based system for fine-grained vehicle classification," **IEEE Trans. Intelligent Transportation Systems**, vol. 19, no. 1, pp. 273–283, Jan. 2018, doi: 10.1109/TITS.2017.2749961.
- [57] Q. Hu, H. Wang, T. Li, and C. Shen, "Deep CNNs with spatially weighted pooling for fine-grained car recognition," **IEEE Trans. Intelligent Transportation Systems**, vol. 18, no. 11, pp. 3147–3156, Nov. 2017, doi: 10.1109/TITS.2017.2679114.
- [58] J. Sochor, A. Herout, and J. Havel, "Boxcars: 3D boxes as CNN input for improved fine-grained vehicle recognition," in **2016 IEEE Conf. Computer Vision and Pattern Recognition (CVPR)**, pp. 3006–3015, Jun. 2016.
- [59] L. Albini, M. Gutoski, and H. Lopes, "Car make and model classification with deep learning methods," **Anais do 14th Congresso Brasileiro de Inteligência Computacional**, 2020, doi: 10.21528/cbic2019-103.

- [60] L. Yang, P. Luo, C. C. Loy, and X. Tang, "A large-scale car dataset for fine-grained categorization and verification," in **Computer Vision and Pattern Recognition (CVPR)**, 2015.
- [61] R. Gurney, V. Lyons, M. Rhead, and S. Ramalingam, "The effect of ANPR camera settings on system performance," in **5th Int. Conf. Imaging for Crime Detection and Prevention (ICDP 2013)**, 2013, doi: 10.1049/ic.2013.0276.
- [62] M. Regus, R. Talar, and R. Labudzki, "Indoor positioning and navigation system for autonomous vehicles based on RFID technology," **IOP Conf. Series: Materials Science and Engineering**, vol. 659, no. 1, p. 012059, 2019, doi: 10.1088/1757-899X/659/1/012059.
- [63] Y. Huo, Y. Lu, W. Cheng, and T. Jing, "Vehicle road distance measurement and maintenance in RFID systems on roads," in **2014 Int. Conf. Connected Vehicles and Expo (ICCVE)**, 2014, doi: 10.1109/ICCVE.2014.7297561.
- [64] **Planhillsborough.org**, "Toll gantry levels of design with estimated costs," 2022. [Online]. Available: <https://planhillsborough.org/wp-content/uploads/2015/12/Toll-Gantry-Levels-of-design-with-estimated-costs.pdf>. [Accessed: Jun. 22, 2022].
- [65] Gantry Costing 210730BP, unpublished.
- [66] "Tag Performance Parameters and Test Methods Version 1.1.3," **Gs1.org**, 2008. [Online]. Available: https://www.gs1.org/docs/epc/uhfg2_1_1_0-TestMethod-TagParameters_1_1_3-20080630.pdf. [Accessed: Aug. 11, 2022].
- [67] M. Foegelle, "1-12," **Ets-lindgren.com**, 2022. [Online]. Available: <https://www.ets-lindgren.com/WhitePapers/APM.pdf>. [Accessed: Aug. 17, 2022].
- [68] T.-Y. Lin, A. RoyChowdhury, and S. Maji, "Bilinear CNN models for fine-grained visual recognition," in **2015 IEEE Int. Conf. Computer Vision (ICCV)**, 2015.
- [69] Licensys, "fleetVALID RCSensor – installation guide," fleetVALID, Queensland, Australia, 2018, unpublished.
- [70] "Speeding violations," **Axxonsoft**, [Online]. Available: <https://www.axxonsoft.com/products/integration/trafficcontrol/speeding-violations>. [Accessed: Dec. 8, 2023].
- [71] S. H. Ali, A. H. Reja, and Y. A. Hachim, "Design of a miniaturized wideband disc monopole antenna used in RFID systems," **Indonesian Journal of Electrical Engineering and Computer Science**, vol. 21, no. 2, p. 994, 2021, doi: 10.11591/ijeecs.v21.i2.pp994-1004.

- [72] "Break on removal RFID tags," *KollaKorn*, [Online]. Available: <https://www.kollakorn.com/tamper-evident-break-on-removal-rfid-tags>. [Accessed: Dec. 8, 2023].
- [73] E. Fernandes, R. Correia, A. Gil, J. Postal, and M. R. Gadelha, "Themes validation tool," *HCI International 2019 – Late Breaking Posters*, pp. 16–22, 2019, doi: 10.1007/978-3-030-30712-7_3.
- [74] P. Kumari, "Mastering transfer learning in computer vision: unleashing the power of pre-trained models," *Medium*, [Online]. Available: <https://medium.com/@PriyankaKumari25102002/mastering-transfer-learning-in-computer-vision-unleashing-the-power-of-pre-trained-models-3593102b278d>. [Accessed: Dec. 8, 2023].
- [75] N. Aslam, "Training a CNN from scratch using data augmentation," *Analytics Vidhya*, [Online]. Available: <https://www.analyticsvidhya.com/blog/2022/09/training-a-cnn-from-scratch-using-data-augmentation/>. [Accessed: Dec. 8, 2023].
- [76] A. Taha, "Bilinear CNN models for fine-grained visual recognition," *Medium*, [Online]. Available: <https://ahmdtaha.medium.com/bilinear-cnn-models-for-fine-grained-visual-recognition-b25ba24d3147>. [Accessed: Dec. 8, 2023].
- [77] I. Gorovyi, "Efficient two-step approach for automatic number plate detection," *Int. Journal of Electronics and Telecommunications*, vol. 61, no. 4, pp. 351–356, 2015, doi: 10.1515/eletel-2015-0046.
- [78] "Best ANPR camera setup: Plate recognizer ALPR," *Plate Recognizer*, [Online]. Available: <https://platerecognizer.com/camera-setup-for-best-anpr/>. [Accessed: Dec. 8, 2023].
- [79] "Everything RF," *everything RF*, [Online]. Available: <https://www.everythingrf.com/community/circular-polarization-vs-linear-polarization-for-rfid-applications>. [Accessed: Dec. 8, 2023].
- [80] A. Robinson and C. Venter, "Mastering transfer learning in computer vision," *RISCOSS*, [Online]. Available: <https://www.riscoss.com/mastering-transfer-learning-in-computer-vision/>. [Accessed: Dec. 8, 2023].
- [81] T. I. Group, "Data retrieval via RFID chip," IDePLATE®: The RFID Licence Plate: Tönnjes International Group, <https://www.toennjes.com/en/evi/ideplate/> (accessed Dec. 8, 2023).
- [82] Y. Li and C. Liu, "An approach to instantly detecting fake plates based on large-scale ANPR data," 2015 12th Web Information System and Application Conference (WISA), vol. 61, pp. 287–292, Sep. 2015. doi: 10.1109/wisa.2015.53

- [83] M. Mwango and E. Musonda, "Designing a vehicle tracking system using RFID technology," **International Research Journal of Modernization in Engineering Technology and Science**, vol. 5, no. 2, Feb. 2023. doi: 10.56726/irjmets33579
- [84] A. Ammar, A. Koubaa, W. Boulila, B. Benjdira, and Y. Alhabashi, "A multi-stage deep-learning-based vehicle and license plate recognition system with real-time edge inference," **Sensors**, vol. 23, no. 4, p. 2120, Feb. 2023. doi: 10.3390/s23042120
- [85] A. R. Unterhuber, S. Iliev, and E. M. Biebl, "Estimation method for high-speed vehicle identification with UHF RFID systems," **IEEE Journal of Radio Frequency Identification**, vol. 4, no. 4, pp. 343–352, Dec. 2020. doi: 10.1109/jrfid.2020.2989900
- [86] L. Zou, "Meta-learning for computer vision," **Meta-Learning**, pp. 91–208, 2023. doi: 10.1016/b978-0-323-89931-4.00012-2
- [87] M. M. Taye, "Understanding of machine learning with Deep Learning: Architectures, workflow, applications and future directions," **Computers**, vol. 12, no. 5, p. 91, Apr. 2023. doi: 10.3390/computers12050091
- [88] R. Nisbet, G. Miner, and K. Yale, "Deep learning," **Handbook of Statistical Analysis and Data Mining Applications**, pp. 741–751, 2018. doi: 10.1016/b978-0-12-416632-5.00019-0
- [89] L. Alzubaidi et al., "Review of Deep Learning: Concepts, CNN Architectures, challenges, applications, Future Directions," **Journal of Big Data**, vol. 8, no. 1, Mar. 2021. doi: 10.1186/s40537-021-00444-8
- [90] G. Dipak, "Technology trends and digital solutions for smart cities development," **International Journal of Advance and Innovative Research**, vol. 6, no. 4, Mar. 2019.
- [91] "What is automatic number plate recognition (ANPR)?," **ANPR Cameras**, <https://www.anprcameras.com/about-us/understanding-anpr/> (accessed Nov. 2, 2024).
- [92] W. Riaz et al., "Yolo based recognition method for automatic license plate recognition," 2020 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA), pp. 87–90, Aug. 2020. doi: 10.1109/aeeca49918.2020.9213506
- [93] F. Tafazzoli, H. Frigui, and K. Nishiyama, "A large and diverse dataset for improved vehicle make and model recognition," 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 874–881, Jul. 2017. doi: 10.1109/cvprw.2017.121
- [94] R. Want, "An introduction to RFID technology," **IEEE Pervasive Computing**, vol. 5, no. 1, pp. 25–33, Jan. 2006. doi: 10.1109/mprv.2006.2

- [95] T. M. Zaman, "How does passive RFID works, briefly explained," Jan. 2016. doi: 10.13140/RG.2.2.12361.34402
- [96] A. R. Unterhuber, S. Iliev, and E. Biebl, "Interferences and shadowing in multi-lane UHF RFID-based Vehicle Identification Systems," 2021 IEEE International Conference on RFID Technology and Applications (RFID-TA), pp. 153–156, Oct. 2021. doi: 10.1109/rfid-ta53372.2021.9617393
- [97] T. Rohit and D. Shailesh, "Importance of spatial hierarchy in convolution neural networks," Jan. 2020.
- [98] "Convolutional Neural Networks," *Deepgram*, <https://deepgram.com/ai-glossary/convolutional-neural-networks> (accessed Aug. 3, 2024).
- [99] A. Goyal and R. Bhatia, "Automated car number plate detection system to detect far number plates," *IOSR Journal of Computer Engineering*, vol. 18, no. 04, pp. 34–40, Apr. 2016. doi: 10.9790/0661-1804033440
- [100] A. Kamran & Hanifa, K. Paul, "RFID Applications: An Introductory and Exploratory Study," *International Journal of Computer Science Issues*, vol. 7, no. 3, Jan. 2010. doi: 10.48550/arXiv.1002.1179
- [101] M. Khosy'i'n, S. A. Prasetyowati, Z. Nawawi, and B. Y. Suprpto, "Review and design of GPS-RFID localization for Autonomous Vehicle Navigation," *Proceedings of the 2019 2nd International Conference on Electronics and Electrical Engineering Technology*, pp. 42–46, Sep. 2019. doi: 10.1145/3362752.3362766
- [102] S. Routray, A. K. Ray, and C. Mishra, "Analysis of various image feature extraction methods against noisy image: SIFT, surf and Hog," 2017 Second International Conference on Electrical, Computer and Communication Technologies (ICECCT), pp. 1–5, Feb. 2017. doi: 10.1109/icecct.2017.8117846
- [103] R. Poojary, R. Raina, and A. Kumar Mondal, "Effect of data-augmentation on fine-tuned CNN model performance," *IAES International Journal of Artificial Intelligence (IJ-AI)*, vol. 10, no. 1, p. 84, Mar. 2021. doi: 10.11591/ijai.v10.i1.pp84-92
- [104] Z. Qin, F. Yu, C. Liu, and X. Chen, "How convolutional neural networks see the world --- a survey of Convolutional Neural Network visualization methods," *Mathematical Foundations of Computing*, vol. 1, no. 2, pp. 149–180, Apr. 2018. doi: 10.3934/mfc.2018008
- [105] Y. Zhu, Y. Zhou, Q. Ye, Q. Qiu, and J. Jiao, "Soft proposal networks for weakly supervised object localization," 2017 IEEE International Conference on Computer Vision (ICCV), pp. 1859–1868, Oct. 2017. doi: 10.1109/iccv.2017.204.

- [106] D. Mauricio, J. Pérez, D. Jimenez, L. Rivera, and P. Castañeda, "Recognition of vehicles with cloned plates using deep learning," *J. Southwest Jiaotong Univ.*, vol. 59, no. 3, 2024, doi: 10.35741/issn.0258-2724.59.3.7.
- [107] ISO/IEC, "Information technology—Radio frequency identification for item management—Part 63: Parameters for air interface communications at 860 MHz to 960 MHz," *ISO/IEC Standard 18000-63:2021*, 2021.
- [108] K. V. S. Rao, P. V. Nikitin, and S. F. Lam, "Antenna design for UHF RFID tags: A review and a practical application," *IEEE Trans. Antennas Propag.*, vol. 53, no. 12, pp. 3870–3876, Dec. 2005, doi: 10.1109/tap.2005.859919.
- [109] GS1, *GS1 General Specifications*, Standard Release 23.0, GS1, Jan. 2023.
- [110] ISO/IEC, "Information technology—Automatic identification and data capture techniques—Digital signature meta structure," *ISO/IEC Standard 20248:2018*, 2018.
- [111] J.-I. Cairo, J. Bonache, F. Paredes, and F. Martin, "Interference sources in congested environments and its effects in UHF-RFID systems: A review," *IEEE J. Radio Frequency Identification*, vol. 2, no. 1, pp. 1–8, Mar. 2018, doi: 10.1109/jrfid.2018.2806738.
- [112] L. Mo and H. Zhang, "RFID antenna near the surface of metal," in *Proc. 2007 Int. Symp. Microwave, Antenna, Propag. EMC Technol. Wireless Commun.*, pp. 803–806, Aug. 2007, doi: 10.1109/mape.2007.4393746.
- [113] T. Ben-Jabeur and A. Kadri, "Enhancing passive UHF RFID communication in multipath fading channels using time-varying continuous wave signals," in *Proc. 2015 IEEE Wireless Commun. Networking Conf. (WCNC)*, pp. 574–579, Mar. 2015, doi: 10.1109/wcnc.2015.7127533.
- [114] ThingMagic® M6e Micro UHF RAIN RFID Series, [Online]. Available: <https://www.jadaktech.com/product/thingmagic-m6e-micro-uhf-rain-rfid-series/>. [Accessed: Oct. 8, 2024].
- [115] "ThingMagic micro LTE Carrier Board," *atlasRFIDstore*, [Online]. Available: <https://www.atlasrfidstore.com/thingmagic-micro-lte-carrier-board/>. [Accessed: Oct. 9, 2024].
- [116] P. R. Everitt, "Prediction of asphalt pavement temperatures in South Africa," Civil Engineering Programme, University of Natal, Durban, South Africa.
- [117] "LEPTON7X4 rain RFID reader module," *Impinj*, [Online]. Available: <https://www.impinj.com/partners/caen-rfid/lepton7x4>. [Accessed: Oct. 9, 2024].

- [118] "SparkFun simultaneous RFID reader - m6e nano," SEN-14066, *SparkFun Electronics*, [Online]. Available: <https://www.sparkfun.com/products/14066>. [Accessed: Oct. 9, 2024].
- [119] "3417 rain RFID Reader Module," *Impinj*, [Online]. Available: [https://www.impinj.com/partners/technology-solutions-\(uk\)-ltd/3417-rain-rfid-reader-module](https://www.impinj.com/partners/technology-solutions-(uk)-ltd/3417-rain-rfid-reader-module). [Accessed: Oct. 9, 2024].
- [120] Ego plus mini sticker tag, [Online]. Available: <https://transcore.com/wp-content/uploads/2017/01/eGo%C2%AE-Plus-Mini-Sticker-Tag.pdf>. [Accessed: Oct. 9, 2024].
- [121] "Omni-ID Exo 750 RFID tag," *atlasRFIDstore*, [Online]. Available: https://www.atlasrfidstore.com/omni-id-exo-750-rfid-tag/?srsltid=AfmBOoqnRhU7VKZbwpqMVVR7Uac2UoqmjUkNDlgKT5721MsHtdfAp_w7. [Accessed: Oct. 9, 2024].
- [122] "CONFIDEX steelwave micro II global m4qt," *Cisper Electronics B.V.*, [Online]. Available: <https://www.cisper.nl/en/confidex-steelwave-micro-ii-global-m4qt>. [Accessed: Oct. 9, 2024].
- [123] "Times-7," *atlasRFIDstore*, [Online]. Available: <https://www.atlasrfidstore.com/times-7-rfid-antennas/?srsltid=AfmBOoq1iN5Y4kwTVaSmB3Mla5GsT5Vcb7t9W9fETNWJMW9bU9UoYg8>. [Accessed: Oct. 9, 2024].
- [124] "Impinj threshold (LP) indoor RFID antenna (FCC)," *atlasRFIDstore*, [Online]. Available: https://www.atlasrfidstore.com/impinj-threshold-lp-indoor-rfid-antenna-fcc/?srsltid=AfmBOopoMwIYxoUZhONMdRAX49BKjNW4_UOgoNvLTEPw3VnC-rhGB8r0. [Accessed: Oct. 9, 2024].
- [125] "6 MP Dual Light MD 2.0 Varifocal Bullet Network Camera," *Hikvision*, [Online]. Available: <https://www.hikvision.com/en/products/HiLook-IP-Products/Network-Cameras/Value-Camera/ipc-b660ha-lzu/>. [Accessed: Oct. 9, 2024].
- [126] J. J. Kunnath, "Integration of an in-road RFID reader with a computer vision camera system," Oct. 1, 2024.
- [127] M. Usman Akram, Z. Bashir, A. Tariq, and S. A. Khan, "Geometric feature points based optical character recognition," in *Proc. 2013 IEEE Symp. Ind. Electron. Appl.*, Kuching, Malaysia, 2013, pp. 86–89, doi: 10.1109/ISIEA.2013.6738973.
- [128] D. Yang et al., "Improved Yolov7 network model for gangue selection robot for gangue and foreign matter detection in coal," *Sensors*, vol. 23, no. 11, p. 5140, May 2023, doi: 10.3390/s23115140.

- [129] "Beginner's Guide to VGG16 implementation in Keras," *Built In*, [Online]. Available: <https://builtin.com/machine-learning/vgg16#:~:text=In%20VGG16%2C%20'VGG'%20refers,has%20about%20138%20million%20parameters>. [Accessed: Jan. 13, 2025].
- [130] "Rain RFID Innovation," *LicenSys*, [Online]. Available: <https://www.licensys.com/rain-rfid-innovation/>. [Accessed: Jan. 16, 2025].
- [131] "Tönnjes International Group GmbH," *Tönnjes / Electronic Vehicle Identification*, [Online]. Available: <https://electronic-vehicle-identification.com/>. [Accessed: Jan. 16, 2025].
- [132] "Oorking 1508-199 UHF license plate tag for access," *Elite Gates*, [Online]. Available: <https://elitegates.net/products/doorking-1508-199-uhf-license-plate-tag-for-vehicle-access-gty-50>. [Accessed: Jan. 16, 2025].
- [133] "How much data can RFID tags hold?," *Nexqo*, [Online]. Available: <https://nexqo.com/2023/03/rfid-tag-memory/#:~:text=Passive%20RFID%20tags%20typically%20store,comes%20exclusively%20from%20the%20reader>. [Accessed: Jan. 17, 2025].
- [134] H. Gupta, O. Kotlyar, H. Andreasson, and A. J. Lilienthal, "Robust object detection in challenging weather conditions," in *Proc. 2024 IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, pp. 7508–7517, Jan. 2024, doi: 10.1109/wacv57701.2024.00735.
- [135] F. Sultan et al., "Towards automatic license plate recognition in challenging conditions," *Appl. Sci.*, vol. 13, no. 6, p. 3956, Mar. 2023, doi: 10.3390/app13063956.
- [136] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, Cambridge, MA, USA: MIT Press, 2016.
- [137] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, Jun. 2016, pp. 770–778.
- [138] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [139] A. Vaswani et al., "Attention is all you need," in *Adv. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 5998–6008.
- [140] C. Szegedy et al., "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, Jun. 2016, pp. 2818–2826.

[141] S. Ren et al., "Faster R-CNN: Towards real-time object detection with region proposal networks," *Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.

[142] "Papers with code - stanford cars dataset," Dataset | Papers With Code, <https://paperswithcode.com/dataset/stanford-cars> (accessed Jan. 25, 2025).

[143] "THC-B340-VF: HiLook 4MP EXIR VF bullet camera, 2.8mm – 12mm vari-Focal Lens, day/night switch, 40m IR, IP66, 4 in 1 video output (switchable TVI/AHD/CVI/CVBS), 1 HD analog output," WCCTV, <https://wcctv.co.za/shop/cctv/hilook/hilook-tvi-cameras/hilook-tvi-bullet-cameras/thc-b340-vf-hilook-4mp-exir-vf-bullet-camera-2-8mm-12mm-vari-focal-lens-day-night-switch-40m-ir-ip66-4-in-1-video-output-switchable-tvi-ahd-cvi-cvbs-1-hd-analog-output/> (accessed Jan. 25, 2025).

9 ANNEX A

9.1 Additional information on CNNs

9.1.1 Convolutional layer

This layer of the CNN is essentially a multistage information refining process, where information is passed through successive filters and comes out increasingly 'purified' [39]. The convolutional layer receives the input image and convolves it with the features created by previous filters. Figure 94 shows the process of image convolution.




Figure 94: Image convolution

We can also notice from Figure 95, that the convolutional network transforms the image into different representations that provide more information about the result. The digit four is filtered by different layers to arrive at the final output.




Figure 95: Layer representations learned by a CNN

An important characteristic that gives CNN an advantage above other vision-based techniques, is that convolutional filters learn patterns in a translation invariant fashion. Once a CNN has learned a pattern in a certain position in a picture, it can recognise the pattern anywhere else. The prediction does not depend on where the features are located, but only if they are present. This also makes CNN data efficient i.e. they need less training samples to learn representations compared to alternative networks that do not re-use filter weights in translated versions of the same filters [39].

Convolving filters are controlled by parameters such as zero-padding, channel depth and stride. Zero-padding is used to control the size of the feature map, while the channel depth is the number of filters in the convolution. Increasing the channel depth means more features will be extracted. The stride is the number of pixels in the filter's sliding window [40].

9.1.2 Activation layer

The activation layer controls how the signals flow between different layers. For example, if you apply a sigmoid activation function to the last layer, the network can only predict values between zero and one. The most common activation function is the ReLu (Rectified Linear Unit), but other activation functions include the PReLU and LeakyReLU [40].

9.1.3 Pooling layer

Noise is a big problem for CNNs because filters can be sensitive to noise. A smoothing technique called pooling has been developed to combat the effects of noise. Pooling is achieved by extracting windows over the input feature map and outputting the average or maximum value of each channel. Figure 96 shows an example of the pooling operation.




Figure 96: Average and max pooling operations

9.1.4 Fully connected layer

The fully connected layer is the final layer in the network. These are the layers where all the inputs from one layer is connected to every activation unit of the next layer. The fully connected layer compiles the data extracted from the previous layers to form the final output. Thus, these layers act as classifiers; it uses the features from the convolutional base to classify the input images into classes. Figure 97 provides a representation of a fully connected layer.




Figure 97: Fully connected layer

9.1.5 The Back propagation algorithm

The Back-Propagation algorithm is used to calculate the gradient of the error with respect to all the weights in the network. The algorithm uses gradient descent to adjust all the filter weights to minimize the output error. Back-propagation starts with the final loss value and works backwards towards the bottom layers to compute the contribution of each parameter to the loss value.

9.1.6 Loss function

Optimizing the CNN involves computing a loss value. That is where the loss function comes into play. The loss function takes the predicted representations of the network and the expected network output (target) and computes a score. The goal of training a network is to find the weights that minimize the loss function. The deep learning algorithm uses the score as a feedback signal to adjust the value of the weights in a direction that will lower the loss score. This is where the optimizer comes into play. It is a central deep learning algorithm that implements the Back-propagation algorithm. Figure 98 shows an example of how the loss function and optimizer is used to adjust the weights of the network [39][40].




Figure 98: The feedback signal is used to adjust the weights

At the start the network is initialized with random weights, that accordingly produces a loss score that is very high. As the network processes new sample data, the weights are adjusted a little in the direction that decreases the loss score. This is called the training loop, and when repeated enough times, would yield a network with minimal loss, which outputs results very close to the targets.

9.1.7 Dropout

Dropout is a commonly used regularization technique for neural networks. Dropout is applied to single layers and consists of randomly dropping out, or setting to zero, several output features during training. This is a very effective method to combat overfitting. For example, consider an output layer returning a vector $[0.2, 0.5, 1.3, 0.8, 1.1]$ for a given input sample during training. After applying dropout, the vector will have a few zero entries distributed at random: for example $[0, 0.5, 1.3, 0, 0]$. The dropout rate is the fraction of the features that are dropped during training. The rate is usually between 0.2 and 0.5 [39].