



# **Inverse Synthetic Aperture Radar Imaging of Automotive Targets**

By

**Neeraj Pandey**

(PhD 18102)

Under the supervision of **Prof. Shobha Sundar Ram**

Department of Electronics and Communication Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

February, 2026

# **Inverse Synthetic Aperture Radar Imaging of Automotive Targets**

By

**Neeraj Pandey**

(PhD 18102)

A Thesis

submitted in partial fulfillment of the requirements for the degree of

**Doctor of Philosophy**



Department of Electronics and Communication Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

February, 2026

# Certificate

This is to certify that the thesis titled "*Inverse Synthetic Aperture Radar Imaging of Automotive Targets*" being submitted by *Neeraj Pandey* to the Indraprastha Institute of Information Technology Delhi, for the award of the degree of Doctor of Philosophy, is an original research work carried out by him under my supervision. In my opinion, the thesis has reached the standard fulfilling the requirements of the regulations relating to the degree.

The results contained in this thesis have not been submitted in part or full to any other university or institute for the award of any degree or diploma.

February, 2026



Prof. Shobha Sundar Ram

Professor,

Department of Electronics & Communication Engineering,

Indraprastha Institute of Information Technology Delhi-110020, India.

# Declaration

This is certified that the thesis entitled "*Inverse Synthetic Aperture Radar Imaging of Automotive Targets*" being submitted by me to the Indraprastha Institute of Information Technology Delhi, for the award of degree of **Doctor of Philosophy**, is a bonafide work carried out by me. This research work has been carried out under the supervision of **Prof. Shobha Sundar Ram**. The study pertaining to this thesis has not been submitted in part or in full, to any other University or Institution for the award of any other degree.

February, 2026

*Neeraj Pandey*

Neeraj Pandey

(PhD 18102),

Department of Electronics & Communication Engineering,  
Indraprastha Institute of Information Technology Delhi-110020, India.

# Abstract

Advanced driver assistance systems (ADAS) aim to improve road safety, reduce fatalities, and enable autonomous driving. Modern vehicles rely on multiple sensors such as lidar, cameras, thermal detectors, and radar to detect and classify road users. Among these, millimetre-wave automotive radars offer robust range and velocity estimation, all-weather operation, and unobtrusive bumper integration. High-resolution two-dimensional radar images, and in particular inverse synthetic aperture radar (ISAR) images, can provide detailed information on target size, shape, and motion. However, existing ISAR studies of ground vehicles at automotive radar frequencies have produced only limited datasets for a few targets under controlled conditions, and these datasets are not publicly available. As a result, there is a lack of realistic, large-scale ISAR data suitable for modern machine learning (ML) algorithms for automotive applications.

This thesis develops a framework for simulating high-fidelity ISAR images of automotive targets at millimetre-wave (mm-wave) frequencies. The simulation model incorporates vehicle kinematics, radar scattering phenomenology, range–Doppler clutter, and receiver noise for a 77 GHz automotive radar. Static and dynamic mm-wave clutter for automotive scenarios is modelled using measurements acquired on Indian roads, taking into account different surface types and roughness conditions. The resulting clutter statistics are used to parameterise phenomenological clutter models in the simulator. The framework is validated qualitatively and quantitatively against measurement data gathered from real automotive radars. The simulated ISAR images are then used as inputs to traditional ML classifiers and deep neural networks for the classification of automotive targets; the results show that ISAR radar images are excellent features for accurately classifying different road vehicles.

The thesis further investigates the reliability and interpretability of these classification decisions. It is shown that misclassifications can occur even when noise and clutter are relatively low. To analyse such cases, a method based on counterfactual explanations is proposed, using generative adversarial networks (GANs) to perturb ISAR images of a query class until they are classified as a distractor class, while enforcing that the perturbations remain realistic and consistent with the original class distribution. The resulting counterfactual images provide

physics-based insight into which target regions and micro-Doppler features drive the classifier's decisions. Finally, two application-oriented studies using ISAR images are presented: (i) an automated parking test framework in which an externally mounted radar generates high-resolution images of a car parking in a designated slot and a polynomial trajectory fit is used to assess parking performance; and (ii) an around-the-corner radar (ACR) proof-of-concept at 77 GHz for non-line-of-sight collision avoidance, where sparsity-based dictionary learning is used to separate overlapping range-Doppler returns and reconstruct the signatures of multiple dynamic targets in an urban NLOS configuration.

# Dedication

This thesis is dedicated to my mentors and teachers, whose wisdom shaped my academic path, Almighty God, for His blessings and guidance, my parents, whose love and sacrifices made this journey possible, and my cherished wife, whose support and encouragement are the pillars of my journey.

# Acknowledgements

I am deeply grateful to all those who have contributed to the completion of this thesis. Their support, guidance, and encouragement have been invaluable throughout this journey.

First and foremost, I extend my sincere thanks to my advisor, **Dr. Shobha Sundar Ram**. Her unwavering support, insightful guidance, and persistent encouragement have been fundamental to my growth as a researcher. From nurturing my research capabilities to improving my communication and technical writing skills, her mentorship has been invaluable. Her belief in my potential, even during challenging times, has kept my optimism alive. I am immensely fortunate to have had her as my advisor.

I am also deeply grateful to my supervisory committee members, **Dr. A V Subramanyam** and **Dr. Saket Anand**, for their valuable advice, support, and constructive feedback. Their expertise and insights have significantly enhanced the quality of this thesis. Additionally, I appreciate the comments and suggestions from **Dr. Debora Pastina** at Sapienza University of Rome, which have further refined my work.

Special thanks are due to my lab mate, **Dr. Shelly Vishwakarma**, for her enthusiastic collaboration and meticulous work during our hardware experiments. Our countless brainstorming sessions have been incredibly productive and enlightening. She has been a wonderful friend and counselor, always providing support and helping me stay calm during stressful times. I have had the pleasure of working with my other lab mates and friends: **Kainat Yasmin**, **Akanksha Sneh**, **Gaurav Duggal**, and **Yoshana**. Conducting research with them has been full of fun and excitement. They have not only been my collaborators but have helped me broaden my horizons, too.

I extend my heartfelt gratitude to my friends, **Ashutosh Vaish**, **Neha Jain**, **Neelam Singh**, **Saumya Chaturvadi**, and **Amina**, for making my journey mem-

orable. The time spent with them has been filled with laughter and camaraderie, making the arduous path of research more bearable. Special acknowledgment goes to **Ashutosh Vaish** for his dedicated support during radar measurement data collection and insightful discussions on various concepts.

I would like to express my deep appreciation to my parents for their sacrifices, strong determination, and unwavering support. Their constant encouragement and unconditional love have shaped me into the person I am today. To my dear wife, **Neha Pandey**, I am forever grateful. Her understanding, love, care, and support have been my anchor through every moment. Her presence in my life is a blessing, and words cannot fully express my appreciation.

I also extend my heartfelt thanks to my extended family—my father-in-law and mother-in-law—for their unconditional love, care, and support.

Special thanks to **Dr. Anand Singh** for generously providing his car for radar measurement data collection and driving during these sessions. Additionally, I express my gratitude to **Dr. Sumeet Patiyal** for his assistance and support during the measurement data collection process.

I would also like to thank all my mentors, teachers, and friends who have guided and supported me throughout my academic journey. Their wisdom and encouragement have been crucial in my development.

Finally, I wish to thank **Indraprastha Institute of Information Technology-Delhi** for providing a conducive environment for research and learning. The resources and support from the institute have been instrumental in the completion of this work.

If I have inadvertently omitted anyone, please accept my sincere apologies. Your contributions, whether big or small, have been deeply appreciated.

# Contents

- Abstract** **i**
  
- Acknowledgements** **iv**
  
- List of Figures** **x**
  
- List of Tables** **xviii**
  
- List of Abbreviations** **xx**
  
- 1 Introduction** **1**
  - 1.1 Motivation . . . . . 1
  - 1.2 Research Objectives . . . . . 10
  - 1.3 List of Publications . . . . . 12
  - 1.4 Notation . . . . . 14
  
- 2 High Fidelity ISAR Image Simulation and Classification of Automotive Targets** **15**
  - 2.1 Simulation Methodology . . . . . 17
    - 2.1.1 Automotive Animation Model . . . . . 17
    - 2.1.2 Advanced Point Scatterer Signal Model . . . . . 21
    - 2.1.3 Noise and Range-Doppler Clutter Models . . . . . 29

2.1.4	Database of Simulated Data . . . . .	34
2.1.5	Study of effects of shadowing on radar signatures . . . . .	47
2.2	Measurement Data . . . . .	49
2.3	Classification Results . . . . .	55
2.3.1	Effect of noise and clutter on classification performance	55
2.3.2	Effect of test and training percentages . . . . .	57
2.3.3	Testing measured data with classifiers trained with simulation data . . . . .	61
2.4	Summary of findings . . . . .	62
<b>3</b>	<b>Characterization of Millimeter Wave Clutter on Indian Roads at Automotive Radar Frequencies</b>	<b>65</b>
3.1	Experimental Setup . . . . .	67
3.2	Analysis and Observations . . . . .	74
3.3	Summary of Findings . . . . .	78
<b>4</b>	<b>Generating Counterfactual Explanations For Misclassification of Automotive Radar Targets</b>	<b>80</b>
4.1	Related work . . . . .	82
4.2	Methodology . . . . .	84
4.2.1	GAN Architecture . . . . .	87
4.2.2	Evaluation metrics . . . . .	90
4.3	Experimental Setup . . . . .	93
4.4	Network Design and Hyperparameter Tuning . . . . .	97
4.4.1	Autorickshaw to Bicycle Transformation . . . . .	100
4.4.2	Car to Autorickshaw Transformation . . . . .	101
4.4.3	Car to Bicycle Transformation . . . . .	101

4.4.4	Classifier training . . . . .	102
4.4.5	Auto-encoder training . . . . .	102
4.5	Experimental Results and Discussion . . . . .	103
4.5.1	Misclassification of autorickshaw ISAR images as bicycle	106
4.5.2	Misclassification of car ISAR images as autorickshaw .	109
4.5.3	Misclassification of car as bicycle . . . . .	113
4.6	Summary of findings . . . . .	115
<b>5</b>	<b>Applications of ISAR Images of Automotive Targets</b>	<b>117</b>
5.1	Automated parking test using ISAR images . . . . .	117
5.1.1	Methodology . . . . .	120
5.1.2	Experimental Setup . . . . .	124
5.1.3	Results . . . . .	126
5.2	Disaggregation and Detection of Multiple Automotive Targets for mm-wave Around-the-Corner Radar . . . . .	128
5.2.1	Methodology . . . . .	133
5.2.2	Experimental Setup . . . . .	136
5.2.3	Measurement Results . . . . .	138
5.3	Summary of findings . . . . .	146
<b>6</b>	<b>Conclusion and Future Work</b>	<b>148</b>
6.1	High-Fidelity ISAR Image Simulation Framework for Automot- ive Targets at mm-wave Radar Frequency . . . . .	148
6.2	Modeling Clutter in Automotive Radar Scenarios . . . . .	149
6.3	Public Database of Automotive Radar Data . . . . .	149
6.4	Classification of Automotive Targets Using ISAR Images . . . . .	150
6.5	Explainable AI in Radar Image Classification . . . . .	151

6.6	Applications . . . . .	151
6.7	Future Research Avenues . . . . .	152
6.7.1	Real-World Testing and Validation . . . . .	152
6.7.2	ISAR Classification and Data Verification . . . . .	153
6.7.3	Advanced Explainable AI Techniques . . . . .	153
6.7.4	Radar-Camera Sensor Fusion . . . . .	154
6.7.5	Enhanced Diversity in Radar Systems and Algorithms . . . . .	154
	<b>References</b>	<b>156</b>

# List of Figures

2.1	Three-dimensional automotive targets with triangular facets used for the simulation. (a) Auto-rickshaw (tuk-tuk), (b) full-sized car (c) mid-sized car, (d) truck (e) bicycle. <i>Color note:</i> Colors denote facet/patch groups in the CAD mesh for visualization only; they do not represent amplitude, material, or any physical parameter. . . . .	18
2.2	(a) Road geometry of four-way traffic junction. (b) Trajectories undertaken by the automotive target in a four-way junction - (i) Right turn, (ii) Left turn, (iii) U-turn and (iv) Straight through. .	19
2.3	Animation model of trajectory: (a) Way points along the trajectory that will be traversed by the center of the vehicle; (b) Wheels rotational angle calculation. . . . .	20
2.4	Radar cross section of flat metallic triangular plate . . . . .	24
2.5	ISAR images of a <i>bicycle</i> with +10 dB SNR at CPI frames corresponding to 1.5, 2.5, 3.5 and 4.5 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from west to south and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm). . . . .	35

2.6	ISAR images of a <i>bicycle</i> with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1.5, 2.5, 3.5 and 4.5 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from west to south and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm). . . . .	36
2.7	ISAR images of an <i>auto-rickshaw</i> with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm). . . . .	37
2.8	ISAR images of an <i>auto-rickshaw</i> with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm). . . . .	38
2.9	ISAR images of a <i>mid-size car</i> with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm). . . . .	39

2.10	ISAR images of a <i>mid-size car</i> with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm). . . . .	40
2.11	ISAR images of a <i>full-size car</i> with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm). . . . .	42
2.12	ISAR images of a <i>full-size car</i> with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm). . . . .	43
2.13	ISAR images of a <i>truck</i> with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm). . . . .	44

2.14	ISAR images of a <i>truck</i> with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).	45
2.15	(i-v) ISAR images of auto-rickshaw for the same CPI frame using Bernoulli’s random variable with visibility of 60 %, 50%, 40 %, 30%, 20 % respectively. (vi) is the ray tracing algorithm ISAR images of auto rickshaw for the same CPI frame. The range axis spans from 0 to 20m while the cross-range axis spans -15 to +15m.	48
2.16	Experimental setup for gathering ISAR images of (a) an auto-rickshaw undertaking a right turn trajectory, (b) a bicycle undertaking a straight trajectory, (c) a small sized car undertaking a left turn along their trajectory (d) Radar hardware.	50
2.17	Experimental setup and trajectories for gathering ISAR images of (a) two straight trajectories before the radar, (b) two right turn trajectories before the radar, (c) a left turn trajectory before the radar.	51
2.18	((i-iv) ISAR images of bicycle along the left turn trajectory at CPI frame corresponding to 1,2,3,4 s, (v-viii) ISAR images of auto-rickshaw along the straight trajectory at CPI frame corresponding to 1,2,3,4 s, and (ix-xii) ISAR images of car along the straight trajectory trajectory at CPI frame CPI frame corresponding to 1,2,3,4 s from measurement data. Range axis spans from 0 to 20m while the cross range axis spans -15 to +15m.	53

2.19	Qualitative comparison of simulated and measured ISAR images of a <i>bicycle</i> . Sub-figures (i) and (ii) show simulated images generated using the proposed ISAR simulator for two representative CPI frames, while sub-figures (iii) and (iv) show the corresponding ISAR images obtained from measurement data. In all four sub-figures, the micro-Doppler streaks due to the rotating wheels and the overall range and cross-range spread are clearly visible. The simulated signatures closely match the measured ones, indicating good qualitative agreement between the simulated and real data. . . . .	54
2.20	Classification accuracy of automotive target ISAR images using SVM, RF and transfer learning algorithms based on Alexnet and Googlenet for (a) differing SNR values and (b) different range-Doppler clutter values. 70% data used for training and remaining data used for validation and testing. . . . .	56
3.1	Experimental setup to calibrate the mm-wave radar at 77GHz using a trihedral corner reflector as a calibration target of known radar cross-section. . . . .	68
3.2	Different locations at which road surface clutter coefficient was estimated. Locations 1, 2, and 3 correspond to 7m wide concrete road, while location 4 is a 9m wide asphalt road. . . . .	71
3.3	Experimental setup to estimate the road roughness using Bosch GLM 500 laser meter mounted on a sliding scale of 1m length and 1.008m height. . . . .	73
3.4	Distribution plots for the $\sigma^0$ with respect to different type of road location (a) location 1 concrete road with roughness 6.4 mm, (b) location 2 concrete road with roughness 7.7 mm, (c) location 3 concrete road with roughness 8.7 mm (d) location 4 asphalt road with roughness 6.7 mm. . . . .	78
3.5	Plots of mean $\sigma^0$ with respect to different radar heights with three different temperatures (a) location 1 concrete road with roughness of 6.4 mm, (b) location 2 concrete road with roughness of 7.7 mm, (c) location 3 concrete road with roughness 8.7 mm (d) location 4 asphalt road with roughness 6.7 mm. . . . .	79

4.1	System diagram of the proposed framework for explaining the misclassification of ISAR images belonging to the $T_1$ class as $T_2$ . It illustrates the flow of input ISAR images of the $T_1$ class to an RWGAN, generating counterfactual images that are now classified as $T_2$ class. Following this, EfficientNetB0 is used to extract low-dimensional features from both electromagnetically simulated images and neural network generated images. Similarity metrics are computed between both sets of images to obtain counterfactual explanations for the misclassification of $T_1$ class images as $T_2$ class images. . . . .	86
4.2	RWGAN Architecture for synthesizing counterfactual images that are realistic but misclassified. . . . .	88
4.3	Computer-aided models of (a) car, (b) autorickshaw, and (c) bicycle, rendered with multiple triangular facets. The facets of the car are divided into seven parts, the autorickshaw into five parts, and the bicycle into three, as indicated by the different colors. Numeric labels (1–7 for the car; 1–5 for the autorickshaw) appear in (a) and (b). The bicycle in (c) is color-coded without numeric labels because it is not included in the physics-based occlusion dataset (see Table 4.1). . . . .	95
4.4	Architecture of the (a) generator and (b) discriminator network. Here, $n, k, s$ are the number of filters, the kernel size, and the stride, respectively. . . . .	100
4.5	Architecture of the classifier $C$ . Here, $n, k, s$ are the number of filters, the kernel size, and the stride, respectively. . . . .	100
4.6	Architecture of the (DAE). In this design, $n, k,$ and $s$ denote the number of filters, the kernel size, and the stride, respectively.	101
4.7	(a) Prediction gain, and (b) Realism score plots for three different cases, input class $T_1$ to $G$ and targeted (generated) class $T_2$ for 1000 iterations. . . . .	105

4.8	(a) Input ISAR images (i-xvi) represent ISAR images of $T_1$ class vehicles (autorickshaw), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as $T_2$ class vehicles (bicycle), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB. . . . .	106
4.9	Enlarged view of representative input–counterfactual pairs for the autorickshaw-to-bicycle misclassification case. Images (i) and (ii) show input ISAR images of $T_1$ class vehicles (autorickshaw) at two different orientations. Images (iii) and (iv) show the corresponding counterfactual ISAR images generated by the RWGAN and classified as $T_2$ class vehicles (bicycle). The range and cross-range span are 5 m with a normalized dynamic range of 60 dB. . . . .	108
4.10	(a) Input ISAR images (i-xvi) represent ISAR images of $T_1$ class vehicles (car), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as $T_2$ class vehicles (autorickshaw), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB. . . . .	111
4.11	Enlarged view of representative input–counterfactual pairs for the car-to-autorickshaw misclassification case. Images (i) and (ii) show input ISAR images of $T_1$ class vehicles (car) at two different orientations. Images (iii) and (iv) show the corresponding counterfactual ISAR images generated by the RWGAN and classified as $T_2$ class vehicles (autorickshaw). The range and cross-range span are 5 m with a normalized dynamic range of 60 dB. . . . .	113
4.12	(a) Input ISAR images (i-xvi) represent ISAR images of $T_1$ class vehicles (car), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as $T_2$ class vehicles (bicycle), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB. . . . .	114
5.1	Three trajectories for the parking test for (a) $45^\circ$ angle parking and (b) perpendicular parking. . . . .	123

5.2	The experiment setup for the parking test radar is situated at the location (0,0,0.5)m, and the car is parking at perpendicular parking.	124
5.3	ISAR images of Ford Figo carrying out (a) perpendicular parking and (b) angle parking at 3.5, 4.0, 4.5, 5s. (1-iv) Top, (v-viii) middle, and (ix-xii) bottom rows in both images are generated for car following correct trajectory, incorrect trajectory-1, and incorrect trajectory-2 respectively. . . . .	125
5.4	Bottom row shows ISAR images of Honda Brio carrying out (a) perpendicular parking and (b) angle parking, at 3.5, 4.0, 4.5, 5s. (i-iv) Top, (v-viii) middle, and (ix-xii) bottom rows in both images are generated for car following correct trajectory, incorrect trajectory-1, and incorrect trajectory-2 respectively. . .	125
5.5	Experimental set up and trajectories of the automotive targets. .	137
5.6	ACR scenario of (a) bike (b) pedestrian and (c) LOS scenario of car and bike. . . . .	137
5.7	Range-Doppler ambiguity plots of automotive targets generated with a 77GHz automotive radar in ACR and LOS conditions. Top row: small size car; second row: bike; third row: human; fourth row: car and bike; fifth row: car and human and sixth row: car and human. The white squares in the figures depict the position of the targets. . . . .	141
5.8	Trajectories followed by (a) car and bike, (b) car and human, (c) human and bike shown in a(i), b(i) and c(i) and their respective range-Doppler ambiguity plots shown in a(ii), b(ii) and c(ii). Ground truth plots of <i>car</i> for cases (a) and (b) shown in a(iii) and b(iii); <i>bike</i> for cases (a) and (c) shown in a(v) and c(iii); <i>human</i> for cases (b) and (c) shown in b(v) and c(v). Reconstructed plots of <i>car</i> , <i>bike</i> and <i>human</i> for the respective cases shown in a(vi)-c(vi).	143

# List of Tables

2.1	Automotive radar parameters for generating ISAR images . . . .	32
2.2	Simulated ISAR image database . . . . .	46
2.3	Structural similarity index measure (SSIM) in ISAR images generated using ray tracing and Bernoulli random variable . . . .	48
2.4	Classification of ISAR images using classical ML algorithms - SVM and RF, and deep learning based algorithms - Alexnet and Googlenet . . . . .	57
2.5	Confusion matrices from SVM, RF, Alexnet and GoogleNet classifiers based on 70% training, 15% validation and 15% test data. . . . .	59
2.6	Confusion matrices from SVM and Alexnet and GoogleNet classifiers, trained with simulation data and test with measured data. . . . .	61
2.7	Hyper-parameters used for the SVM and Random Forest classifiers.	62
3.1	TI-AWR 1843 radar parameters . . . . .	69
3.2	Road roughness variance . . . . .	74
3.3	Negative log likelihood (NLL) scores and BIC scores for different road surfaces. . . . .	77
4.1	Total Number of Occluded Images for Car and autorickshaw . .	96
4.2	Classifier training results . . . . .	103
4.3	Autoencoder training . . . . .	103

4.4	Counterfactual explanations for misclassification of autorickshaw as a bicycle. . . . .	110
4.5	Counterfactual explanations for misclassification of a car as an autorickshaw. . . . .	112
4.6	Counterfactual explanations for misclassification of a car as a bicycle. . . . .	114
5.1	Result of parking test algorithm using Ford Figo data for training and Honda Brio data for test. . . . .	128
5.2	Number of detections and false alarms for two target scenario. . . . .	145

# List of Abbreviations

<b>SAE</b>	Society of Automotive Engineers
<b>ADAS</b>	Advanced Driver Assistance Systems
<b>NLOS</b>	Non-Line-of-Sight
<b>ACR</b>	Around-the-Corner Radar
<b>CNN</b>	Convolutional Neural Networks
<b>SAR</b>	Synthetic Aperture Radar
<b>ISAR</b>	Inverse Synthetic Aperture Radar
<b>DCNN</b>	Deep Convolutional Neural Networks
<b>FFT</b>	Fast Fourier Transform
<b>FMCW</b>	Frequency Modulated Continuous Wave
<b>ATR</b>	Automatic Target Recognition
<b>FDTD</b>	Finite Difference Time Domain
<b>AI</b>	Artificial Intelligence
<b>SVM</b>	Support Vector Machine
<b>RF</b>	Random Forest
<b>LFM</b>	Linear Frequency Modulated
<b>RCS</b>	Radar Cross Section
<b>GMM</b>	Gaussian mixture modal
<b>NLL</b>	Negative Log-Likelihood
<b>BIC</b>	Bayesian Information Criteria

**CPI** Coherent Processing Interval

**PRI** Pulse Repetition Interval

**SNR** Signal-to-Noise Ratio

**SICNR** Signal-to-Interference-plus-Clutter-and-Noise Ratio

**SSIM** Structural Similarity Index Measure

**NMSE** Normalized Mean Square Error

**GAN** Generative Adversarial Network

**WGAN** Wasserstein Generative Adversarial Network

**DAE** Denoising Autoencoder

**VUT** Vehicle Under Test

# Chapter 1

## Introduction

### 1.1 Motivation

**Road safety and motivation for ADAS.** Road safety is a significant public health issue. The most recently published global status report on road safety in 2023 by the World Health Organization (WHO) highlighted that the number of road traffic deaths reached 1.19 million in 2021, with nearly 3700 people dying every day on the road, and almost 50 million more injured or disabled every year [1,2]. Road traffic accidents are the 12<sup>th</sup> leading cause of death and the leading cause of death for children and young adults between 5-29 years of age [2]. The factors that potentially impact road safety are unsafe vehicles, speeding, driving under the influence, alcohol and other psychoactive substances or drugs, non-use of seat belts, helmets, child restraints, distracted driving, and unsafe road infrastructure [3]. Due to these factors, government regulations for vehicle safety are rising across different parts of the globe [4–6], alongside a surge in the demands for new safety features in vehicles. The strategies adopted in the

last few decades for improving road safety are (1) The adoption of technology-driven solutions using advanced driver assistance systems (ADAS) leveraging sensors such as radars, lidars, cameras, and new algorithms for enforcing safety mechanisms while driving; (2) The second strategy is to ensure that better drivers use the roads, who are well-trained and thoroughly tested; (3) and the third strategy is to improve the road infrastructure.

**Sensors for ADAS and environment perception.** The Society of Automotive Engineers (SAE) defines six-level industry standards for driving autonomy. These standards range from level 0 fully manual to level 5 fully autonomous [7], and there has been significant research on ADAS for supporting the evolution of these standards [8,9]. One critical challenge in ADAS is detecting and recognizing static automotive targets such as poles, traffic signs, overhead bridges, guardrails, buildings, and parked vehicles and dynamics targets such as pedestrians, bicycles, bikes, auto-rickshaws, cars, and trucks. This task is assisted with information from one or more sensors on the ego vehicle, such as lidar, camera, ultrasonic, and radar. Applications that commonly rely on vision-based sensors include detecting road signs and traffic lights, surround-view systems, lane changing, and driver monitoring systems. Lidar is used with high precision and accuracy for three-dimensional high-resolution point cloud data of the surrounding environment. The point cloud data is used for accurate localization and mapping. The limitation of lidar is its high cost and maintenance [10], and it is hard to integrate into the bumper of cars.

**Role of automotive radar in ADAS.** The automotive radar is a low-cost, highly

reliable perceptive sensor in ADAS systems that enables range measurement and target motion detection. Automotive radars work on 24 GHz and 77 GHz bands. They are classified as short-range radar (up to 40 m), mid-range radar (up to 100 m), and long-range radar (up to 250 m) [10, 11]. These radars are used for multiple applications such as automatic cruise control, pedestrian detection, cross-traffic alert, blind-spot detection, and parking assistance [12–14]. Due to the widespread use of automotive radar sensors in ADAS systems, the global automotive radar market size is projected to reach USD 3.13 billion by 2027 [15]. The main advantages of automotive radar over cameras for object detection and classification are: First, the radar can operate in low visibility conditions, such as rain and fog. Secondly, these sensors are typically cheaper than cameras; hence, multiple of them can be mounted around the vehicle’s periphery, usually behind the bumpers. These sensors can be used in non-line-of-sight (NLOS) conditions, specifically in around-the-corner radar (ACR) scenarios [16]. Finally, automotive radars operate at millimeter-wave (mm-wave) frequencies with high bandwidths and electrically large antenna apertures. Hence, they offer excellent range, Doppler, and azimuth resolution [11].

**From point clouds to radar images.** Usually, the raw radar data cube in these systems is processed to provide a collection of point scatterers corresponding to both vehicles and clutter with range, azimuth, elevation, and Doppler information. Direct object detection and classification based on this data type can be challenging since it is difficult to correctly cluster the point scatterers that correspond to the same object [17, 18]. We explore the hypothesis that radar

images/signatures directly processed from raw radar data provide more valuable features for automatic target recognition.

**Radar signatures and ISAR imaging.** Classification of radar targets for a variety of applications has been researched over the last few decades [19–25]. Many different types of radar signatures have been studied. For example, radar micro-Doppler spectrograms, the joint time-frequency representations of time-domain narrowband radar data, have proven to provide excellent features for classification. They have been used for distinguishing between different types of human activities [26–29]; armed and unarmed personnel [30]; ground vehicles and pedestrians [31,32]; and different types of airborne targets such as unmanned aerial vehicles and birds. [33,34] for diverse applications like assisted living, elderly care and security, and surveillance. With broadband radar data from frequency-modulated continuous waveforms, range-Doppler plots have been generated that have also served as excellent features for target recognition [35]. Other works have used range-crossrange images generated through synthetic aperture radar (SAR) imaging for classification purposes [36–38]. However, SAR images are typically more suited for classifying static targets since dynamic targets may distort the radar images. The alternative is to use antenna array processing to obtain fine cross-range resolution. However, this requires a large array with many antenna elements and precise phase synchronization across the multiple channel data [39]. A third method for obtaining fine cross-range resolution is to use inverse synthetic aperture radar (ISAR) processing of single-channel broadband data. When a dynamic target travels along a complex

trajectory, the target undergoes a combination of translational and rotational motion. If the translational motion of a target can be correctly estimated and compensated, then the Doppler dimension can be mapped to cross-range to obtain ISAR images [40,41].

**ISAR for automotive targets.** The traditional arrangement of ISAR is to consider a stationary monostatic broadband radar and a dynamic target undergoing rotational motion. The large bandwidth enables fine downrange resolution, and the rotational motion provides the radar with a large synthetic aperture for obtaining fine cross-range resolution. The cross-range resolution is inversely proportional to the extent of the target aspect presented to the radar during the rotational motion and the coherent processing time interval. The advantage of ISAR is that high-resolution images of targets are realized with single-channel sensor data without the requirement of costly and complex hardware with large antenna arrays with lots of elements and receiver channels. ISAR images, generated from range-Doppler plots of dynamic targets, have been extensively researched over the last two decades - especially for the detection and classification of airborne targets and ships [42–44]. There is growing interest in ISAR imaging of ground-moving vehicles, especially at automotive radar frequencies. ISAR imaging of ground-moving vehicles with a stationary radar was realized with experimental data in [45–47] and high-resolution ISAR imaging of a car was carried out by capturing turntable measurement data in [48], ISAR images of ground vehicles were realized with airborne sensors in [49–53]. In the automotive target scenario, vehicles undergo complex turns, which can result in a large radar aspect angle.

Further, even while moving along a straight path, a slight offset of the target vehicular trajectory from the ego radar can result in significant radar aspects to get a fine cross-range resolution. Since automotive short-range radars are characterized by large bandwidths (above 2GHz) and high carrier frequencies (77GHz) that result in fine range and Doppler (or cross-range) resolution, they are particularly suited for generating high-resolution ISAR images of vehicles. In ADAS systems, multiple auxiliary sensors (gyrometers, accelerometers, other radars) are deployed on the ego vehicle. Therefore, the translational motion compensation of both the ego and target vehicles can be carried out easily. In all of the above works, ISAR images were generated for a few automotive targets under controlled laboratory conditions. However, these limited data sets are not available to the radar community and are insufficient to support state-of-the-art target recognition algorithms.

**Need for large datasets and simulation models.** Of late, the radar community has been increasingly turning to learning algorithms based on deep convolutional neural networks (DCNN) for automatic target recognition (ATR) [37, 54–56]. In multiple applications, these algorithms have demonstrated their superiority in feature extraction and classification performance, especially with respect to traditional approaches of heuristic parameter selection for machine learning (ML)-based target recognition. However, the performance of DCNNs relies on the availability of large training databases gathered in diverse scenarios. The data set must comprise different types of automotive targets, such as vehicles of diverse sizes, moving along different trajectories and orientations with respect

to radar. Similarly, there must be diversity in the road conditions in terms of clutter and interference. There are two possible methods for generating training data sets. The first method is to collect the data in real traffic and target scenarios using actual automotive radar sensors. The advantage is that the training data is realistic and can be gathered in laboratory conditions and during test drives. However, the disadvantage is that data collection and labeling are time-consuming and laborious. Further, the database must be updated based on hardware modifications to the sensor or software upgrades to the signal processing. Also, the captured data may be corrupted in the presence of clutter in the environment and by sensor hardware limitations. The second method to generate large training data is to simulate the sensor, target, and traffic conditions. The advantage of simulation is that the training data set can be generated rapidly for diverse conditions. Also, the simulator can be easily integrated with the radar test bed and signal processing platforms for rapid prototyping and validations. Finally, the simulator can be helpful to study the effects of different radar phenomenology, including sensor position and parameters, clutter, and road conditions on the radar signatures. While several prior works have described simulation models of pedestrians [57, 58], there are very few works that model automotive vehicles along complex trajectories [59, 60]. These works have confined their scope to simulating high-range resolution profiles and micro-Doppler spectrograms. We propose a single channel-based, high-resolution range-Doppler realistic ISAR images simulation model in this thesis in chapter 2.

**Motivation for automotive clutter modelling.** To make the ISAR simulation model optimum, we require an in-depth understanding of the radar phenomenology of the targets in automotive scenarios. So, we need an accurate model for the static and dynamic clutters present in the road environment. In the literature, there have been few measurements conducted at 77 GHz for the different types of road surfaces (Asphalts, concrete, and gravel surface) [61–64] for the automotive clutters. However, very little work has been done on the range-Doppler-based land clutter model in the literature for the different automotive scenarios. This thesis will investigate the automotive cluster model for Indian roads in varying road roughness and wet conditions to improve the radar simulation model and make it more realistic in chapter 3.

**Motivation for explainable AI in automotive radar.** The explainability of AI-based models is a critical consideration for the broad commercial adoption of artificial intelligence (AI) techniques in radar, particularly within the automotive context. Explainability denotes the process of making a classifier’s decision-making understandable to humans. Several motivations drive the development of explainable AI models: First, the inherent opacity of deep learning classifiers raises legal implications since misclassification can lead to severe consequences. Therefore, explaining the reasons for classifications, especially misclassifications, facilitates responsibility to the correct actor. Second, the widespread adoption of AI-based decisions will only be accepted if there is trust that these decision-making processes are fair and accountable. Finally, enhancing the interpretability of machine-made decisions enables detecting and rectifying errors or bugs in

models, contributing to improved classification performance. Recent efforts have developed various approaches for explainable AI algorithms [65–67]. One approach involves quantifying the importance of different data features contributing to a model’s overall prediction performance [68, 69]. Another strategy creates simplified interpretable models that explain the performance of opaque deep neural network classifiers [70, 71]. Some methods aim to generate logical or mathematical explanations for AI models [66]. Counterfactual explanations, which illustrate how slight alterations to input data can affect a model’s behavior, provide another perspective [72]. For example, in automotive radar, counterfactual explanations can address whether the occlusion of certain vehicle parts will result in a different classification. Note that explainable AI differs from studying the general degradation of classification performance due to deteriorating signal-to-interference, clutter, and noise ratio (SICNR). For instance, while high SICNR allows for clear class separation, low SICNR mixes data, degrading classification. In this thesis, we examine how specific perturbations can shift data points between classes, even under high SICNR, providing counterfactual explanations for changes in classifier decision-making in chapter 4.

**Emerging applications and thesis overview.** As automotive radar technology advances, there is a growing need to explore its potential beyond traditional applications like navigation and collision avoidance. In this thesis, we examine two applications of radar imaging to address emerging challenges in urban environments. The first application explores the use of ISAR images generated by a single-channel externally mounted automotive radar system. We propose

a novel approach for automating parking tests, specifically for angle and perpendicular parking scenarios, by developing an algorithm that leverages these ISAR images to assess and verify the correctness of parking maneuvers, thereby improving parking accuracy and efficiency. The second application investigates the implementation of ACR sensing using radar technology, particularly W-band mm-wave radars, in NLOS and ACR scenarios. We demonstrate the feasibility and effectiveness of this approach in enhancing situational awareness, collision avoidance, and hazard detection in complex urban environments. These applications are discussed in chapter 5.

Chapter 6 summarizes the key findings and contributions of the research. It discusses potential avenues for future research and development in automotive radar technology, emphasizing advancements in ISAR simulation, explainable AI, and radar-based applications.

## **1.2 Research Objectives**

The research objectives for this thesis are the following.

1. Develop a high-fidelity simulation framework to generate realistic ISAR images of automotive targets, incorporating the effects of additive noise in the radar receiver: The main objective is to create a simulation framework for rapidly generating large volumes of radar data without the cost and labor involved in collecting measurement data. This data can be used for training deep neural networks, which are now the algorithm of choice for classifying

radar images [27, 29, 37]. The simulation framework can be integrated with software test beds for rapid prototype development and validation. Finally, the simulation models can aid in understanding radar propagation phenomenology in the environment and pinpointing cause and effect.

2. Incorporate detailed range-Doppler clutter based on Indian road conditions in the radar model to simulate realistic radar images: In real-world conditions, considerable clutter can arise from the rough road surface at mm-wave frequencies, which is proportional to the radar coverage area [73]. Hence, this range-based clutter cross-section increases with the radar range. Further, Doppler-based clutter can arise due to wind [64]. The range and Doppler-based clutter combination manifests as speckle noise and can significantly distort the ISAR images. These simulated images are benchmarked against ISAR images generated with measurement data gathered from Texas Instruments' AWR 1843 77GHz automotive radar
3. Release a curated database of realistic ISAR images to the radar community.
4. Analyze the ISAR images for detailed information about the type of vehicle, its dimensions, the number of wheels, and the trajectory followed by the vehicle, and test the effectiveness of these features for classification of automotive radar targets.
5. Apply explainable artificial intelligence (AI) techniques to enhance the interpretability of radar-based decision-making processes in ADAS based on counterfactual explanations. Specifically, we explore a generative AI

framework that introduces realistic perturbations to ISAR images of a vehicular class, which eventually result in misclassification to a second vehicular class. These perturbations simulate shadowing/occlusion phenomena of specific parts due to environmental conditions.

6. Explore novel applications of ISAR imaging of automotive targets:
  - (a) Automated parking test using ISAR images from automotive radar
  - (b) Investigate and implement around-the-corner sensing capabilities using radar technology to enhance situational awareness and hazard detection in urban environments.

### 1.3 List of Publications

1. **N. Pandey**, D. Mathur, D. Sarkar, and S. S. Ram, “Generating counterfactual explanations for misclassification of automotive radar targets,” *IEEE Transactions on Radar Systems*, 2025.
2. **N. Pandey** and S. S. Ram, “Explainable Artificial Intelligence based Classification of Automotive Radar Targets,” in *2023 IEEE Radar Conference (RadarConf23)*, IEEE, 2023, (pp.1-6).
3. **N. Pandey** and S. S. Ram, “Characterization of millimeter wave clutter on Indian roads at automotive radar frequencies,” in *International Conference on Radar Systems (RADAR 2022)*, IET, 2022, (pp.641-646).
4. **N. Pandey** A. Sneha, and S. S. Ram, “Disaggregation and detection of

- multiple automotive targets for millimeter-wave around-the-corner radar,” in *International Conference on Radar Systems (RADAR 2022)*, IET, 2022, (pp.647-652).
5. **N. Pandey** and S. S. Ram, “Dataset of simulated inverse synthetic aperture radar (ISAR) images of automotive targets,” in *IEEE Dataport (2021)*, IEEE, 2021.
  6. **N. Pandey** and S. S. Ram, “Automated parking test using ISAR images from automotive radar,” in *IEEE International Conference on Autonomous Systems (ICAS)*, IEEE, 2021, pp. 1–5.
  7. **N. Pandey** and S. S. Ram, “Classification of automotive targets using inverse synthetic aperture radar images,” in *IEEE Transactions on Intelligent Vehicles*, IEEE, 2022,(pp. 675-689).
  8. **N. Pandey**, G. Duggal and S. S. Ram, “Database of simulated inverse synthetic aperture radar images for short range automotive radar,” in *2020 IEEE International Radar Conference (RADAR)*., IEEE, 2020, (pp. 238-243).
  9. S. Vishwakarma, **N. Pandey** and S. S. Ram, “Clutter mitigation in range enhanced radar images using sparsity based denoising autoencoders,” *2019 International Radar Conference (RADAR)*, IEEE, 2019, (pp. 1-6).

## 1.4 Notation

*Notation:* Scalar ( $x$ ) and vector variables ( $\mathbf{x}$ ) are denoted with regular and bold font, respectively. Matrices ( $\mathbf{X}$ ) are denoted with bold capital letters. The expectation operator is denoted by  $\mathbb{E}[\cdot]$ .

## Chapter 2

# High Fidelity ISAR Image Simulation and Classification of Automotive Targets

*In this chapter, we developed a comprehensive ISAR simulation framework to generate realistic radar images of automotive targets. The primary objective was to create a detailed simulation environment that incorporated the effects of additive noise in the radar receiver and speckle noise due to ground clutter, providing a rapid and cost-effective alternative to extensive real-world data collection.*

*We began by detailing the simulation model in section 2.1 , starting with the animation of targets along a trajectory and progressing to the advanced signal processing model, where the radar signal processing steps for ISAR image simulation are thoroughly discussed. A significant focus was placed on the impact of real-world clutter on radar imaging, particularly from rough road surfaces at mm-wave frequencies and Doppler-based clutter caused by environmental factors.*

*In section 2.1.3, we modeled static and dynamic clutter using simulation and measurement data from a 77 GHz radar in various automotive scenarios. We then integrated range-Doppler clutter and receiver noise into the radar models, which is crucial for training ML and deep neural networks for radar image classification.*

*In section 2.1.4, we generated a robust database of over 30,000 ISAR images and shared them with the radar research community. These images captured intricate details such as vehicle types, dimensions, wheel counts, and trajectories, highlighting the unique features of ISAR imaging compared to traditional radar data representations.*

*In section 2.2 the simulated images are validated using measurement data from Texas Instrument's AWR 1843 77GHz automotive radar, ensuring the accuracy and applicability of our models. Furthermore, we demonstrate the effectiveness of ISAR features in classifying automotive targets using traditional ML algorithms, such as support vector machine (SVM) and random forest (RF), and deep learning techniques like AlexNet and GoogLeNet through transfer learning in section 2.3.*

*This chapter lays a strong foundation for understanding and leveraging ISAR simulation in automotive radar applications, contributing significantly to radar image processing and target classification.*

## 2.1 Simulation Methodology

In this section, we discuss the simulation methodology for modeling the scattered radar signals from five common automotive targets - bicycle, auto-rickshaw, mid-size car, full-sized car, and truck. Then we describe how these models can be integrated into the radar waveform to obtain ISAR images. Finally, we present the method to incorporate noise and clutter in the images.

### 2.1.1 Automotive Animation Model

We imported freely available three-dimensional (3D) computer aided design (CAD) models<sup>1</sup> of the automotive targets into Blender software. Then, we rendered the metallic parts of the automotive into triangular facets. An accurate rendering of the target capturing the diversity of features on the chassis of the vehicle is realized by using a large number of facets. In our work, the bicycle, and auto-rickshaw are rendered with 3919 and 6949 facets, respectively; the mid-size and full-size cars with 6905 and 19964 facets, respectively, while the truck is rendered with 7206 facets as shown in Fig.2.1.

*Facet-count choice.* We begin with a high-resolution triangulated CAD model and progressively reduce the number of facets. Final counts (bicycle: 3919; autorickshaw: 6949; mid-size car: 6905; full-size car: 19964; truck: 7206) were selected so that the overall outline and key edges are preserved on visual inspection, and the simulated ISAR image for a reference trajectory remains

---

<sup>1</sup><https://free3d.com/3d-models/>

stable under further decimation. High-curvature/specular regions (e.g., edges, wheel arches) are kept denser, whereas large flat body panels are meshed more coarsely. Very small geometric details that are below the ISAR image resolution (i.e., too small to be distinguished) are omitted to avoid unnecessary runtime.

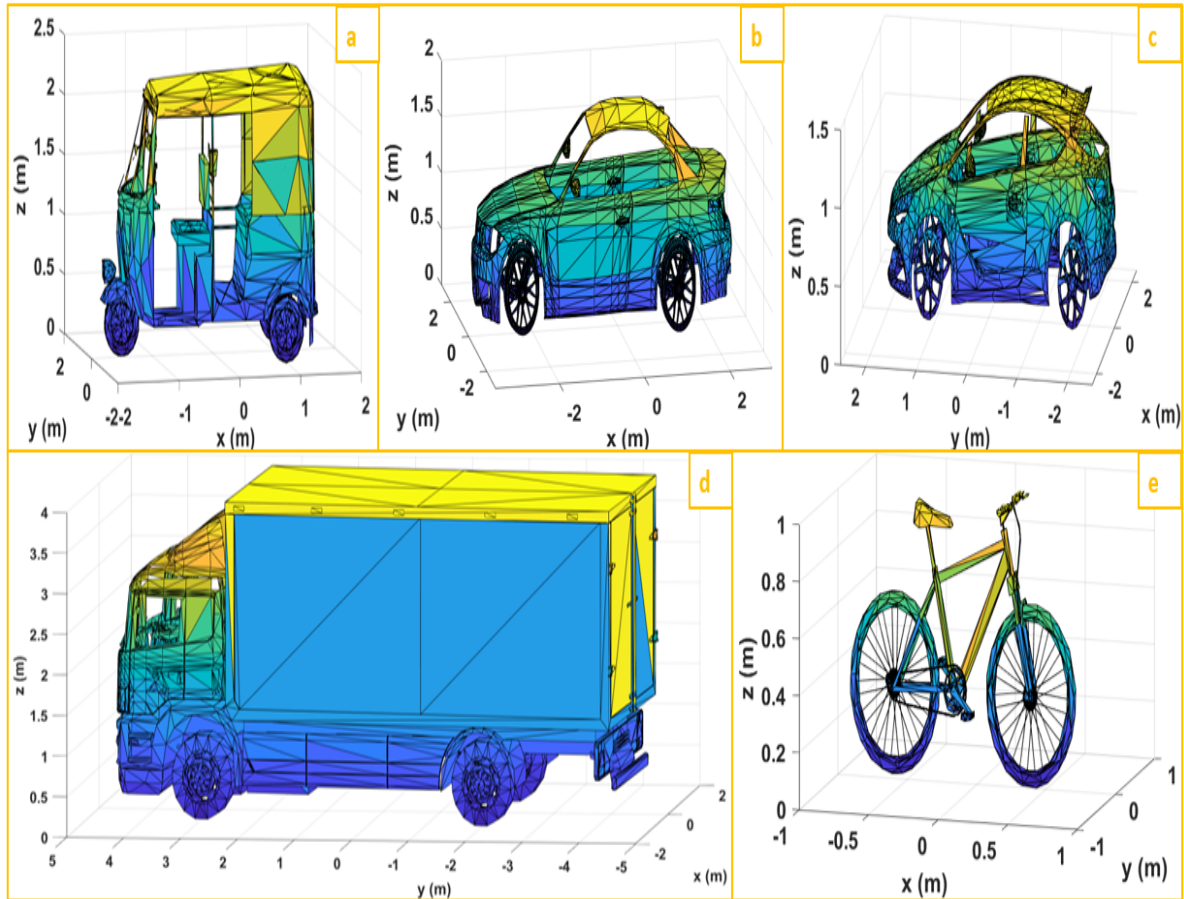


Figure 2.1: Three-dimensional automotive targets with triangular facets used for the simulation. (a) Auto-rickshaw (tuk-tuk), (b) full-sized car (c) mid-sized car, (d) truck (e) bicycle. *Color note:* Colors denote facet/patch groups in the CAD mesh for visualization only; they do not represent amplitude, material, or any physical parameter.

We have considered a four-way traffic junction, where lanes from the north (N), south (S), east (E), and west (W) meet as shown in Fig2.2a.

The targets are assumed to stand on the XY ground plane which is aligned with the N-S and E-W directions with the height along the Z-axis. The ego radar is assumed to be static and fixed at  $(0, 0, 0.5)$ m along the south road. A total of




Figure 2.2: (a) Road geometry of four-way traffic junction. (b) Trajectories undertaken by the automotive target in a four-way junction - (i) Right turn, (ii) Left turn, (iii) U-turn and (iv) Straight through.

16 different trajectories are possible at this junction. They are the four right turns (S to E, E to N, N to W, W to S), four left turns (S to W, W to N, N to E, E to S), four U-turns (S to S, E to E, N to N, W to W) and four straight through (S to N, N to S, W to E, E to W) as shown in Fig.2.2b. Now, we describe the method for animating the vehicle along the desired trajectory at a specified speed. We first identify the center of the vehicle  $\vec{r}_C$  and fix it at the starting position along a trajectory. Then, we identify way points along the distance of the trajectory such that the time taken for the vehicle's center to travel between any two way points is fixed ( $t_f$ ) as shown in Fig.2.3a. The sampling time instants corresponding to these way points are indicated by  $f = 1, 2, \dots, F$ . Therefore, the center of the vehicle undergoes translational displacement  $\vec{\nabla}r[f] = \vec{r}_C[f] - \vec{r}_C[f - 1]$  at each  $f$ . Now, the target is composed of  $B$  triangular facets where the centroid of each facet is  $\vec{r}_b, b = 1 : B$ . Based on the rendering of the vehicle, we obtain displacement vectors ( $\vec{\nabla}r_b$ ) of these centroids from the center of the vehicle.

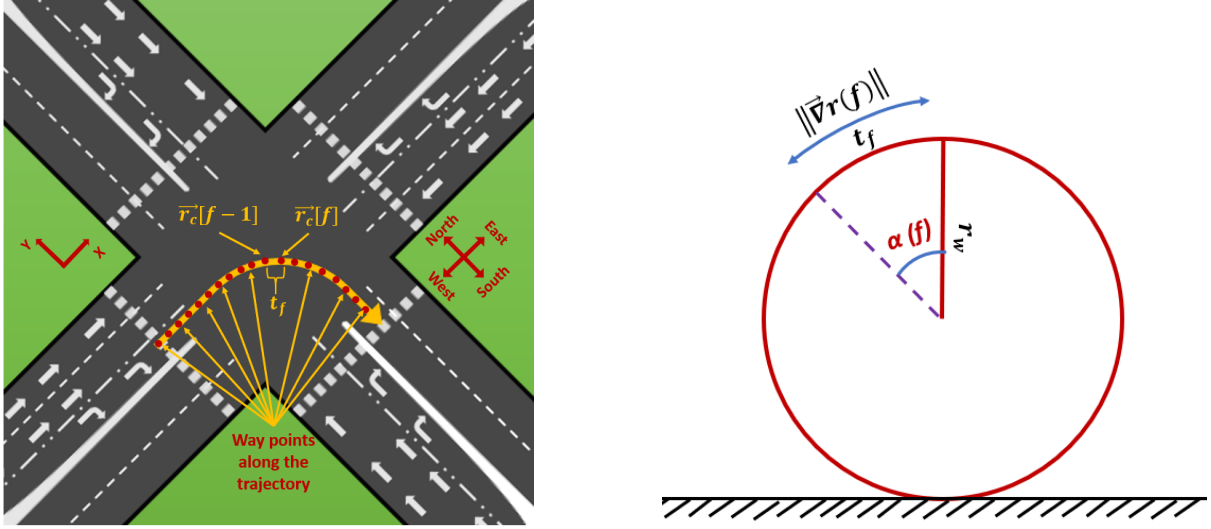


Figure 2.3: Animation model of trajectory: (a) Way points along the trajectory that will be traversed by the center of the vehicle; (b) Wheels rotational angle calculation.

These displacement vectors are fixed with time since the chassis of the vehicle behaves as a rigid body (excepting the wheels). Now, the vehicle must undergo rotational motion along with translational motion. This rotational motion at each frame  $f$  is described in terms of the yaw (rotation angle about the  $Z$  axis,  $\theta_b[f]$ ) which is computed from

$$\theta[f] = \arctan \left( \frac{y_C[f] - y_C[f-1]}{x_C[f] - x_C[f-1]} \right). \quad (2.1)$$

where  $y_C$  and  $x_C$  are  $y$  and  $x$  coordinates of centroid. We restrict attention to planar motion with *yaw-only* rotation about the  $Z$  axis; roll and pitch, including small rotations induced by road unevenness, are not modeled.

The position of the facet centroid on the chassis of the vehicle is

$$\vec{r}_b[f] = (R_{\theta[f]} \vec{r}_b[f-1]) + \vec{\nabla}r[f], \quad (2.2)$$

where  $R_\theta$  is the Euler rotation matrix for a yaw of  $\theta$ . In the case of wheels, a

facet centroid on the wheel undergoes additional rotation due to the motion of the wheel. The angular displacement of the wheel is proportional to the distance travelled by the wheel and the radius of the wheel  $r_w$  as shown in

$$\alpha[f] = \frac{\|\vec{\nabla}r[f]\|}{r_w}. \quad (2.3)$$

The axis of this wheel rotation is obtained by the cross product of the direction of translational displacement and the height axis. The total displacement of a point on the wheel is

$$\vec{r}_b[f] = (R_{\alpha[f]}R_{\theta[f]}\vec{r}_b[f-1]) + \vec{\nabla}r[f], \quad (2.4)$$

where  $R_{\alpha}$  is the Euler rotation matrix corresponding to a pitch angle of  $\alpha$ . The entire algorithm describing the animation motion modeling is summarized in [Algorithm.1](#).

### 2.1.2 Advanced Point Scatterer Signal Model

Automotive radars use linear frequency modulated (LFM) waveforms for the detection and tracking of targets. We consider a radar transmitting an LFM waveform,

$$s^{tx}(\tau) = \text{rect}\left(\frac{\tau}{T_{PRI}}\right) e^{j2\pi f_c \tau} e^{j\pi \gamma \tau^2}, \quad (2.5)$$

with  $f_c$  carrier frequency (and wavelength  $\lambda_c$ ) and a chirp rate  $\gamma$ . In the above expression,  $\text{rect}(\cdot)$  indicates that the transmitting signal is defined for a pulse

---

**Algorithm 1:** Animation model of vehicle along desired trajectory

---

**Data:** Fixed displacement vectors corresponding to facet centroids on the chassis and wheels of the vehicle ( $\vec{\nabla}r_b, b = 1 : B$ )

1 with respect to the center of the vehicle ( $\vec{r}_C$ ). **Data:** Specify way point positions for center of vehicle along trajectory:  $\vec{r}_C[f], f = 1 : F$ . Time duration between two way points is fixed ( $t_f$ ).

**Result:** Time-varying position coordinates of facet centroids on vehicle  
( $\vec{r}_b[f], b = 1 : B, f = 1 : F$ )

2 Initialization: Initialize positions of all the facet centroids  $\vec{r}_b[1] = \vec{r}_C[1] + \vec{\nabla}r_b, b = 1 : B$ ;

3 **for**  $f=2:F$  **do**

4     Compute displacement between two consecutive way points  $\vec{\nabla}r[f] = \vec{r}_C[f] - \vec{r}_C[f - 1]$ ;

5     Compute vehicle yaw rotation:  $\theta[f] = \frac{y_C[f]-y_C[f-1]}{x_C[f]-x_C[f-1]}$ ;

6     **if** *Facet centroids on chassis* **then**

7         Perform Euler rotation on facet centroids  $\vec{r}_b[f] = (R_{\theta[f]}\vec{r}_b[f - 1]) + \vec{\nabla}r[f], b = 1 : B$ ;

8     **else if** *Facet centroids on wheels* **then**

9         Compute rotation of wheel  $\alpha[f] = \frac{\|\vec{\nabla}r[f]\|}{r_w}$  where  $r_w$  is the radius of the wheel.;

10         Compute axis of wheel rotation which is perpendicular to the plane defined by height axis and vehicle translational motion direction.;

11         Perform Euler rotation on facet centroids on wheel based on wheel rotation axis  
            $\vec{r}_b[f] = (R_{\alpha[f]}R_{\theta[f]}\vec{r}_b[f - 1]) + \vec{\nabla}r[f], b = 1 : B$

12     **end**

13     **else**

14     **end**

15 **end**

---

repetition interval(PRI) of  $T_{PRI}$ . We model the automotive target as a collection of moving scattering centers,  $b = 1 : B$ , located at the centroids of each of the facets of the vehicle. We only model the reflections from the metallic surfaces of the automotive targets since the non-metallic surfaces are known to be poor reflectors [74, 75]. We have strictly considered scenarios where there are single targets in the images. This is usually ensured in real-world conditions through pre-processing of raw radar data before classification. In real-world deployments, single-channel source separation and temporal association *can* isolate returns from multiple extended targets *when* they are sufficiently separated in range/angle/Doppler and maintain stable tracks over time; this is not guaranteed in dense traffic [40, 76, 77]. Since these algorithms are well known to the research community, they are not included in the manuscript due to space constraints. The time-varying range for each scattering center is  $r_b = R_b + v_b t$ , where  $R_b$  is the starting distance from the radar. The radar signal scattered back from a single point target is Doppler shifted by  $f_{D_b} = \frac{2v_b}{\lambda_c}$  due to the target's relative radial velocity ( $v_b$ ) with respect to the radar. The received radar signal, after downconversion to the baseband, is written in terms of slow time  $t$  and fast time,  $\tau$  as

$$s_b^{rx}(\tau, t) = a_b(t) \text{rect} \left( \frac{\tau - \tau_b}{T_{PRI}} \right) e^{-j \frac{4\pi f_c}{c} R_b} e^{-j 2\pi f_{D_b} t} e^{j \pi \gamma (\tau - \tau_b)^2} + \nu, \quad (2.6)$$

where  $\tau_b(t) = \frac{2r_b(t)}{c}$  is the time delay to the target. In the equation above,  $\nu$  denotes the additive noise that will be discussed in greater detail in the following section. The strength of the received returns, denoted by  $a_b$ , is obtained through

the radar range equation by incorporating the transmitted power ( $P^{tx}$ ), the gains of the transmitting ( $G^{tx}$ ), and receiving radar antennas ( $G^{rx}$ ), the radar cross-section of each scatterer center ( $\sigma_b$ ) and the distance of the point scatterer from the radar, as shown in

$$a_b^2 = \frac{P^{tx} G^{tx} G^{rx} \sigma_b \lambda_c^2}{(4\pi)^3 r_b^4}. \quad (2.7)$$

In (2.6) and (2.7), we have assumed stationary channel conditions and direct path target returns without any type of multipath. The Radar Cross Section (RCS) of a flat metallic triangular plate is a function of the radar aspect angle ( $\theta_b$ ), the plate area ( $A_b$ ), and long dimension ( $d_b$ ), [78], as shown in Fig. 2.4, and (2.8)

$$\sigma_b = \eta \frac{4\pi A_b^2 \cos^2 \theta_b \sin^4 \left( \frac{2\pi}{\lambda_c} d_b \sin \theta_b \right)}{\lambda_c^2 \left( \frac{2\pi}{\lambda_c} d_b \sin \theta_b \right)^4}. \quad (2.8)$$

The aspect angle is computed from the dot product of the incident vector from

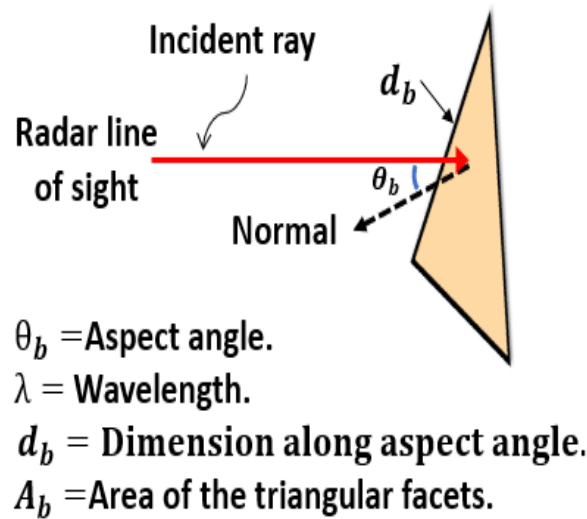


Figure 2.4: Radar cross section of flat metallic triangular plate

the radar to the plate and the normal vector of the plate. Since  $\theta_b$  changes along the target trajectory,  $\sigma_b$  fluctuates. Not all scattering centers may be visible to the radar. One method for incorporating shadowing effects is to implement a shadowing algorithm on the radar data. However, since these algorithms utilize ray tracing methods, they are computationally expensive, especially for spatially large dynamic objects such as automobiles at mm-wave frequencies. For example, a ray-tracing algorithm for a 4mm wavelength would require the rays to be spaced one-tenth wavelength apart and emanate from an aperture that is at least the size of the largest automotive target that we consider. In our case, that would correspond to a truck of  $8.5 \times 2.6m \times 5m$ . Further, these computations would have to be repeated at the radar frame rate of several gigahertz as opposed to the video frame rates of computer animation algorithms typically used in vision. Instead, we opt for a low complexity algorithm by introducing Bernoulli's random variable,  $\eta$ , to model a 20% visibility of each point scatterer [60]. The detailed comparison (accuracy and computational complexity) of the ray tracing and probabilistic method is presented in the section 2.1.5.

The radar data is sampled at a frequency of  $F_s = 1/\delta t$ , and the fast time samples are numbered from  $1 : N$ . Similarly, if we consider a  $p^{th}$  coherent processing interval (CPI) consisting of  $M$  PRIs, then the discrete representation

of (2.6) is

$$S_b^{rx}[n, m] = a_b[m] \text{rect} \left[ \frac{n - n_b}{N} \right] e^{-j\frac{4\pi fc}{c} R_b} e^{-j2\pi m f_{D_b} T_{PRI}} e^{j\pi\gamma\delta t^2 (n-n_b)^2} + \nu, \quad (2.9)$$

where  $n_b$  is the integer rounded from  $\frac{\tau_b(t)}{\delta t}$ .

We process the received signal using *stretch processing*, a variation of matched filtering, which is especially suited for low sampling frequencies [79]. The maximum unambiguous range of the radar,  $R_{max}$ , is equal to  $\frac{cT_{PRI}}{2}$ . For every CPI, we consider a radar range span of interest,  $R_0 - \frac{R_{span}}{2} : R_0 + \frac{R_{span}}{2}$  within  $R_{max}$  where  $R_0$  is called the central reference position (CRP). The time delay to the  $b^{th}$  point scatterer can be expressed as  $\tau_b = \tau_0 + \delta\tau_b(t)$  where  $\tau_0 = \frac{2R_0}{c}$  corresponds to the time delay to the CRP. Since, the target motion is known, the CRP is chosen to correspond to the mean range to  $\vec{r}_C$  in every CPI. In stretch processing, the received signal is multiplied with  $e^{-j\pi\gamma\delta t^2 (n-n_0)^2}$ , where  $n_0$  is the integer rounded from  $\frac{\tau_0}{\delta t}$  over every PRI. Thus, we obtain

$$S_b^{rx}[n, m] = a_b[m] e^{-j\frac{4\pi fc}{c} R_b} e^{-\pi\gamma\delta t^2 (n_0-n_b)^2} e^{-j2\pi m f_{D_b} T_{PRI}} e^{-j4\pi\gamma\delta t^2 n(n_0-n_b)}. \quad (2.10)$$

We carefully compensate for the translational motion of the vehicle, and only consider the rotational motion of the point scatterer within a CPI. Then, the first two exponential terms in (2.10) are constant phase terms and are absorbed into the amplitude during further processing. Since translation motion is compensated and the rotational motion within a CPI is modelled as constant velocity, the

Doppler  $f_{D_b}$  is constant over slow time; consequently, the slow-time factor  $e^{-j2\pi m f_{D_b} T_{PRI}}$  is *linear* in  $m$ . Significant acceleration would instead manifest as slow-time phase curvature, which is negligible for the trajectories considered.

The last two terms show the variation of the two-dimensional (2D) signal over slow and fast times, as shown in

$$S_b^{rx}[m, n] = a_b(\cdot) e^{-j2\pi f_{D_b} m T_{PRI}} e^{-j4\pi\gamma\delta t^2 n(n_0 - n_b)}. \quad (2.11)$$

The fast time sampling frequency ( $F_s = 1/\delta t$ ) is obtained from twice the stretch bandwidth which is  $\frac{2R_{span}\gamma}{c}$  where  $R_{span}$  is much lower than  $R_{max}$ . Note that the stretch bandwidth is much lower than the radar bandwidth ( $\gamma T_{PRI} = \frac{2R_{max}\gamma}{c}$ ). Hence, the sampling frequency requirements for stretch processing results in lower than that of the ordinary matched filtering.

*Ambiguity limits.* For a linear FMCW chirp with slope  $\gamma$  and IF/ADC constraints  $\{B_{IF}, F_s\}$ , the maximum unambiguous beat frequency is  $f_{b,max} = \min\{B_{IF}, F_s/2\}$ , giving the unambiguous range

$$R_u = \frac{c f_{b,max}}{2\gamma}. \quad (2.12)$$

Here,  $B_{IF}$  denotes the analog IF bandwidth and  $F_s$  is the ADC sampling frequency. Slow-time Doppler is sampled at a rate  $1/T_{PRI}$ ; all simulation scenarios here are designed to remain within the corresponding unambiguous Doppler band set by  $T_{PRI}$ .

When the target is an extended target with multiple point scatterers ( $B$ ), then the received signal is obtained by the sum of the returns from each scatterer.

$$S^{rx}[m, n] = \sum_{b=1}^B S_b^{rx}[m, n]. \quad (2.13)$$

Here, we have ignored the multiple scattering between the different parts of the target. The output is processed using 2D Fourier transform to obtain range-Doppler ambiguity plots,

$$\chi[f_D, r] = \mathcal{DFT}_{2D}\{S^{rx}[m, n]\}, \quad (2.14)$$

where the range dimension  $r$  spans  $N$  steps from  $R_0 - \frac{R_{span}}{2}$  to  $R_0 + \frac{R_{span}}{2}$ ; and  $f_D$  spans  $M$  steps from  $-\frac{1}{2T_{PRI}}$  to  $\frac{1}{2T_{PRI}}$ . The 2D plot can also be interpreted as a range-cross plot ( $\chi[r, cr]$ ) provided an accurate estimate of the angular velocity ( $\omega$ ) of the target is available, since translational motion has been compensated. We estimate  $\omega$  for every  $p^{th}$  CPI by the change in yaw ( $\Theta$ ) of the vehicle as shown in

$$\omega = \frac{\Theta[p] - \Theta[p - 1]}{T_{CPI}}. \quad (2.15)$$

Then the Doppler axis is converted to the cross-range axis by

$$cr[m] = f_D[m] \times \frac{\lambda_c}{2\omega}, \text{ for } m = 0 : M - 1. \quad (2.16)$$

Depending on  $\omega$ , the cross-range spans across images may vary even when the pixel dimensions of the plot remain unchanged.

*Unambiguous Doppler for cross-range.* Since  $f_D$  lies within the unambiguous slow-time band set by  $T_{\text{PRI}}$  for our settings, all ISAR images reported here fall within the unambiguous Doppler region.

*Scope note.* In this work, frequency-dependent diffraction and environment-induced multipath are not explicitly modelled; instead, their effect on target visibility is approximated using a calibrated phenomenological visibility process. Extending the simulator to incorporate full high-frequency electromagnetic models (e.g., ray-based methods including diffraction) is left as future work.

### 2.1.3 Noise and Range-Doppler Clutter Models

In this section, we discuss how we incorporated ground-based clutter along the range and wind-based clutter along the Doppler dimensions along with additive noise in the time-domain data.

*Ground clutter:* For a rough surface, the clutter cross-section is proportional to the surface clutter coefficient,  $\sigma^0$ , and the radar coverage area. A stable component - due to static road conditions such as road material - and a fluctuating component due to wind contribute to  $\sigma^0$  [64, 80]:

$$\sigma^0 = \sigma_{stable}^0 + \sigma_{fluctuating}^0. \quad (2.17)$$

We model  $\sigma^0$  as an exponential random variable with a mean of -23.6dB which corresponds to asphalt at mm-wave frequencies in chapter 3, and [64, 81]. *Mapping from measured statistics to simulator.* Section 3.3 shows that  $\sigma_0$  across

*sites/conditions* is multi-modal, well captured by GMM; here, for a single scene, we adopt the standard exponential (power) speckle model with a *scenario-specific mean*  $\sigma_0$  selected from the measured set for the chosen road type, together with geometry-based range scaling. The radar coverage area is proportional to the antenna beamwidth ( $\theta_{BW}$ ), grazing angle ( $\psi$ ), radar range resolution ( $\delta r = \frac{c}{2\gamma T_{PRI}}$ ) and range. Therefore,  $\sigma_c$  is

$$\sigma_c = \sigma^0 r \theta_{BW} \delta r \sec \psi. \quad (2.18)$$

This results in a ground clutter that is a function of range as shown in

$$C_0[r] = \frac{P^{tx} G^{tx} G^{rx} \lambda_c^2 \sigma^0 \theta_{BW} \delta r \sec \psi}{(4\pi)^2 r^3}. \quad (2.19)$$

Here  $C_0[r]$  denotes the received clutter power for range bin  $r$  in watts (W).

We have maintained the radar at the height of 0.5m above the ground from which the grazing angle can be computed for every range  $r$ .

*Doppler clutter:* Based on [64], the power spectrum of the Doppler clutter can be modeled as a low pass filter response. When combined with the range-related clutter, we obtain

$$C[f_D, r] = C_0[r] \left[ 1 + \left( \frac{f_D}{\Delta f_D} \right)^s \right]^{-1}. \quad (2.20)$$

where  $s$  is a function of the average wind velocity  $U$  (m/s) as shown in

$$s = \frac{2(U + 2)}{(U + 1)} \left( \frac{100}{2\pi f_c} \right)^{0.2}. \quad (2.21)$$

In (2.20),  $\Delta f_D$  is the  $-3dB$  width of the spectrum and is given by

$$\Delta f_D = 1.23 \left( \frac{3.2}{\lambda_c} \right) U^{1.3}. \quad (2.22)$$

Based on local meteorological reports [82],  $U$  can vary from 0m/s to 10m/s. We consider four possible wind speeds (2.5, 5, 7.5, 10m/s) in our simulations. Finally, we convert the power values obtained from (2.20) to voltage values for each range-Doppler pixel ( $c[f_D, r]$ ). We multiply the voltage with a phase modeled as a complex circularly symmetric random variable ( $\phi[f_D, r]$ ). This complex clutter signal ( $c[f_D, r] \exp(j\phi[f_D, r])$ ) is then added to each pixel of the range-Doppler ISAR images  $\chi[f_D, r]$  in (2.14). In this work, we use radar parameters that closely correspond to an actual automotive radar and list them in Table.2.1.

*Choice of parameters and practical constraints.* The parameter values in Table 2.1 are chosen to satisfy practical constraints of the AWR-1843 measurement chain (IF/ADC bandwidth and sampling), and to keep the radar configuration consistent across all simulation and measurement results reported in this thesis. The carrier frequency  $f_c = 77$  GHz follows the standard automotive mm-wave band. The sampling frequency  $F_s$  and the stretch/IF bandwidth ( $1/T_{SBW}$ ) are selected such that the de-chirped (stretch-processed) beat signal remains within the usable IF/ADC band for the ranges of interest, thereby avoiding aliasing in the range FFT.

Although the AWR-1843 platform can support a larger RF sweep bandwidth

(up to  $\sim 4$  GHz), this work uses an effective sweep corresponding to the reported range resolution of 0.075 m in Table 2.1 (i.e.,  $B \approx c/(2 \times 0.075) \approx 2$  GHz). While using  $B \approx 4$  GHz would improve range resolution, it would also tighten the practical ambiguity limits under the configured  $\{B_{\text{IF}}, F_s\}$  constraints. In particular, for fixed chirp duration  $T_{\text{PRI}}$ , increasing the sweep bandwidth requires increasing the chirp slope  $\gamma$ , which increases the beat frequency at a given range and reduces the unambiguous range  $R_u$  (as quantified in the ambiguity-limit expression (2.12)). Conversely, maintaining the same  $\gamma$  while increasing the sweep bandwidth would require a larger  $T_{\text{PRI}}$ , which impacts slow-time sampling ( $1/T_{\text{PRI}}$ ) and Doppler processing for a fixed  $T_{\text{CPI}}$ . Therefore, the chosen configuration (Table 2.1) provides a practical balance: sufficient range resolution for vehicle-scale scattering structures while maintaining unambiguous range compatibility with the measurement sampling/IF constraints and ensuring consistent parameterisation across the thesis.

Table 2.1: Automotive radar parameters for generating ISAR images

Parameters	Values
Carrier frequency ( $2\pi f_c$ )	77GHz
stretch Bandwidth ( $\frac{1}{T_{\text{SBW}}}$ )	8MHz
Sampling Frequency ( $F_s$ )	5MHz
Chirp rate ( $\gamma$ )	$60 \times 10^{12} \text{ Hz}^2$
Chirp duration ( $T_{\text{PRI}}$ )	$83.33 \mu\text{s}$
Coherent processing interval ( $T_{\text{CPI}}$ )	0.1s
Doppler resolution	10 Hz
Range resolution	0.075 m
Minimum cross-range resolution	0.19 m
Transmitted power ( $P_t$ )	25dBm
Transmitted power (TI AWR-1843)	14dBm

*Noise:* While clutter was modeled as a speckle noise in the radar images based on the above description, we modeled noise as an additive white Gaussian

$\mathcal{N}(0, N_p)$  in the time-domain radar returns in (2.6). Based on our radar range equation in (2.7), the minimum received signal at the radar is  $-80$  dBm. Based on the ratio between the minimum received signal at the radar and the mean noise floor, we considered four different signal to noise ratio (SNR) scenarios from  $-5$  dB to  $+10$  dB in our simulations.

*SNR definition (time-domain returns).* Throughout this chapter, SNR refers to the signal SNR at the stage of (2.6) (post-stretch, pre-2D processing).

$$S^{rx}[n, m] = S_{\text{sig}}^{rx}[n, m] + \nu[n, m], \quad \mathbb{E}[\nu[n, m]] = 0, \quad (2.23)$$

$$\text{SNR} \triangleq \frac{\mathbb{E}[|S_{\text{sig}}^{rx}[n, m]|^2]}{\mathbb{E}[|\nu[n, m]|^2]}, \quad (2.24)$$

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \left( \frac{\mathbb{E}[|S_{\text{sig}}^{rx}[n, m]|^2]}{\mathbb{E}[|\nu[n, m]|^2]} \right). \quad (2.25)$$

The scenarios “SNR =  $\{-5, 0, 5, 10\}$  dB” are realized by selecting  $\mathbb{E}[|\nu[n, m]|^2]$  from the ratio between the minimum received signal level ( $-80$  dBm) and the mean noise floor at this stage.

### 2.1.4 Database of Simulated Data

We present examples of ISAR images of each of the targets below Figs 2.5, 2.6, 2.7, 2.8, 2.9, 2.10, 2.11, 2.12, 2.13, and 2.14. For all five automotive targets (bicycle, auto-rickshaw, mid-size car, full-size car, and truck), we present two sets of results. The first set, shown in Figs. 2.5, 2.7, 2.9, 2.11, and 2.13, contains ISAR images corrupted only by additive noise in the receiver data due to receiver electronics; we present the set of images corresponding to an SNR of +10 dB. The second set, shown in Figs. 2.6, 2.8, 2.10, 2.12, and 2.14, contains ISAR images corrupted by range–Doppler clutter that gives rise to speckle noise. In this scenario, the clutter strength is a function of the surface reflection coefficient of the road and the wind speed; we present figures corresponding to a wind speed of 2.5 m/s. In all of the figures, each row corresponds to images simulated for a distinct trajectory. The top row is obtained when the vehicle moves along an almost straight trajectory from N to S; the second row shows the trajectory of a target taking a left turn from E to S; the third row shows the trajectory of a target taking a right turn from S to E; while the fourth row shows the trajectory of a target doing a U-turn from W to W. The range span (Y-axis) in all the figures is 20 m and is centered along the CRP corresponding to the specific CPI. The cross-range span (X-axis) in all of the figures may vary from 10 m to 20 m and is centered at 0 m. The cross-range axis correlates to the Doppler axis of  $\chi[f_D, r]$ . The noisy images (SNR +10 dB) are plotted with a dynamic range of 50 dB from –40 dBm to –90 dBm, while the cluttered images are plotted with a dynamic

range of 80 dB from  $-40$  dBm to  $-120$  dBm.

We first present the ISAR images of a *bicycle* in Figs.2.5, and 2.6. As mentioned earlier, the images in Fig.2.5 correspond to the ISAR images corrupted by noise, while the figures in 2.6 correspond to the ISAR images corrupted by range-Doppler based clutter.




Figure 2.5: ISAR images of a *bicycle* with +10 dB SNR at CPI frames corresponding to 1.5, 2.5, 3.5 and 4.5 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from west to south and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB ( $-40$  dBm to  $-90$  dBm).

The bicycle is a spatially narrow target and hence appears as a cluster of very closely spaced scattering centers, almost like a single-point scatterer. The dimensions of the target can be estimated from some of the images (for example, sub-figure (vi) in Fig. 2.5). The noisy images in Fig. 2.5 show that at some

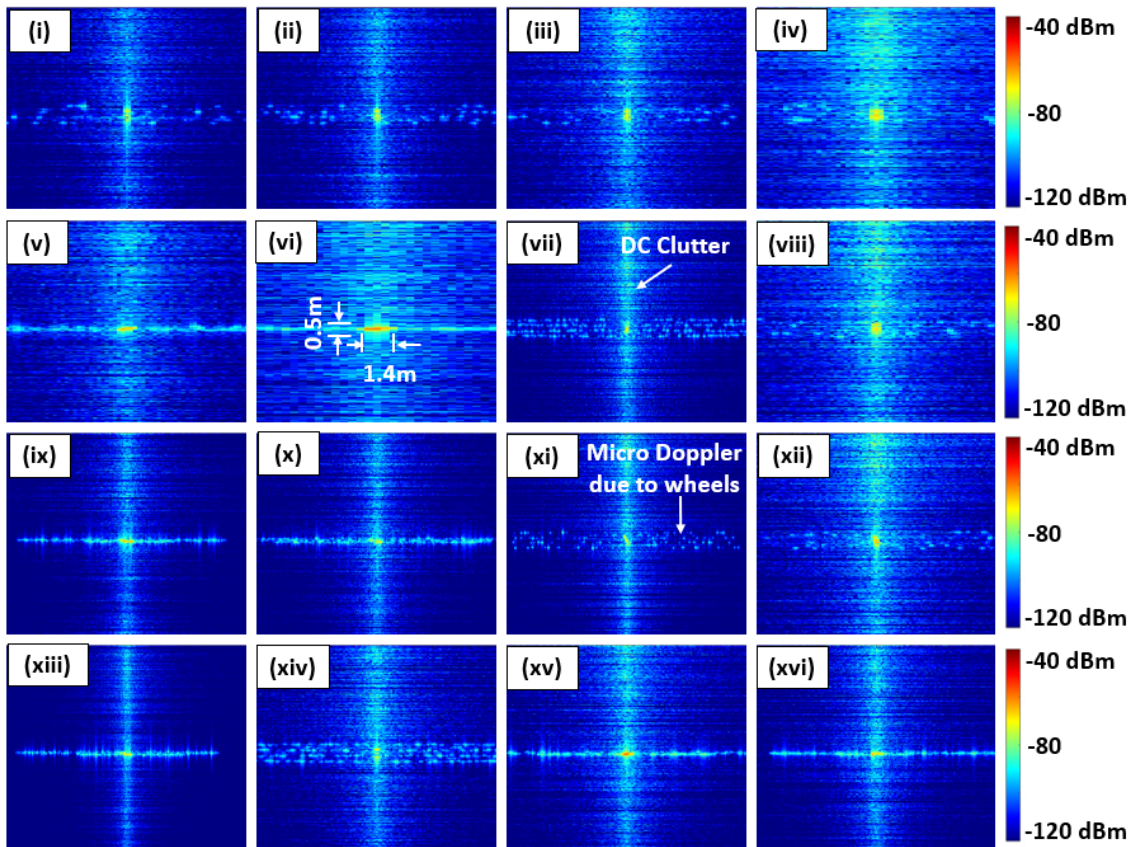


Figure 2.6: ISAR images of a *bicycle* with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1.5, 2.5, 3.5 and 4.5 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from west to south and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).

range positions, the target becomes difficult to discern due to low returns from the bicycle, especially when it is far from the radar. This is because of the low RCS of the bicycle. The cluttered images in Fig. 2.6 show strong clutter at DC (corresponding to 0m along the cross-range). The width of the Doppler spectrum and the strength of the clutter returns change depending on the wind speed. The bicycle can still be discerned in some of the images along with the micro-Doppler tracks due to its wheels (sub-figures (xi) and (xiv) in Fig. 2.6). Note that our model does not consider the motion of the bicycle handlebars or rider's body, and hence those effects cannot be observed in the images.

Next, we present the ISAR images of an *auto-rickshaw* in Figs.2.7, and 2.8.




Figure 2.7: ISAR images of an *auto-rickshaw* with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm).

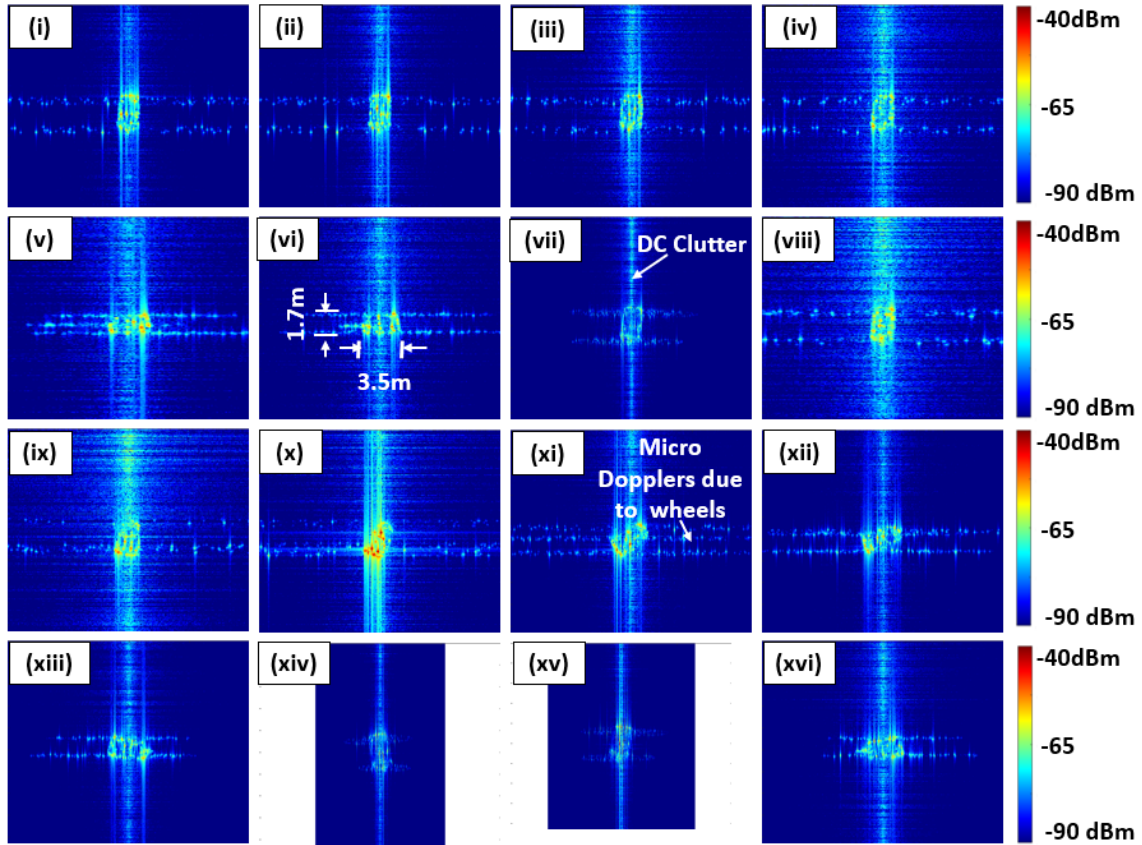


Figure 2.8: ISAR images of an *auto-rickshaw* with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).

The images show that the auto-rickshaw is a spatially larger target than the bicycle. The shape of the vehicle is triangular in the top-view. In fact, in some of the top-view images (sub-figure (vi) in Fig. 2.7), we can clearly see the dimensions of the vehicle. We also see considerable distortions along the Doppler (cross-range) dimension due to the micro-Doppler from the rotation of the wheels. Interestingly, in some images, we can see three distinct micro-Doppler tracks from the three wheels (sub-figure (xi) in Fig. 2.8). On the top row, we observe that the longer dimension of the target’s top view is oriented along the range dimension when the car moves from N to S. In the second row,

the target is first oriented laterally and then turns length wise. This is because the target did a left turn from E to S. Similarly, in the third row, the target was first oriented along the long direction and then along the lateral direction as the target moved from S to E. Finally, in the last row, the target is always along the lateral direction since it does a U-turn from W to W. Therefore, the ISAR images offer some indication of the type of trajectory undertaken by a target.

Figure.?? presents the results of the mid-size car.

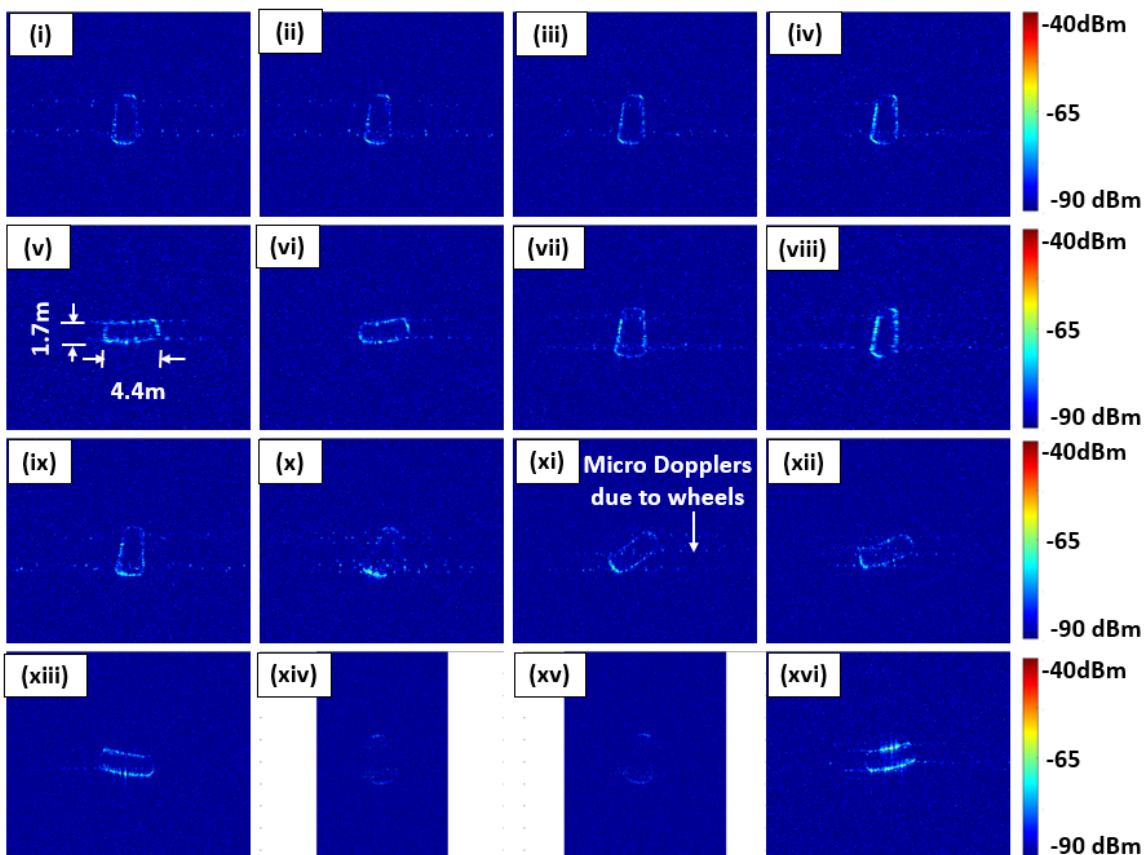


Figure 2.9: ISAR images of a *mid-size car* with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm).

Since this is a larger target than the auto-rickshaw, the returns are stronger. We are able to see the top view of the target with all four sides. Again, we

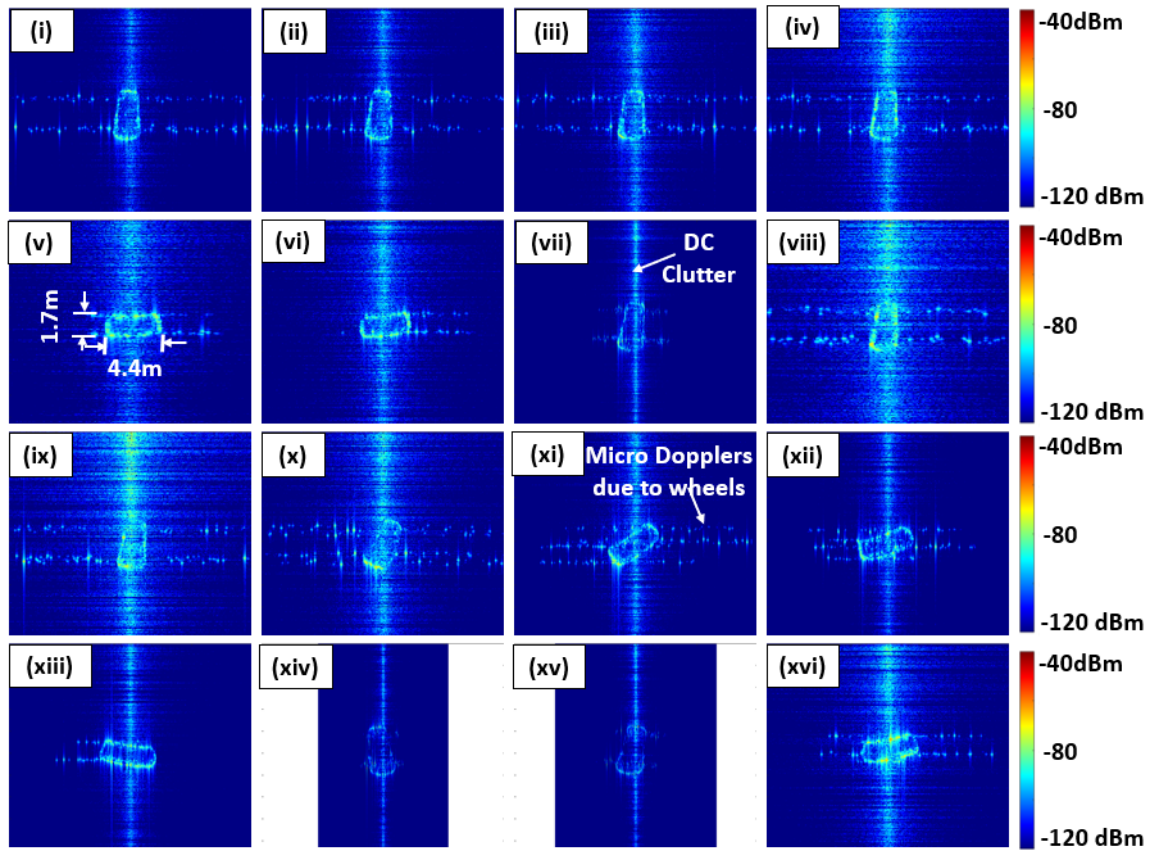


Figure 2.10: ISAR images of a *mid-size car* with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).

observe some micro-Doppler based distortions along the cross-range due to the micro-Doppler returns from the four wheels. Four distinct micro-Doppler tracks are observed in the sub-figure (x) and (xi) in Fig. 2.10. Again, we observe the longer dimension of the car oriented along the range dimension when it is moving from either N to S or S to N. But the longer dimension of the car is oriented along the lateral dimension when the car is moving from E to W or vice versa.

The results in Figs. 2.11 and 2.12, corresponding to the *full-size car*, look similar to those from the mid-size car in Figs. 2.9 and 2.10, except for the larger dimensions of the car in the top-view. The dimensions of the full-size car are  $5.7\text{m} \times 2.4\text{m}$  whereas those of the mid-size car were  $4.4\text{m} \times 1.7\text{m}$ .

The larger target also has stronger returns and is thus easily discerned in the images. Again, we are able to observe four distinct micro-Doppler tracks from the four wheels in some of the images (sub-figures (ix)–(xii) in Fig. 2.12). Also, we are able to see the changes in the orientation of the images as the car undergoes turns along its trajectory.

The largest automotive target that we have considered is the *four wheel truck*, for which the results are presented in Figs. 2.13 and 2.14.

Due to its large size, the top-view obtained from the ISAR images clearly presents the dimensions of the target, which are  $8.5\text{m} \times 2.6\text{m}$ . We are also able to observe the changes in the target orientation along the four distinct trajectories. The micro-Doppler distortions are considerably greater in this case due to the

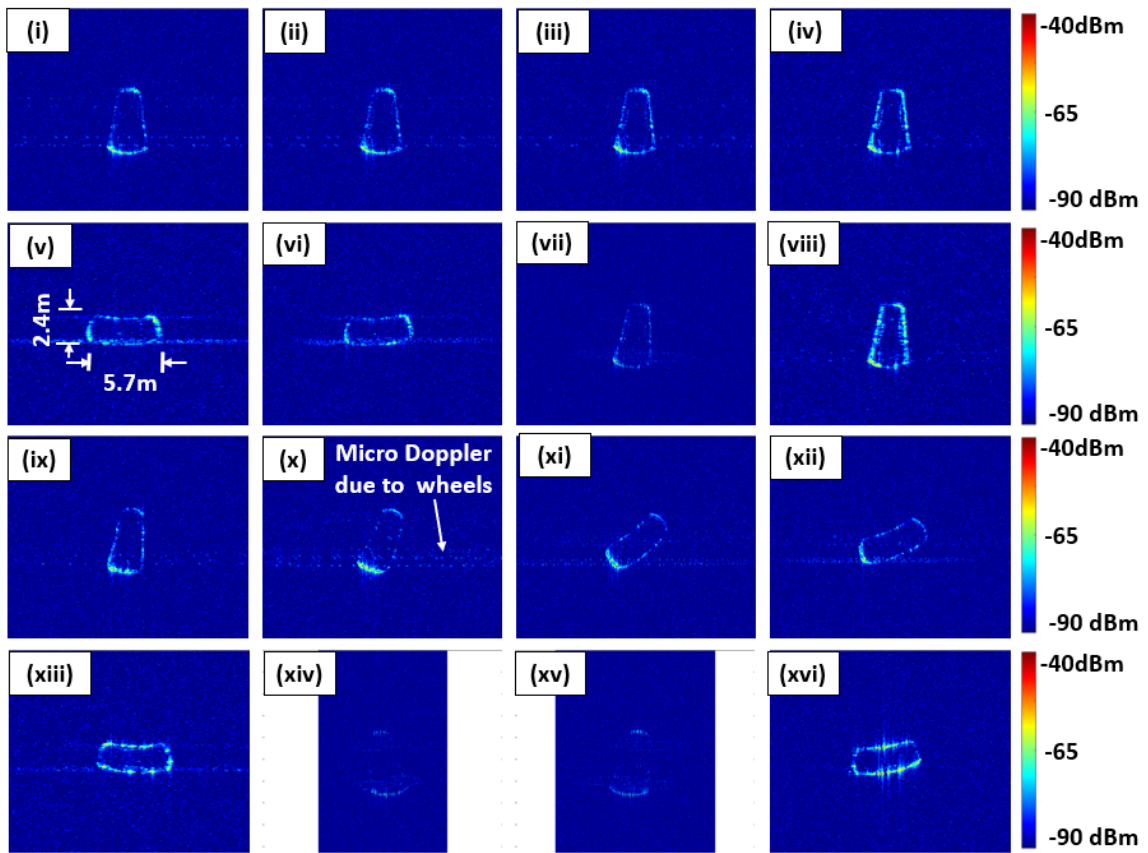


Figure 2.11: ISAR images of a *full-size car* with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm).

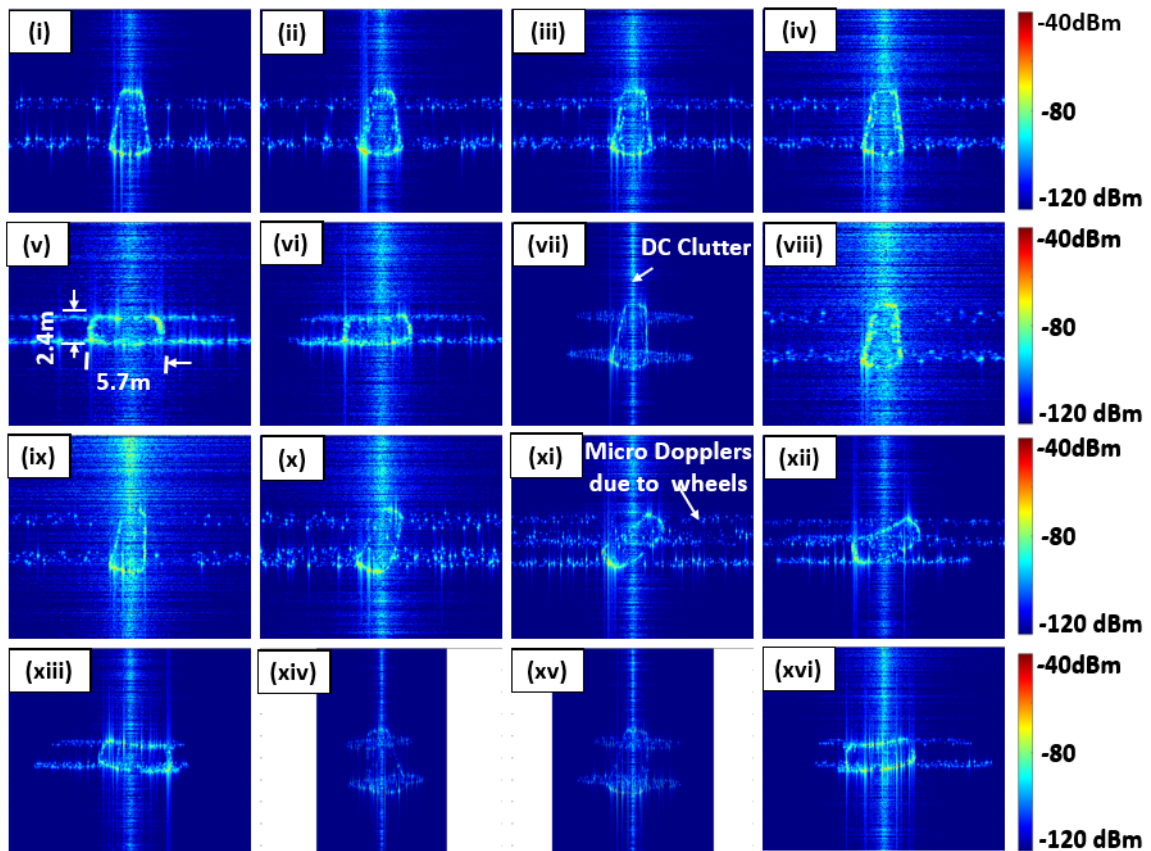


Figure 2.12: ISAR images of a *full-size car* with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).

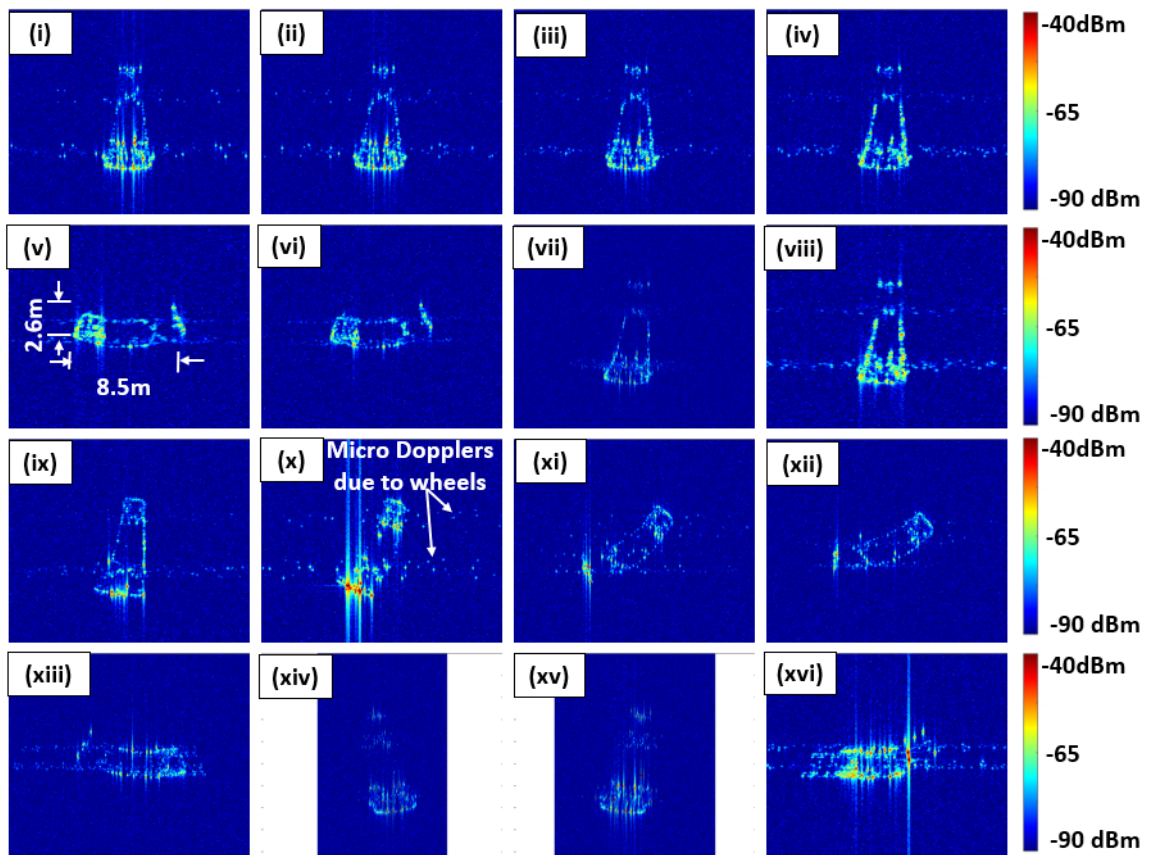


Figure 2.13: ISAR images of a *truck* with +10 dB SNR at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 50 dB (–40 dBm to –90 dBm).

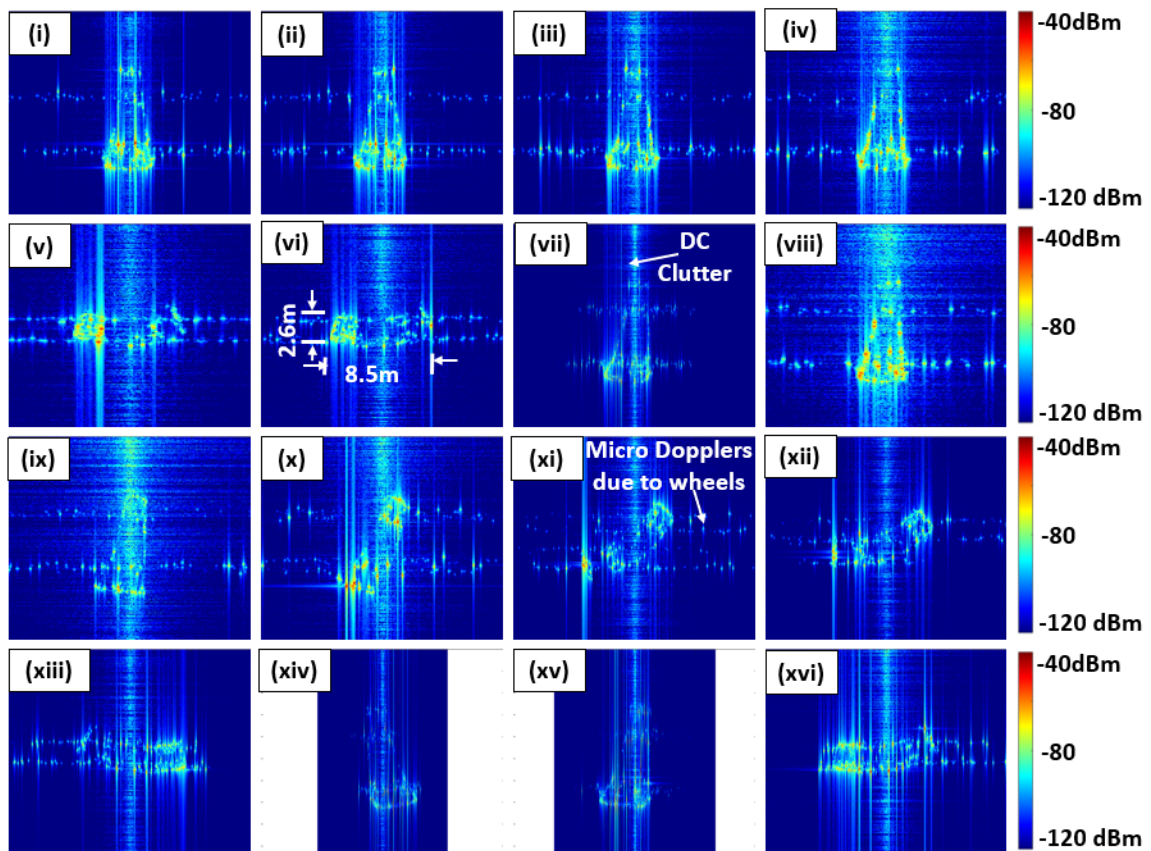


Figure 2.14: ISAR images of a *truck* with range-Doppler clutter for a wind velocity of 2.5 m/s at CPI frames corresponding to 1, 2, 3 and 4 s along the following trajectories: (i–iv) straight path from north to south, (v–viii) left turn from east to south, (ix–xii) right turn from south to east and (xiii–xvi) U-turn from west to west. The range span is 20 m while the cross-range span varies from 10 m to 20 m. The dynamic range is 80 dB (–40 dBm to –120 dBm).

large wheel radii, and four micro-Doppler tracks in sub-figures xi in Fig.2.13, and 2.14 are well resolved in this case.

We list the complete set of simulated ISAR images in Table.2.2. To sum-

Table 2.2: Simulated ISAR image database

Type of images	Types of Targets (#)	Trajectories (#)	Images per trajectory	Total images (#)
Ideal Images	5	16	45-49	3750
Noisy images of SNR (+10,+5,0,-5 dB)	5	16	45-49	14976
Cluttered Images with wind velocities (2.5,5,7.5,10 m/s)	5	16	45-49	14976

marize, we have considered five automotive targets - full-size car, mid-size car, truck, auto-rickshaw, and bicycle. Each target undergoes 16 trajectories, and each trajectory is of 5 seconds duration. Since each CPI is 0.1 seconds, we obtain between 45 and 49 images from each trajectory. The resulting range and minimum cross-range resolution in the ISAR images are  $0.075 \times 0.19\text{m}$  respectively. We obtain 3750 clean images that are free of noise and clutter. Then, we corrupt these images with additive white Gaussian noise in the time-domain to obtain noisy images with SNR ranging from  $-5$  to  $+10\text{dB}$ , resulting in 14976 noisy images. Similarly, we introduce range-Doppler clutter with four different wind speeds ( $U$ ) ranging from 2.5 m/s to 10 m/s to obtain 14976 cluttered images. We publicly released this data set to the research community at <https://tinyurl.com/msu6aj7y>.

### 2.1.5 Study of effects of shadowing on radar signatures

In our simulation model, we resolve the target into triangular facets and identify a scattering center at the centroid of each of the facets. Then we sum the returns from all the facets to compute the time-domain radar returns. Naturally, some of these facets would not be visible to the radar due to shadowing. The visibility of a facet can be determined through ray tracing. However, the process is computationally complex and scales both with the carrier frequency and sampling frequency of the radar data. To compare the computational complexity, the ray tracing algorithm (based on the surface normals of the facets) was implemented on a Intel(R) Xeon(R) CPU E5-2620 V3 processor running at 2.40 GHz using 16 GB of RAM. The duration of data collection for a single frame (shown in Fig.2.15(vi)) was 179.33 minutes. The computation time of ray tracing algorithm using advanced GPU reduces significantly, but the cost of the system increases which makes the technique challenging for generating large volumes of data [57]. Instead, we incorporate a Bernoulli's random variable ( $\eta$ ) to account for the visibility. The computational time for generating the ISAR image using the probabilistic method was a few seconds. To study the effectiveness/accuracy of this probabilistic approach, we compare the ISAR signature results with that obtained from ray tracing in Fig.2.15. The first five subfigures show the ISAR images of the auto-rickshaw with visibility of  $\eta$  fixed at 60%, 50%, 40%, 30%, and 20% respectively, and (vi) is the ISAR image of the same frame using ray tracing. We observe in the result from ray tracing, that the nearer portions of the




Figure 2.15: (i-v) ISAR images of auto-rickshaw for the same CPI frame using Bernoulli's random variable with visibility of 60 %, 50%, 40 %, 30%, 20 % respectively. (vi) is the ray tracing algorithm ISAR images of auto rickshaw for the same CPI frame. The range axis spans from 0 to 20m while the cross-range axis spans -15 to +15m.

target are more distinctive than the distant portions though they are somewhat visible. Visually, the image with 20% visibility resembles the ray tracing result most closely. Quantitatively, we compare the results using two metrics - the structural similarity index measure (SSIM) and the normalized mean square error (NMSE) - and present it in Table.2.3. The SSIM compares gross features like luminosity and contrast between two images while NMSE does pixel-wise comparison.

Table 2.3: Structural similarity index measure (SSIM) in ISAR images generated using ray tracing and Bernoulli random variable

% of facets visible	SSIM	NSME
60	1	0.8636
50	1	0.4470
40	1	0.4337
30	1	0.3465
20	1	0.3411

The table shows that the SSIM is high for all five cases but the NMSE is

lowest for a visibility of 20%. Further, in the ISAR images generated using the measurement data in Fig.2.18, we observe that some of the obscured parts of the vehicle are not fully shadowed. Due to these reasons, we use the approximate and less computationally expensive probabilistic method, with a visibility coefficient of 20%, to account for the shadowing effect in place of the ray tracing algorithm.

In the next section, we experimentally demonstrate the validity of our simulation framework by comparing the results with those generated from measurement data.

## 2.2 Measurement Data

The experimental set up for the measurements is shown in Fig.2.16. We configured the TI AWR-1843 radar to operate in an ultra-short range radar (USRR) mode and the corresponding parameters are listed in Table.2.1. The transmitted power from the radar is 14dBm, and the noise floor of the receiver is approximately -110dBm. We considered an auto-rickshaw (Fig.2.16a) of  $2.6m \times 1.3m \times 1.7m$  dimensions, a bicycle (Fig.2.16b) of  $1.6m \times 0.5m \times 1.5m$  dimensions, and a mid-sized car - Honda Brio (Fig.2.16c) of  $3.6m \times 1.7m \times 1.5m$  dimensions for the measurement data collection. Note that these targets resemble the ones used in the simulations in terms of gross size. However, the exact make and model are not the same in both scenarios. We could not carry out experiments with the truck and large size car due to the lack of availability of such targets in our premises. The simulation scenario considers 16 possible trajectories. Not

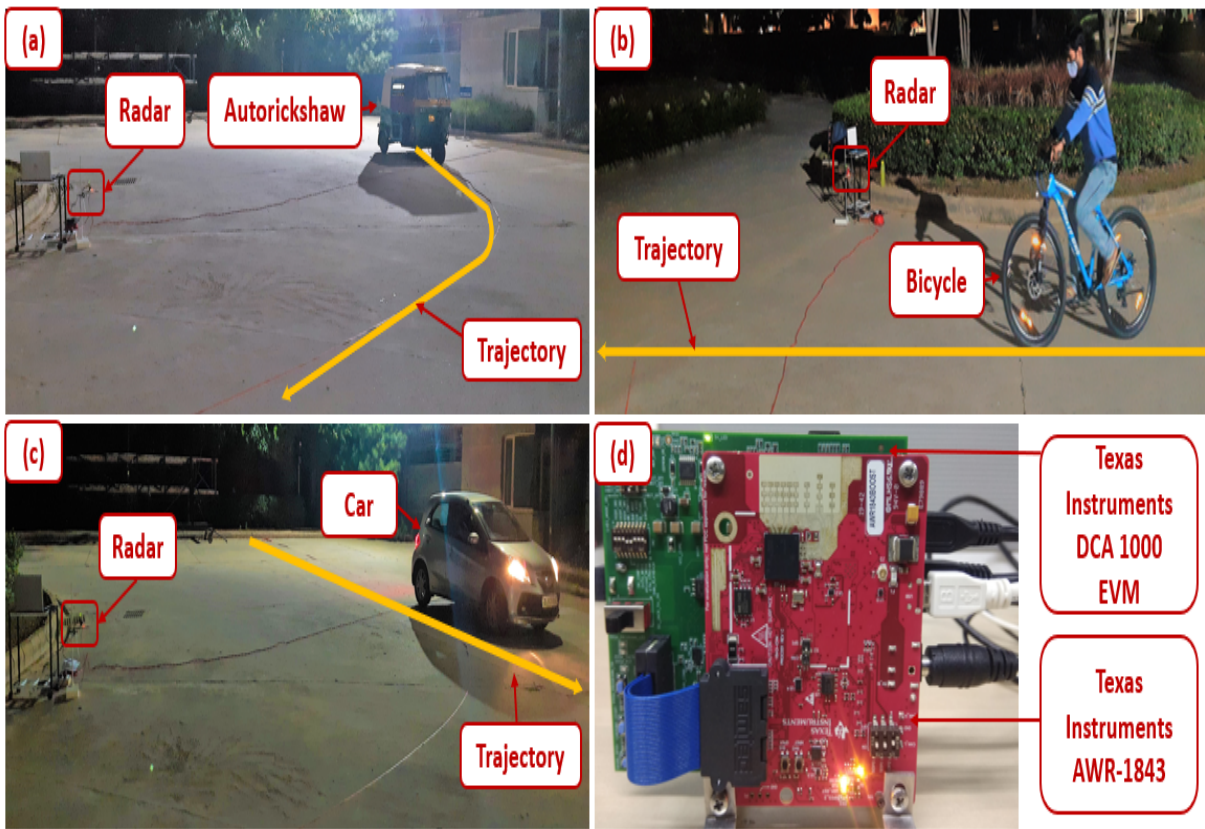


Figure 2.16: Experimental setup for gathering ISAR images of (a) an auto-rickshaw undertaking a right turn trajectory, (b) a bicycle undertaking a straight trajectory, (c) a small sized car undertaking a left turn along their trajectory (d) Radar hardware.

all of these trajectories could be replicated in our institute premises. Hence, we considered five trajectories as shown in Fig.2.17 - two straight trajectories, two right turn trajectories, and a left turn trajectory. These trajectories are similar to




Figure 2.17: Experimental setup and trajectories for gathering ISAR images of (a) two straight trajectories before the radar, (b) two right turn trajectories before the radar, (c) a left turn trajectory before the radar.

the trajectories chosen for the simulation data. In the two straight trajectories, the vehicles move along a straight path from left to right and right to left at an average speed of 15 kmph as shown in Fig.2.17b. In the first right turn trajectory, vehicles first move on the straight path and then execute a right turn in front of the radar, whereas in the second right turn trajectory, the vehicles first move on the straight path right to the radar and then execute a right turn as shown in Fig.2.17c at an average speed of 15 kmph. In the left turn trajectory, the vehicle first move on the straight path and then execute a left turn as shown in Fig.2.17d at an average speed of 15 kmph.

We performed matched filtering along the fast time and Doppler processing along the slow time to obtain ISAR images of the target for every CPI. We center the target in each measurement image by performing coarse and fine

range compensation based on [83]. The radar data is carefully calibrated based on the radar range equation, and the dynamic range of the images are fixed from -40dBm to -120dBm. The ISAR images of a bicycle moving on left turn trajectory are shown in Fig.2.18 (i-iv) at CPI frame corresponding to 1,2,3,4 seconds. The Fig. 2.18(v-viii) are ISAR images of auto-rickshaw moving along the straight trajectory at CPI frame corresponding the 1,2,3,4 seconds. The Fig. 2.18(ix-xii) are the ISAR images of the car executing the straight trajectory at the CPI frame corresponding to 1,2,3,4 seconds. These images show the top-view dimensions of the vehicle along the range and cross-range dimensions from which we can infer the size of the vehicle. Interestingly, in the measurement results, we can observe the back of the car, which is not shadowed entirely by the front of the vehicles in some time intervals. We are able to see the obscured part of the vehicle due to the inherent advantages of ISAR processing, wherein reflections from different aspect angles get captured due to the relative rotational motion of the vehicle with respect to the radar. Another interesting observation is that the movie of the ISAR images provides insights on how the car is turning based on the orientation of the longer dimension of the car. In some frames, we observe a large spread along the cross-range dimension due to the micro-Dopplers arising due to the wheel motions. The strength of the radar returns are also fairly comparable between the two sets of results. We are also able to observe range based clutter in the measurement results increasing with for higher values of range as observed in the simulations. We also see Doppler-based clutter with the highest returns at cross-range values close to  $0m$  corresponding to static

and low-frequency clutter. Hence, there is, overall, a good qualitative agreement between the measured and simulated images. For a representative case, Fig. 2.19 shows a side-by-side comparison between simulated and experimental ISAR images of the bicycle target, where the simulated images closely reproduce the main micro-Doppler features due to the rotating wheels and the overall Doppler spread observed in the measurements. A detailed quantitative comparison is not reported since it was not possible to measure the exact clutter conditions in the real data or to replicate identical targets and trajectories in the two scenarios.

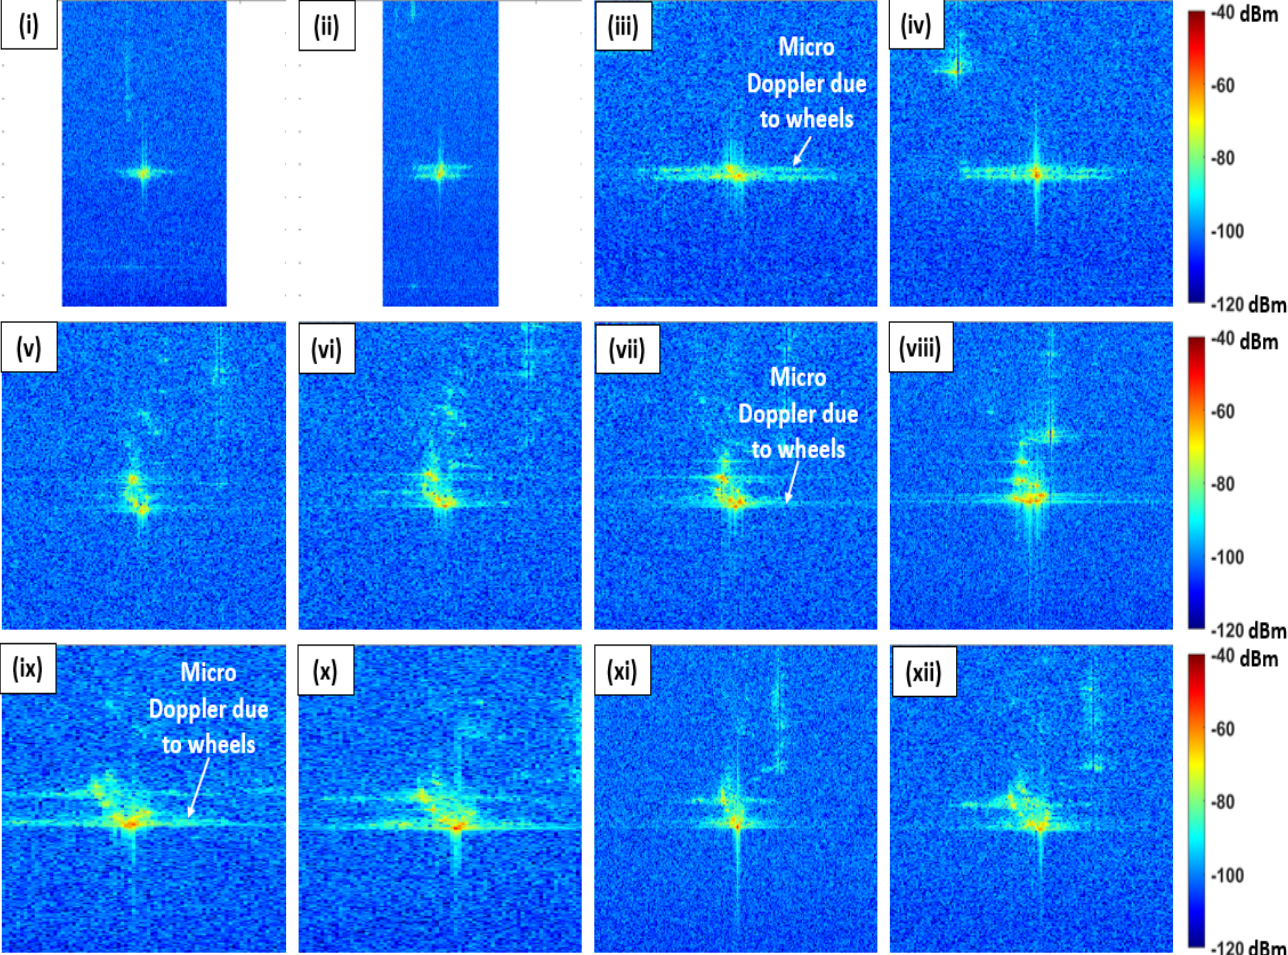


Figure 2.18: ((i-iv) ISAR images of bicycle along the left turn trajectory at CPI frame corresponding to 1,2,3,4 s, (v-viii) ISAR images of auto-rickshaw along the straight trajectory at CPI frame corresponding to 1,2,3,4 s, and (ix-xii) ISAR images of car along the straight trajectory trajectory at CPI frame CPI frame corresponding to 1,2,3,4 s from measurement data. Range axis spans from 0 to 20m while the cross range axis spans -15 to +15m.

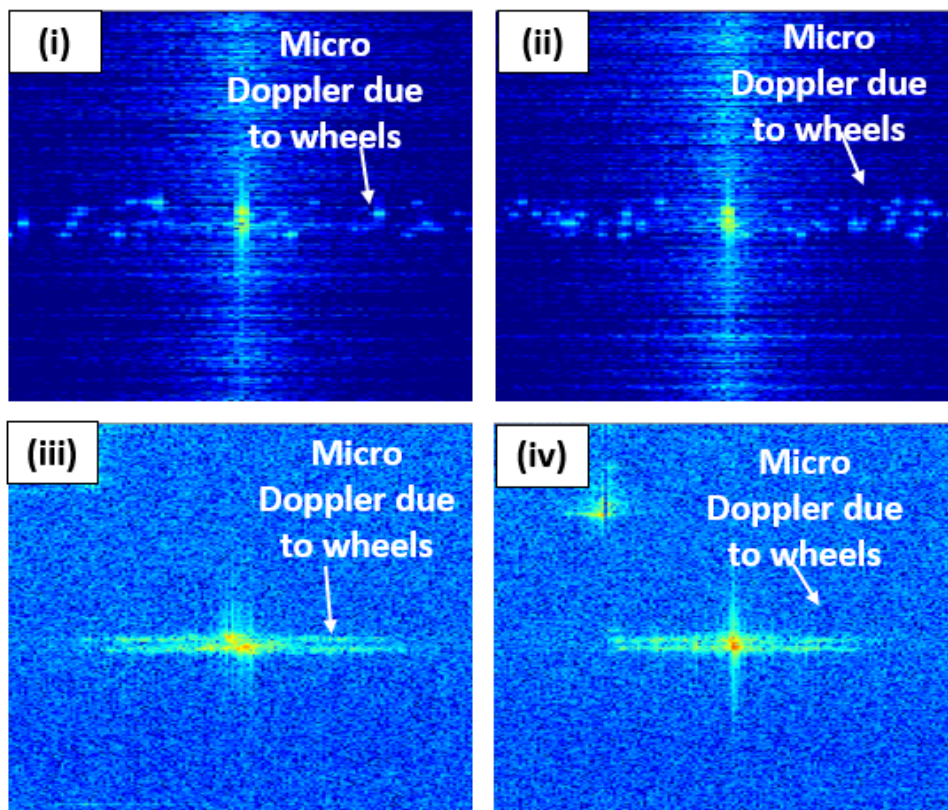


Figure 2.19: Qualitative comparison of simulated and measured ISAR images of a *bicycle*. Sub-figures (i) and (ii) show simulated images generated using the proposed ISAR simulator for two representative CPI frames, while sub-figures (iii) and (iv) show the corresponding ISAR images obtained from measurement data. In all four sub-figures, the micro-Doppler streaks due to the rotating wheels and the overall range and cross-range spread are clearly visible. The simulated signatures closely match the measured ones, indicating good qualitative agreement between the simulated and real data.

## 2.3 Classification Results

In this section, we use classical ML techniques - SVM and RF - and more recent deep learning algorithms based on transfer learning - Alexnet and Googlenet - for classifying the five automotive targets on the basis of their ISAR images. We will examine the impact of noise and clutter and the volume of test and training data on the classification performance. We will present the classification result using measurement data ISAR images as test data for the classifier trained on the simulation data. We use ImageNet-pretrained *AlexNet* and *GoogLeNet*. ISAR images are resized to the network input size and adapted to three channels. For **AlexNet**, we replace the final fully connected and classification layers with a 5-class head and then fine-tune the network end-to-end using stochastic gradient descent with momentum (SGDM). For **GoogLeNet**, we replace the last learnable and classification layers with a 5-class head, freeze the early layers, and fine-tune the remaining layers using the same optimiser; light geometric augmentation is applied during training.

### 2.3.1 Effect of noise and clutter on classification performance

Based on Table.2.2, of the total volume of 14976 images for different SNR values, 70% are used for training and 30% for testing in the case of SVM and RF. In the case of Alexnet and Googlenet, we split the 30% data that are not used for training between validation and testing. The resulting classification accuracy for different SNR values is shown in Fig.2.20a.

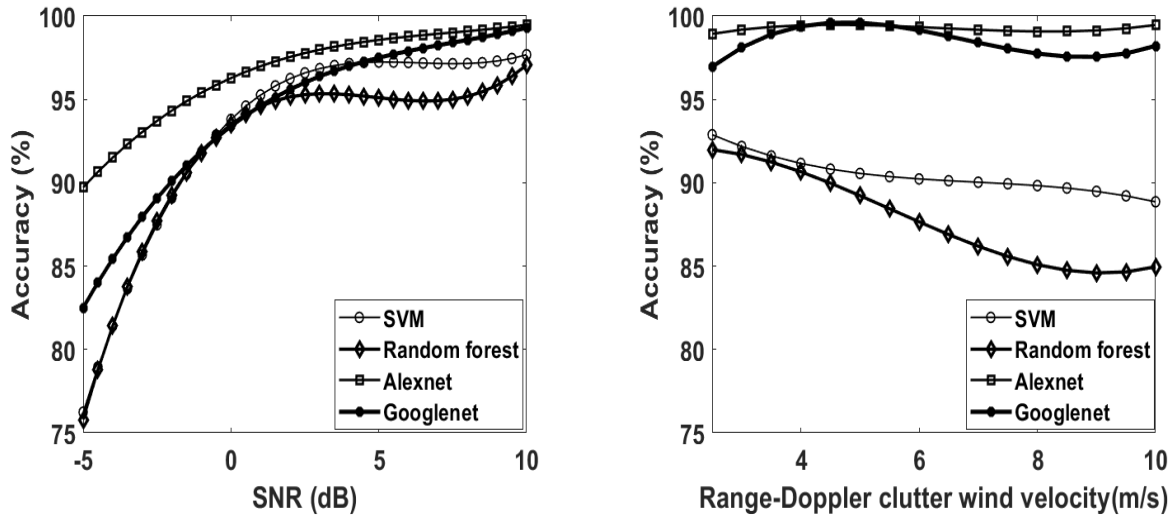


Figure 2.20: Classification accuracy of automotive target ISAR images using SVM, RF and transfer learning algorithms based on Alexnet and Googlenet for (a) differing SNR values and (b) different range-Doppler clutter values. 70% data used for training and remaining data used for validation and testing.

We first observe that the classification accuracy for all algorithms is above 75% even for low SNR of  $-5dB$ . The accuracy for SVM and RF are significantly poorer than those obtained from Alexnet and Googlenet at low SNR ( $-5dB$ ). The classification accuracy improves for all cases as the SNR increases. The performances of Alexnet and Googlenet hold steady (above 80% for Googlenet and 90% for Alexnet) for all cases.

We perform a similar study where we examine the effect of clutter on the classification performance of the ISAR images. The clutter along the range is modeled using a mean surface clutter coefficient. As the range increases, the area of coverage increases resulting in greater clutter. Wind gives rise to Doppler based clutter along the cross-range dimension. Higher wind velocities ( $U$ ) give rise to greater clutter. Again, we have assumed a 70-30 split between training and test data for SVM and RF and a 70-15-15 split between training, validation, and test for Alexnet and Googlenet. We show the variation of the

classification accuracy with respect to mean wind velocity in Fig.2.20b. We observe that the classification performance for all the algorithms is fairly high (above 85%) even for high values of clutter arising from high wind speeds (10m/s). The performance of the two transfer learning-based algorithms (Alexnet and Googlenet) remains consistent even for the high values of clutter. On the other hand, we observe a slight deterioration in the performance of the two classical ML techniques with higher clutter values. The results indicate that the ISAR images offer highly discriminatory features for classification, even in the presence of high noise and clutter.

### 2.3.2 Effect of test and training percentages

In Table.2.4, we report the classification results for the four algorithms - SVM, RF, Alexnet, and Googlenet - for different percentages of training, testing, and validation data.

Classifier	Training (%)	Testing (%)	Validation (%)	SNR	Range-Doppler Clutter	Combined
SVM	70	30	-	$89.8 \pm 0.6$	$93.1 \pm 0.3$	$86.5 \pm 0.1$
	60	40	-	$89.7 \pm 0.2$	$93.0 \pm 0.3$	$86.3 \pm 0.3$
	50	50	-	$89.3 \pm 0.3$	$92.6 \pm 0.2$	$86.2 \pm 0.1$
Random Forest	70	30	-	$88.8 \pm 0.3$	$94.6 \pm 0.1$	$91.2 \pm 0.2$
	60	40	-	$88.3 \pm 0.1$	$93.9 \pm 0.2$	$90.7 \pm 0.2$
	50	50	-	$97.7 \pm 0.2$	$93.7 \pm 0.5$	$90.2 \pm 0.2$
Alexnet	70	15	15	$96.9 \pm 0.7$	$99.1 \pm 0.4$	$98.1 \pm 0.5$
	60	20	20	$96.6 \pm 0.7$	$99.1 \pm 0.3$	$97.9 \pm 0.3$
	50	25	25	$96.2 \pm 0.8$	$98.3 \pm 0.7$	$97.3 \pm 0.2$
Googlenet	70	15	15	$94.6 \pm 0.4$	$98.9 \pm 0.2$	$97.2 \pm 0.2$
	60	20	20	$94.4 \pm 0.5$	$98.6 \pm 0.3$	$97.1 \pm 0.1$
	50	25	25	$94.1 \pm 0.3$	$98.3 \pm 0.4$	$96.9 \pm 0.3$

Table 2.4: Classification of ISAR images using classical ML algorithms - SVM and RF, and deep learning based algorithms - Alexnet and Googlenet

For each case, we have performed a 5-fold cross-validation. We first consider the data that are just corrupted by noise (SNR varying from  $+10dB$  to  $-5dB$ ). Then we repeat the tests on data that are just corrupted by only clutter (wind velocities varying from 2.5 to 10m/s). Finally, we repeat the tests on data that are corrupted by both noise and clutter. We observe that the performances for all the cases are above 87%. Since the training data is very large, the algorithms perform well even when the training and testing are split evenly. We do not see a significant improvement in the performance with an increase in training data. The transfer learning algorithms like Alexnet and Googlenet perform very well, even for low SNR and high clutter.

In the following section, we present the confusion matrices obtained from the classification of data combining both noise and clutter. These results are obtained using 70% training data. For each case, the rows show the true labels of the data, while columns show the labels of the predictions. We present three metrics with each confusion matrix. They are *precision*, *recall*, and the  $F_1$  score. Precision is defined as the ratio of the true positives (the highlighted number along the diagonal) over the sum of the true positives and false positives (column sum); while recall is defined as the ratio of the true positives over the sum of the true positives and false negatives (row sum). For each case, we also provide the  $F_1$  score, which is defined as the harmonic mean of average precision and average recall,

$$F_1 = 2 \frac{avg.precision \times avg.recall}{avg.precision + avg.recall}. \quad (2.26)$$

The first confusion matrix is presented for the SVM in Table.2.5. We observe

SVM		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	1655	3	41	96	2	92.1
	Bicycle	24	1687	17	72	3	93.6
	Full size car	48	39	1478	193	34	82.5
	Mid size car	111	112	146	1463	2	79.8
	Truck	14	2	56	15	1673	95.1
	Precision	89.4	91.5	85.0	79.6	97.6	
RF		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	1674	14	13	96	0	93.2
	Bicycle	4	1723	11	63	2	95.6
	Full size car	29	34	1577	145	7	88.0
	Mid size car	65	130	67	1570	2	85.6
	Truck	3	4	24	18	1711	97.2
	Precision	94.3	90.4	93.2	83.0	99.4	
Alexnet		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	895	1	0	2	0	99.7
	Bicycle	0	897	1	1	0	99.8
	Full size car	3	1	876	14	5	97.4
	Mid size car	5	40	5	849	0	94.4
	Truck	2	0	4	0	893	99.3
	Precision	98.9	95.5	98.9	98.0	99.4	
Googlenet		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	881	0	8	7	2	98.1
	Bicycle	0	890	0	9	0	99.0
	Full size car	0	3	863	32	1	96.0
	Mid size car	0	37	2	860	0	95.7
	Truck	0	0	8	1	891	99.0
	Precision	100	95.7	98.0	94.6	99.7	

Table 2.5: Confusion matrices from SVM, RF, Alexnet and GoogleNet classifiers based on 70% training, 15% validation and 15% test data.

the least confusion for the truck. Due to its large dimensions and strong returns, the truck is rarely mistaken for any of the other targets or vice versa. Similarly, the bicycle is very small and hence not easily confused with the other targets. However, due to its weak returns and small size, sometimes, the bicycle is not easily discernible in noisy images. The two cars are often confused by each

other due to the similarity in their dimensions, the number of wheels, and their strength of returns. The mid-size car, especially, has the lowest precision and recall because it is most similar to both the full-size car and the auto-rickshaw. The  $F_1$  score for SVM is 88.6%.

We observe a similar result for the RF classifier in the second matrix in Table.2.5. Again, the best precision and recall are observed for the large truck and the small bicycle. The results of the bicycle are slightly worse than the truck because of its weak returns, which get affected when the noise floor is high. Again, the two cars are confused by each other. However, this time, the results for the full-size car have significantly improved while those of the mid-size car have only slightly improved. There is a noticeable improvement in the performance of the Alexnet classifier, compared to the traditional ML algorithms for all cases, as reported in the third matrix in Table.2.5. Here, both the precision and recall for all the cases are above 95%. Thus the two cars are no longer significantly confused by each other. The confusion between the auto-rickshaw and mid-size car has also substantially decreased. The same improvement is also observed for the Googlenet classifier as seen in the fourth matrix in Table.2.5. Again, the accuracy is above 95% for all five cases, both precision and recall. The  $F_1$  scores for RF, Alexnet, and Googlenet are 92%, 98.1%, and 97.6%, respectively. Note that in all of the cases discussed above, we use only a single target instance per class. The limitation is due to the limited availability of open data of computer animated models of vehicles provided by their manufacturers. The diversity is incorporated by varying the trajectories, speeds, noise, and clutter conditions in

the radar data.

These results are obtained with the simulator’s clutter modeled as exponential (power) speckle using a scenario-specific mean  $\sigma_0$  and a separable range–Doppler spectrum as defined in Section 2.1.3. More robust classification tests with multiple instances from the same class will be presented in future studies.

### 2.3.3 Testing measured data with classifiers trained with simulation data

We present the classification results of two classifiers - SVM based on classical ML and Alexnet transfer learning-based deep learning technique. Both classifiers are trained on simulation data. We use 99 images for each of the three automotive targets - bicycle, auto-rickshaw, and medium-sized car - as test images for the classifier performance evaluation. In Table.2.6 we present confusion matrices for both the classifiers. For the SVM classifier, we observe that the least confusion

SVM		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	76	22	0	2	0	76.0
	Bicycle	4	92	0	4	0	93.9
	Full size car	0	0	0	0	0	-
	Mid size car	22	25	1	50	2	51.0
	Truck	0	0	0	0	0	-
	Precision	75.5	66.2	-	89.3	-	
Alexnet		Predicted Labels					
	Vehicle	Auto-rickshaw	Bicycle	Full size car	Mid size car	Truck	Recall
True Labels	Auto-rickshaw	80	0	0	18	1	80.8
	Bicycle	0	99	0	0	0	100
	Full size car	0	0	0	0	0	-
	Mid size car	0	0	0	99	0	100
	Truck	0	0	0	0	0	-
	Precision	100	100	-	84.6	-	

Table 2.6: Confusion matrices from SVM and Alexnet and GoogleNet classifiers, trained with simulation data and test with measured data.

is for the bicycle because it is smaller (in terms of spatial extent) than the other targets. Confusion arises between the car and auto-rickshaw because of their comparable sizes. The performance of the Alexnet classifier is better than the SVM classifier and the confusion between the auto-rickshaw and the mid-size car is reduced. The  $F_1$  scores for SVM and Alexnet are 75.3 %, and 94.2% respectively. The results indicate that the simulated ISAR images are of high fidelity and can be used to train classifiers that can subsequently be used on real test data.

The classifier hyper-parameters are summarised in Table 2.7.

Table 2.7: Hyper-parameters used for the SVM and Random Forest classifiers.

Classifier	Parameter	Value
SVM (SVC)	kernel	linear
	$C$	1.0
	probability calibration	enabled
	random_state	42
Random Forest	$n_{estimators}$	100
	max_depth	14
	criterion	gini
	max_features	sqrt
	random_state	0

## 2.4 Summary of findings

- We have demonstrated an automotive radar simulation framework that incorporates radar scattering phenomenology of commonly found road vehicles as well as range-based surface clutter and Doppler-based wind clutter and additive receiver noise.
- Using this simulation framework, we have demonstrated that high-resolution

ISAR radar images, characterized by the fine range and cross-range resolution of dynamic automotive targets, can be generated with mm-wave automotive radars.

- A large database of over **30,000** images has been publicly released to the radar community.
- The simulation framework has been verified through experimental data gathered with a real automotive mm-wave radar from Texas Instruments.
- These images provide meaningful information about the dimensions of the vehicle along the top-view as well as the number of wheels and the trajectory undertaken by the vehicle in the case of larger vehicles such as auto-rickshaws, cars, and trucks. Smaller targets such as bicycles, on the other hand, more closely resemble single-point scatterers.
- These images indicate the robustness of ISAR images as automotive target classification features for both traditional ML techniques as well as the more recent deep neural networks.
- Due to restrictions in the data collection facilities within the institute premises, we have limited our data collection to single target instances of each target class. This may be satisfactory in the cases of auto-rickshaw where the models and makes are fairly standard. However, in other types of vehicles, there may be greater diversity in each target class. The testing of the classifiers on multiple instances of the same target class will form the focus of future works.

- *Target diversity and scope.* Due to constraints in the data-collection facilities, measured data were limited to a single physical instance per class. This reduces—but does not eliminate—intra-class variability captured here (e.g., accessories, body variants, etc.). Accordingly, the claims in this chapter are limited to *class-level* discrimination under these data constraints.
- Also, note that the specific problem addressed in the work falls in the category of supervised learning — where the classifiers are trained with data from specific targets and tested with data from the same set of targets. In real world scenarios — especially in automotive radar settings — we may encounter new targets that may have previously not been seen during training. Such problems would fall under the category of unsupervised learning. The exploration of new and better algorithms for handling unknown radar data will form the basis of our future work.

## Chapter 3

# Characterization of Millimeter Wave Clutter on Indian Roads at Automotive Radar Frequencies

*In this chapter, we presented the statistical analysis of road clutter gathered from experimental data collected on four different road surfaces of varying roughness across diverse environmental conditions - temperatures, wind speeds, and surface moisture. Then, we have estimated the statistics of the surface clutter coefficient,  $\sigma^0$ , from the measurement data by fitting the data with popular land clutter models in current literature such as lognormal, Weibull [64, 84], Rician Gamma and Gaussian mixture models. Based on two different scores - negative log-likelihood (NLL), and Bayesian Information Criteria (BIC) score, [85–87] - derived from statistical principles [88, 89], we find that the Gaussian mixture model is the best fit for the measured*

*data to capture the variations with temperature and grazing angle.*

*System simulation studies and experiments have demonstrated that the performance of automotive radars is significantly influenced by road surface backscatter [90–93]. When the signal-to-clutter and noise ratio is poor, the probability of detection of true targets reduces while the probability of false alarms increases. Road clutter can show significant variation based on the type of road material, the roughness of the road, wind, moisture and the presence of buildings in the neighborhood. Hence, the study of the statistical characteristics of road backscatter at automotive radar frequencies is crucial for several reasons. First, a detailed study of clutter enables the determination of appropriate thresholds for radar detectors. Second, in the development of radar simulators for automotive targets, which are used to generate large volumes of training data for ML algorithms, it is crucial to incorporate accurate clutter characteristics within these models. This ensures that classifiers trained on simulated data perform effectively when tested with real-world data. There has been extensive research to model the statistical parameters of land [64, 94, 95] and sea clutter [84, 96], over the years, at microwave frequencies. In [61, 81], the authors presented a theoretical model for estimating the backscatter from different types of roads with validation in laboratory conditions. In [63], empirical studies were conducted at 77 GHz using a vector network analyzer for various types of road surfaces in wet, dry, and snow-covered con-*

*ditions. However, statistical analyses of these results were limited. Experimental studies on European roads at 24GHz were carried out in [97] while in [98], statistical analyses of radar data were carried out to identify different types of road structures such as tunnels and guard rails. In this work, we carry out a systematic examination of the clutter characteristics in urban Indian road conditions using a commercial off-the-shelf frequency modulated continuous wave (FMCW) radar at 77GHz. Based on our experimental studies, Indian roads are significantly rough. Hence, a detailed examination of their clutter statistics is required before the widespread deployment of automotive radars.*

*This chapter is organized as follows. In the following section, we present the experiments for collecting the radar data along with the radar signal processing for calculating the surface clutter coefficient. We present the measurement results along with the insights obtained from statistical analysis in section.3.2. We discuss summary of findings of experiments in section 3.3*

### **3.1 Experimental Setup**

In this section, we describe the procedure for estimating the surface clutter coefficient ( $\sigma^0$ ) of different types of road surfaces using Texas Instruments AWR 1843 FMCW radar sensor. The process consists of three steps: first, we calibrate the mm-wave radar with a target of known radar cross-section; second, we

collect a large volume of radar backscatter of the road in real-world conditions and estimate the  $\sigma^0$  and third, we estimate the roughness of the surface. We elaborate on these steps below. We perform the calibration experiment using the setup shown in Fig.3.1.



Figure 3.1: Experimental setup to calibrate the mm-wave radar at 77GHz using a trihedral corner reflector as a calibration target of known radar cross-section.

The radar is mounted on a static platform and is configured for short-range operation with the parameters given in Table.3.1. This configuration is selected to be consistent with the ISAR simulation/measurement parameter set in Chapter 2 (Table 2.1) wherever possible (same platform and carrier frequency, and the same  $T_{CPI}$  to keep the coherent integration duration comparable). Table 3.1 is used specifically for the static short-range clutter measurement campaign, and therefore some parameters differ from Table 2.1 to satisfy practical IF/ADC constraints under stretch processing and to support clutter characterisation. In particular, a larger sweep bandwidth (4 GHz) and sampling frequency (10 MHz) are used here to enable finer range gating of the near-field clutter patches and to reduce mixing/leakage between neighbouring range bins during statistical estimation. The chirp rate  $\gamma$  and sampling  $F_s$  are chosen such

that the resulting beat frequencies remain within the available IF/ADC band, i.e.,  $f_b(R) \approx \frac{2\gamma R}{c} \leq f_{b,\max} = \min\{B_{\text{IF}}, F_s/2\}$ , for the short-range geometry considered in this chapter.

Table 3.1: TI-AWR 1843 radar parameters

Parameters	Values
Carrier frequency	77GHz
Bandwidth	4GHz
Sampling Frequency	10MHz
Chirp rate	$49.97 \times 10^{12}$ Hz <sup>2</sup>
Chirp duration	400 $\mu$ s
Coherent processing interval	0.1s
Ramp time	80.04 $\mu$ s
Idle time	319.95 $\mu$ s
Transmitted power	14dBm

The radar transmits a FMCW radar signal with  $f_c$  carrier frequency (and wavelength  $\lambda_c$ ) and a chirp rate  $\gamma$  as given by 2.5.

The down-converted radar received signal from the calibration target is

$$s_{rx}(\tau) = A(\cdot) \text{rect}\left(\frac{\tau - \frac{2R}{c}}{T_{PRI}}\right) e^{j\pi\gamma(\tau - \frac{2R}{c})^2} + \nu, \quad (3.1)$$

where  $\nu$  is the additive receiver noise,  $A$  is the strength of the received signal and  $R$  is the distance from the target. To find  $A$ , we perform stretch processing along  $\tau$  to obtain,

$$S_{rx}(\tau) = A(\cdot) \text{sinc}\left(\tau - \frac{2R}{c}\right). \quad (3.2)$$

Here  $\text{sinc}(x)$  is the well known expression  $\sin x/x$ . Note that through Friis radar

range equation [99] shown in

$$A^2 = \frac{P^{tx} G^2 \lambda^2 \sigma_{Tri}}{(4\pi)^3 L_s R^4}, \quad (3.3)$$

$A$  is a function of the transmitted power  $P^{tx}$ , antenna gain  $G$ , system loss  $L_s$  and the radar cross-section of the calibration target. In this case, the target is a trihedral corner reflector and  $\sigma_{Tri} = \frac{4\pi a^4}{3\lambda^2}$  [100] where  $\lambda$  is the wavelength and  $a$  is the inner edge of trihedral and equal to 0.7m. Thus, all the radar parameters in (3.3) can be grouped into one common term  $K$  which is estimated from

$$K = \frac{A^2 R^4}{\sigma_{Tri}}. \quad (3.4)$$

We perform this experiment out in the field to reduce multipath from walls and ceilings. We repeat this experiment with five different values of  $R$ : 5, 10, 15, 20, 25m and calculate  $K$  for each case. Then the average value of  $K$  is used for the remaining experiments. In order to estimate  $\sigma^0$ , we consider four different locations as shown in Fig.3.2. The first three locations consist of concrete roads while, the fourth is an asphalt road. The road width for the first three locations is 7m, and for the fourth location, it is 9m. In locations 1 and 2, the road is banked by a high wall of 2.5m height and a tall building on the other side. In location 3, one side is banked by a wall, but the other side is close to free space conditions. In the fourth location, both sides of the road are banked by walls of low height (1.82m). We captured the radar backscatter data from the road surfaces at the four locations indicated in Fig.3.2 over several days and at several different times of the day - morning, noon, evening, and midnight. The temperature variation



Figure 3.2: Different locations at which road surface clutter coefficient was estimated. Locations 1, 2, and 3 correspond to 7m wide concrete road, while location 4 is a 9m wide asphalt road.

during the measurements spanned from 24°C to 46°C, while the wind velocity varied from 2m/s to 5.8m/s (gust). The surface moisture varied from 0.45% corresponding to dry conditions to 2.1% corresponding to wet conditions after rain. Each measurement was carried out at four different heights of the radar - 0.3, 0.5, 0.75, and 1m corresponding to different grazing angles. Across five field days, this resulted in  $4 \times 4 \times 4 \times 5 = 320$  radar backscatter acquisitions. In each case, the radar is mounted on a static platform and faces forward (90°) from the ground to resemble the scenario of the actual deployment of automotive

radar. The radar backscatter from the road is given by

$$P_{rC} = \frac{K\sigma^0 A_c}{R_c^4}, \quad (3.5)$$

where  $K$  is the calibration constant that we estimated earlier,  $A_c$  is the clutter resolution area, and  $R_c$  is the distance to the road surface. Note that  $K$  was computed for a scenario when the target was directly facing the radar at boresight. Hence, in the above expression, we only consider the radar returns from that part of the road section that falls within the main lobe of the radar antennas. If  $\Theta_b$  is the radar antenna's azimuth beamwidth, and  $\Delta_r$  is the range resolution and  $\psi$  is the grazing angle, then  $A_c$  is  $R_c\Theta_b\Delta_r\sec(\psi)$ . By repeating the identical steps from (2)-(3), we estimate  $\sigma^0$  for the different road surfaces. Next, we estimate the roughness of the roads using the setup shown in Fig.3.3. Here, we mount the Bosch GLM 500 laser meter on a sliding scale of 1m in length and 1.008m in height for calculating the distance of the ground from the laser. Measurements are made at regular intervals along the ground by sliding the laser meter along the scale. We rotate the setup by its axis (about the stand) and repeat the measurements at every 15 degrees to complete an entire circle. In this manner, we obtain 144 polar measurements (counting both radial and angular positions) for a fixed position of the stand. Since the stand is of a fixed height, the difference between each laser depth measurement,  $h_i$ , and the height of the stand provides the local variation of the ground surface. The standard deviation of these variations provides the roughness of the surface for each location. The road roughness values for the different locations are provided in Table.3.2. It is




Figure 3.3: Experimental setup to estimate the road roughness using Bosch GLM 500 laser meter mounted on a sliding scale of 1m length and 1.008m height.

evident from these measurements that all the roads are significantly rough based on the international road standards [101]. We use the Stanley STHT770030

Table 3.2: Road roughness variance

Location	Road material	Roughness (mm)
Location 1	Concrete	6.4
Location 2	Concrete	7.7
Location 3	Concrete	8.7
Location-4	Asphalt	6.7

moisture meter for measuring the surface moisture along several points in the same four locations.

### 3.2 Analysis and Observations

In this section, we present the histograms of the  $\sigma^0$  estimated from the radar data as a function of the four road locations. Then we fit five different probability distribution functions - lognormal, Gamma, Rician, Weibull, and GMM - that have been used in radar literature to model the  $\sigma^0$  values [64, 84]. We examine the best-fit model based on the negative log-likelihood (NLL) and BIC scores. The lognormal distribution is a probability distribution function whose logarithm has a normal distribution [102]. It is characterized by two parameters - the mean of the logarithmic values,  $\mu$ , and the standard deviation of the logarithmic values,  $\sigma$ . The Gamma distribution is a two-parameter family of distributions used to model sums of exponentially distributed random variables [103]. The exponential distribution and the chi-square distributions are special cases of the Gamma distribution. Due to its high generalization, it has been used to model different types of land and sea clutter. It is characterized by the shape parameter,

$k$ , and the scale parameter,  $\theta$ . The Rician distribution is used to model the magnitude/envelope of the vector sum of in-phase and quadrature components of complex radar data [102]. It is characterized by a parameter that indicates the non-centrality of the mean,  $\nu$ , and a scale parameter,  $s$ . Another popular distribution function that has been used to model clutter data is the Weibull which is characterized by two parameters - the mean of the distribution,  $\eta$ , and the shape parameter,  $\beta$  [102]. When  $\beta$  is one, the distribution corresponds to the exponential distribution, which has been used extensively in radar literature to model discrete targets. When  $\beta$  is two, it corresponds to Rayleigh distribution which arises when several independent scatterers of similar strengths occupy the same radar clutter cell. Usually, in real-world conditions,  $\beta$  lies somewhere between 1 and 2. A GMM is a mixture of a finite number ( $B$ ) of Gaussian distributions characterized by the component means ( $\mu_b, b = 1 : B$ ) and standard deviations ( $\sigma_b, b = 1 : B$ ), and the mixture component weights ( $w_b, b = 1 : B$ ).

A statistical model is represented by the probability distribution function  $p(d|\Theta)$  where  $d \in R^N$  are data, and  $\Theta$  is the set of all the parameters. If we have  $d_1, d_2 \dots d_F$  samples, then the NLL score is  $L(\Theta|d) = -\frac{1}{F} \sum_{i=1}^F \log p(d_i|\Theta)$ . So, the lower the score, the better the fit of the statistical model to the data. The BIC score is  $BIC = 2L(\Theta|d) + \log(F) \times g$ , where  $F$  is the number of the samples, and  $g$  is the number of the parameters. We fit each of the five distributions to the measured data by optimizing for the distribution parameters that result in the minimum NLL and BIC scores.

Based on all the experiments, the mean value of  $\sigma^0$  for the concrete and

asphalt roads are  $-25dB$  and  $-23.6dB$ , respectively. In Fig.3.4 we present the histograms of the estimated  $\sigma^0$  for four different roads with varying surface roughness. Fig.3.4(a), (b) and (c) correspond to a concrete road with an average roughness of 6.4 mm, 7.7mm and 8.7mm respectively, while Fig.3.4(d) is for an asphalt road with a roughness of 6.7mm. The measurement data are captured with varying environmental conditions such as temperature and radar height; due to this, histograms show more than one clustering of the data. We show the best-fitted statistical models for Gamma, log-normal, Rician, Weibull and GMM for each of these data. Based on qualitative assessment, none of the statistical models except for GMM satisfactorily capture the histogram envelope. For the GMM, we evaluated  $B \in \{1, 2, 3, 4\}$  and selected  $B=3$  by the minimum BIC averaged over the sites/conditions; larger  $B$  showed negligible change in  $BIC$  and smaller  $B$  underfit. In Table.3.3 we show the NLL and BIC scores for the six different distributions along with their optimized distribution parameters. Based on both the scores, the Gaussian mixture model (GMM) distribution is the best fit for the  $\sigma^0$  values. The GMM model suggests that multiple factors affect the surface clutter. To study this further, we show the mean  $\sigma^0$  values at different road locations with respect to different radar heights and varying temperatures in Fig.3.5. We observe that the mean  $\sigma^0$  decreases with an increase in the radar height while the mean  $\sigma^0$  varies significantly with the temperature, with a maximum in mid-temperature ranges. From these plots, it is clearly seen that the  $\sigma^0$  values increase with the surface roughness.

*Use in simulation.* The multi-modal  $\sigma_0$  behavior across sites/conditions

Road location	Distribution	Distribution parameters	NLL score ( $\times 10^4$ )	BIC score ( $\times 10^4$ )
Location 1- Concrete	Exponential	$\mu_0=3.1 \times 10^{-3}$	-2.87	-5.74
	Lognormal	$\mu=-5.79, \sigma=6.55 \times 10^{-2}$	-4.25	-8.51
	Gamma	$k=258.28, \theta=1.18 \times 10^{-5}$	-3.90	-7.80
	Rician	$\nu=3.04 \times 10^{-3}, s=1.8 \times 10^{-4}$	-3.93	-7.86
	Weibull	$\eta=3.13 \times 10^{-3}, \beta=18.57$	-4.30	-8.60
	GMM	$\mu_1=2.90 \times 10^{-3}, \sigma_1=5.78 \times 10^{-9}, w_1=0.50$ $\mu_2=3.20 \times 10^{-3}, \sigma_2=1.17 \times 10^{-8}, w_2=0.02$ $\mu_3=2.10 \times 10^{-3}, \sigma_1=6.26 \times 10^{-7}, w_3=0.48$	-4.40	-8.80
Location 2- Concrete	Exponential	$\mu_0=3.5 \times 10^{-3}$	-2.79	-5.59
	Lognormal	$\mu=-5.67, \sigma=0.12$	-3.84	-7.67
	Gamma	$k=92.27, \theta=3.77 \times 10^{-5}$	-3.93	-7.86
	Rician	$\nu=3.46 \times 10^{-3}, s=3.44 \times 10^{-4}$	-3.90	-7.80
	Weibull	$\eta=3.63 \times 10^{-3}, \beta=10.71$	-3.90	-7.80
	GMM	$\mu_1=3.40 \times 10^{-3}, \sigma_1=7.58 \times 10^{-8}, w_1=0.11$ $\mu_2=4.00 \times 10^{-3}, \sigma_2=2.21 \times 10^{-8}, w_2=0.65$ $\mu_3=3.40 \times 10^{-3}, \sigma_1=9.73 \times 10^{-8}, w_3=0.24$	-3.94	-7.88
Location 3- Concrete	Exponential	$\mu_0=3.7 \times 10^{-3}$	-2.76	-5.52
	Lognormal	$\mu=-5.61, \sigma=0.11$	-3.83	-7.66
	Gamma	$k=80.27, \theta=4.57 \times 10^{-5}$	-3.82	-7.65
	Rician	$\nu=3.65 \times 10^{-3}, s=0.413 \times 10^{-4}$	-3.81	-7.64
	Weibull	$\eta=3.86 \times 10^{-3}, \beta=9.37$	-3.79	-7.59
	GMM	$\mu_1=4.20 \times 10^{-3}, \sigma_1=5.95 \times 10^{-8}, w_1=0.35$ $\mu_2=3.80 \times 10^{-3}, \sigma_2=1.16 \times 10^{-7}, w_2=0.17$ $\mu_3=3.50 \times 10^{-3}, \sigma_1=6.81 \times 10^{-8}, w_3=0.48$	-3.92	-7.84
Location-4- Asphalt	Exponential	$\mu_0=4.4 \times 10^{-3}$	-2.65	-5.30
	Lognormal	$\mu=-5.41, \sigma=0.15$	-3.76	-7.53
	Gamma	$k=39.31, \theta=1.15 \times 10^{-43}$	-3.75	-7.51
	Rician	$\nu=4.44 \times 10^{-3}, s=8.75 \times 10^{-4}$	-3.73	-7.47
	Weibull	$\eta=3.90 \times 10^{-3}, \beta=3.90$	-3.66	-7.32
	GMM	$\mu_1=4.70 \times 10^{-3}, \sigma_1=3.29 \times 10^{-8}, w_1=0.60$ $\mu_2=5.60 \times 10^{-3}, \sigma_2=1.53 \times 10^{-8}, w_1=0.09$ $\mu_3=4.10 \times 10^{-3}, \sigma_1=3.15 \times 10^{-8}, w_1=0.31$	-3.86	-7.72

Table 3.3: Negative log likelihood (NLL) scores and BIC scores for different road surfaces.

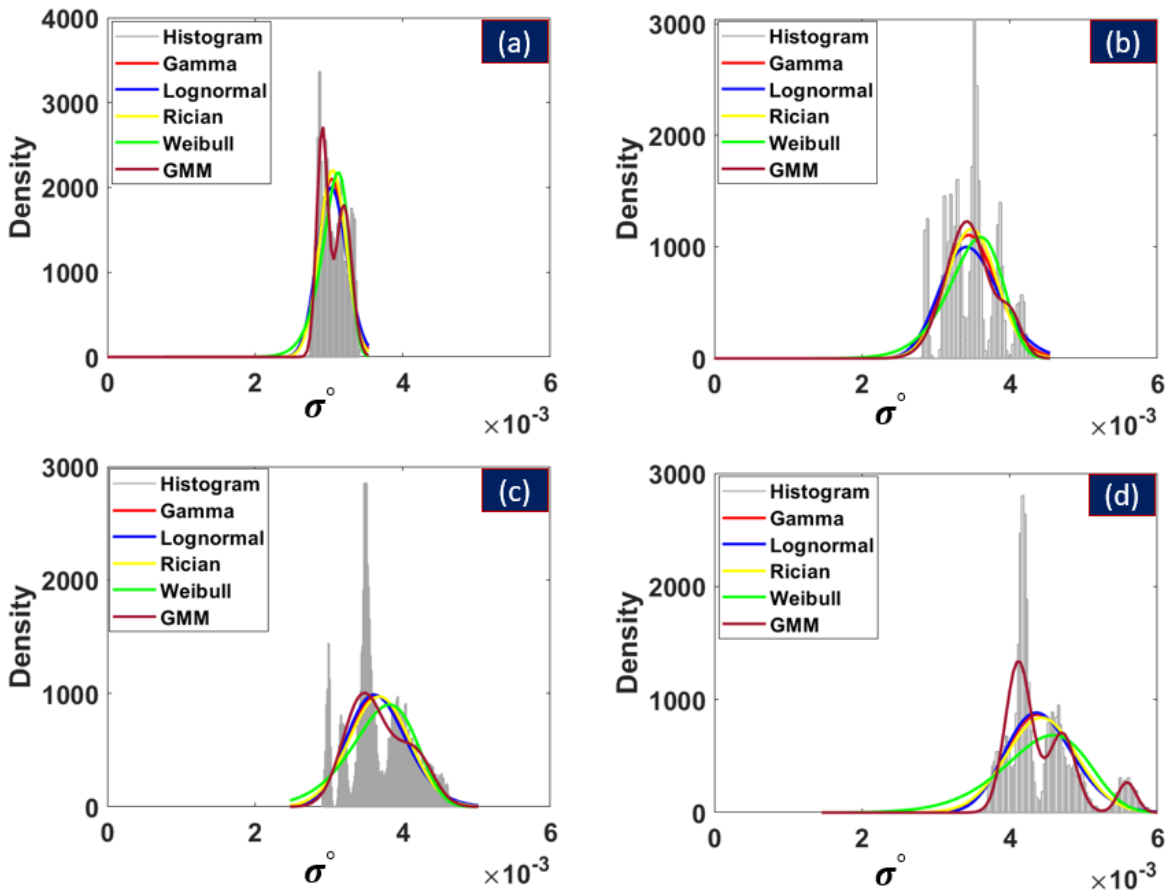


Figure 3.4: Distribution plots for the  $\sigma_0$  with respect to different type of road location (a) location 1 concrete road with roughness 6.4 mm, (b) location 2 concrete road with roughness 7.7 mm, (c) location 3 concrete road with roughness 8.7 mm (d) location 4 asphalt road with roughness 6.7 mm.

motivates a GMM at the dataset level. In the simulator (Chapter 2, Section 2.1.3), a single-scene exponential (power) speckle model is used with  $\sigma_0$  chosen from the measured set corresponding to the selected road type.

### 3.3 Summary of Findings

In summary, the main findings of this chapter are as follows:

- Paved Indian roads are typically made of concrete or asphalt, and the experimental measurements show that these roads are significantly rough




Figure 3.5: Plots of mean  $\sigma^0$  with respect to different radar heights with three different temperatures (a) location 1 concrete road with roughness of 6.4 mm, (b) location 2 concrete road with roughness of 7.7 mm, (c) location 3 concrete road with roughness 8.7 mm (d) location 4 asphalt road with roughness 6.7 mm.

according to international standards.

- Using the 77 GHz automotive radar, the mean values of  $\sigma_0$  for the concrete and asphalt roads are estimated to be  $-25$  dB and  $-23.6$  dB, respectively. The standard deviation of the  $\sigma_0$  values increases with grazing angle and surface roughness.
- Five commonly used probability distribution functions (lognormal, Gamma, Rician, Weibull, and Gaussian mixture model (GMM)) are fitted to the measured  $\sigma_0$  data, and their goodness-of-fit is assessed using the NLL and BIC scores.
- Based on these scores, the GMM provides the best fit to the  $\sigma_0$  values and captures the change in clutter characteristics with temperature, roughness, and grazing angle. For each road location, a mixture with  $B = 3$  Gaussian components is used, and the corresponding component means, standard deviations, and mixture weights are reported in Table 3.3.

## Chapter 4

# Generating Counterfactual Explanations For Misclassification of Automotive Radar Targets

*Building upon the classification framework established in the previous chapter, this chapter focuses on enhancing the interpretability and transparency of radar-based decision-making processes in ADAS through explainable AI techniques. The objective is to address the challenge of misclassification in automotive radar images by leveraging counterfactual explanations derived from generative adversarial networks (GANs). Our approach entails perturbing ISAR images belonging to a specific vehicle class ( $T_1$ ) to generate counterfactual images classified as an alternative class ( $T_2$ ). These perturbations are meticulously designed to maintain fidelity to the distribution of ( $T_1$ ) ISAR images, thereby ensuring realism while providing insightful*

*physics-based explanations for misclassification events. By enhancing the interpretability of radar-based decisions, this methodology supports regulatory considerations and fosters trust among stakeholders in ADAS deployment scenarios. This chapter presents a comprehensive exploration of our methodology, demonstrating its efficacy in refining radar-based target classification for critical automotive applications.*

*The specific contributions of this chapter are as follows:*

- **Counterfactual framework:** *a generative AI approach that introduces realistic, minimal perturbations to ISAR images of a vehicular class  $T_1$  to induce misclassification to  $T_2$ , mirroring physics-based shadowing/occlusions of specific parts of  $T_1$ .*
- **Realism and diversity:** *training objectives and architectural choices that keep counterfactuals on-manifold (realistic) while mitigating modal collapse for diverse generations.*
- **Physics-aligned explanations:** *a matching step that relates learned counterfactual perturbations to simulated part-occlusions, yielding actionable causes of misclassification.*

*This chapter is organized as follows. Section 4.1 reviews related work. The subsequent section describes the GAN-based architecture and the evaluation metrics. In Section 4.3, we detail the experimental setup for generating ISAR images from high-fidelity simulations, and in Section 4.5, we present the experimental outcomes and the physics-*

*based insights into the causes for misclassification. We conclude the chapter with a discussion of future work in Section 4.6.*

## **4.1 Related work**

ML algorithms have been crucial in processing information in various vehicular radar applications. For example, [104] demonstrates detection of parked vehicles using RF classifiers based on occupancy grids built on radar data. Dynamic targets like pedestrians and automobiles were effectively classified using micro-Doppler spectrograms [27]. Deep learning techniques were utilized to classify diverse targets from radar range-velocity spectra in [105], while [106, 107] contributed a high-resolution radar dataset for object classification. Convolutional neural networks (CNNs) were successfully deployed in [108] to extract semantic information regarding static and dynamic road targets from radar occupancy grid maps and point cloud data. Our prior work leveraged ISAR images to distinguish different vehicles using deep neural networks for a variety of applications [41, 91, 109].

However, for the broad commercial adoption of artificial intelligence (AI) techniques in radar, particularly within the automotive context, the explainability of AI-based models is a critical consideration. The term *explainability* denotes the process of rendering a classifier’s decision-making understandable to humans. Several motivations drive the development of explainable AI models: First, the inherent opacity of deep learning classifiers raises legal implications since

misclassification can lead to severe consequences. Therefore, explaining the reasons for classifications and, more importantly, misclassification facilitates responsibility to the correct actor. Second, the widespread adoption of AI-based decisions faces resistance unless there is trust that these decision-making processes are held accountable for fairness. Finally, enhancing the interpretability of machine-made decisions enables detecting and rectifying errors or bugs in models, contributing to improved classification performance.

Recent efforts have developed different approaches for explainable AI algorithms [65–67]. One approach involves quantifying the importance of different data features contributing to a model’s overall prediction performance [68, 69]. Another strategy creates simplified interpretable models that explain the performance of opaque deep neural network classifiers [70, 71]. Some methods aim to generate logical or mathematical explanations for AI models [66]. In [110], the authors proposed a black-box approach that generates importance maps to explain the predictions of deep neural networks by empirically estimating pixel importance through randomized masking. Building on this strategy, the authors later proposed a method that explains both localization and classification aspects of object detection [111]. In [112], the authors explored Bayesian strategies for target detection from SAR images and introduced an occlusion-based explanation method for highlighting the target scattering features contributing to model uncertainty. While the above list is not a comprehensive overview of explainable AI algorithms, these approaches broadly aim to reveal how the internal state of neural networks leads to specific decision outcomes. Contrastingly,

*counterfactual explanations* provide an alternative strategy for understanding a model’s classification behavior [72]. Rather than explaining the internal workings of the network, counterfactuals focus on *what-if* scenarios that illustrate how small changes or disturbances to the input data can shift the model’s prediction. This approach is particularly relevant in radar-based classification, where small distortions can lead to unexpected misclassifications.

- Feature/importance maps and masking-based saliency [110, 111].
- Surrogate/interpretable models for black-box explanations [66, 70].
- Uncertainty/occlusion analyses in radar [67, 112].
- *This chapter*: counterfactual explanations via GAN-based minimal, realistic perturbations aligned with physics-based occlusions.

## 4.2 Methodology

In this work on automotive radar, counterfactual explanations are used to address whether the occlusion of some parts of the body of a vehicle (say, one of the wheels) will result in a vehicle being misclassified in a different target class. Specifically, we examine scenarios where selective occlusions (e.g., one of the wheels of a four-wheeled vehicle) could result in a misclassification into another vehicle class. For example, a potential counterfactual explanation would be: *If the radar scattering from the front wheel of a four-wheeler is occluded, then the classification would shift from a car to a three-wheel autorickshaw, highlighting*

*the dependency on detected wheel structures for accurate classification.* Through such counterfactual explanations, this study demonstrates how minor input changes impact the classifier's predictions, allowing us to understand better and ultimately improve the robustness of radar-based automotive target identification.

This approach thus focuses the scope of the research study to the specific causes of misclassification rather than a generalized study of the deterioration of classification performance as a result of noise and clutter that was offered in previous works such as [91]. To illustrate this distinction, let us assume that two automotive radar target classes are well separated in some hyper-plane - at high signal to noise ratio (SNR) - resulting in excellent classification metrics. When the SNR falls, we observe that the radar data from the two classes mix in the hyperplane, resulting in degraded classification. On the other hand, in this work, we explain how a specific and minimum perturbation could cause a data point from one class to shift into the region of a second class even under high SNR.

Our approach involves leveraging short-range ISAR radar images generated from the simulation methodology outlined in [91, 109]. However, we believe the ideas presented in this chapter can be extended to other types of radar signatures - such as micro-Doppler spectrograms or high-range resolution profiles - used for target classification. Simulation models are specifically selected over measurement data collection to enable the occlusion of specific parts of the vehicles by design. ISAR images of dynamic vehicles provide information regarding their shape, size, and trajectory. These attributes form robust features for effective target classification using deep learning networks [91]. However, the classifiers'

performance proved susceptible to noise and clutter. This chapter presents a generative AI framework developed to explain misclassifications within automotive radar data. The framework has undergone significant advancements, including modifications to standard GAN architecture to address modal collapse problems commonly encountered in generative AI. Additionally, we have integrated shadowing phenomena into the high-fidelity simulations that generate the input radar data for the networks. Furthermore, new metrics have been introduced to facilitate detailed studies of causes of misclassification.

The system diagram illustrating our approach is depicted in Fig.4.1.




Figure 4.1: System diagram of the proposed framework for explaining the misclassification of ISAR images belonging to the  $T_1$  class as  $T_2$ . It illustrates the flow of input ISAR images of the  $T_1$  class to an RWGAN, generating counterfactual images that are now classified as  $T_2$  class. Following this, EfficientNetB0 is used to extract low-dimensional features from both electromagnetically simulated images and neural network generated images. Similarity metrics are computed between both sets of images to obtain counterfactual explanations for the misclassification of  $T_1$  class images as  $T_2$  class images.

At the core of our framework lies a Residual Wasserstein Generative Adversarial Network (RWGAN) that is trained to minimally perturb ISAR images of  $T_1$  class vehicles (for example, *car*) such that they are misclassified as images of  $T_2$  class (say, *autorickshaw*). *Experimental pairs evaluated*: {*car*→*autorickshaw*, *autorickshaw*→*bicycle*, *car*→*bicycle*}. These perturbations aim to mirror occlusions of parts of vehicles due to environmental conditions encountered in real-world scenarios. The generated radar images must ensure two specific conditions: One, the images must be realistic - belong to the distribution of  $T_1$

class; Second, the images must now be misclassified as  $T_2$  class. We hereafter refer to this class of images as  $T_{12}$  or counterfactual images.

Next, we extract low-dimensional image features from the generated  $T_{12}$  images using EfficientNetB0, a widely adopted baseline known for its efficiency in image classification [113]. These features are compared with similar features extracted from ISAR images generated from electromagnetic simulations (using the same EM simulation framework as Chapter 2) by occluding specific vehicular parts of a  $T_1$  class vehicle. Specifically, we compute two similarity metrics - the cosine similarity and Euclidian distance - to help identify the simulated images with specific occlusions most similar to the generated counterfactual images through which we determine which occlusions are responsible for misclassifying the  $T_1$  image as  $T_2$ . This approach provides physics-based insights into the misclassification patterns. We explain each of these steps in detail in this section.

#### 4.2.1 GAN Architecture

The standard GAN, proposed by [114], for image synthesis, comprises two neural networks - a generator denoted by  $\mathbf{G}$  and a discriminator denoted by  $\mathbf{D}$ . The GAN aims to find an equilibrium between these two networks using the objective function

$$\begin{aligned} \min_{\mathbf{G}} \max_{\mathbf{D}} V(\mathbf{D}, \mathbf{G}) = & \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log \mathbf{D}(\mathbf{x})] \\ & + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - \mathbf{D}(\mathbf{G}(\mathbf{z})))]. \end{aligned} \quad (4.1)$$

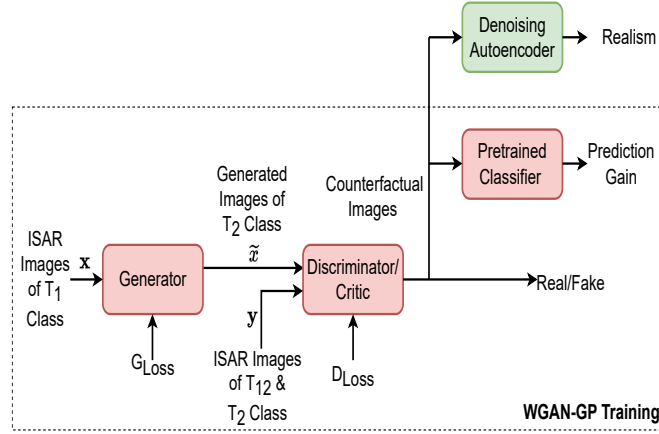


Figure 4.2: RWGAN Architecture for synthesizing counterfactual images that are realistic but misclassified.

Here,  $p_{\text{data}}$  and  $p_z$  signify the distributions of real data ( $\mathbf{x}$ ) and input noise ( $\mathbf{z}$ ). The objective is to minimize the difference between the distributions of real and fake/counterfactual samples while maximizing the discriminator’s ability to distinguish between them. The zero-sum framework of the problem helps train the network to create realistic fake samples.

Unlike the standard GAN that utilizes a noise vector  $\mathbf{z}$ , we consider a residual GAN where simulated ISAR images,  $\mathbf{x}$ , sampled from the distribution  $p_{T_1}$  are fed as input to the generator as shown in Fig. 4.2.  $p_{T_1}$  encompass various orientations and trajectories of  $T_1$  to the radar while assuming an unobstructed view of all vehicle parts. The objective is to train a residual generative framework to introduce minor perturbations to these images that mirror shadowing so that the counterfactual ISAR images,  $\tilde{\mathbf{x}}$ , will be classified as belonging to the second vehicular class  $T_2$ . These counterfactual images are provided as inputs to  $\mathbf{D}$ , along with  $\mathbf{y}$ , sampled from a combination of samples from two distributions:  $p_{T_{12}}$  and  $p_{T_2}$ , which are the respective distributions of  $T_{12}$  and  $T_2$  classes. Note that  $\mathbf{y}$  are generated from electromagnetic simulations and not neural networks

as  $\tilde{\mathbf{x}}$ . The primary goal of  $\mathbf{D}$  is to discriminate between  $\mathbf{y}$ , and  $\tilde{\mathbf{x}}$ .

A commonly encountered problem in the standard GAN in (4.1) is modal collapse, where the generator produces only a few samples/patterns from the target distribution, resulting in poor sample quality and low diversity in the generated data, as observed in [115]. Hence, we modify the standard GAN further to the RWGAN with the objective function as shown as

$$\begin{aligned} \mathcal{V}(\mathbf{D}, \mathbf{G}) = & \mathbb{E}_{\mathbf{y} \sim p_{\hat{T}_1}, p_{\hat{T}_2}} [\mathbf{D}(\mathbf{y})] - \mathbb{E}_{\mathbf{x} \sim p_{T_1}} [\mathbf{D}(\mathbf{G}(\mathbf{x}))] \\ & + \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}} [(\|\nabla_{\hat{\mathbf{x}}} \mathbf{D}(\hat{\mathbf{x}})\|_2 - 1)^2]. \end{aligned} \quad (4.2)$$

Here, the probability-based divergence in (4.1) is changed to the Wasserstein distance-based metric, which directly measures the *cost* of transforming one distribution into another.

The modified architecture results in smoother gradients and more stable training dynamics. In (4.2),  $\hat{\mathbf{x}}$  denotes perturbed samples generated by linearly interpolating real and generated samples by  $\hat{\mathbf{x}} = \epsilon \mathbf{y} + (1 - \epsilon) \tilde{\mathbf{x}}$ , where  $\epsilon$  is a random number sampled from a uniform distribution  $\epsilon \sim U(0, 1)$ . This leads to more stable and robust training. Also, the 1-Lipschitz constraint,  $\lambda \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}} [(\|\nabla_{\hat{\mathbf{x}}} \mathbf{D}(\hat{\mathbf{x}})\|_2 - 1)^2]$ , is enforced with a penalty term to ensure that the gradients of the discriminator concerning  $\hat{\mathbf{x}}$  have a norm close to 1. Here,  $\lambda$  controls the strength of the penalty and  $\|\nabla_{\hat{\mathbf{x}}} \mathbf{D}(\hat{\mathbf{x}})\|_2$  is the  $\ell_2$ -norm of the gradient of the critic  $\mathbf{D}$  to  $\hat{\mathbf{x}}$ .

To study the concept of misclassification, we introduce a third neural network,

a pre-trained classifier  $\mathbf{C}$ , to the existing architecture as shown in the Figure 4.2. This three-class classifier categorizes input into  $T_1$ ,  $T_2$ , or a third class representing neither  $T_1$  nor  $T_2$ . The incorporation of  $\mathbf{C}$  into the architecture leads to a modification of the objective function, as shown in

$$\begin{aligned} \mathcal{V}(\mathbf{D}, \mathbf{G}) = & \mathbb{E}_{\mathbf{y} \sim p_{\hat{T}_1}, p_{\hat{T}_2}} [\mathbf{D}(\mathbf{y})] - \mathbb{E}_{\mathbf{x} \sim p_{T_1}} [\mathbf{D}(\mathbf{G}(\mathbf{x}))] \\ & + \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}} [(\|\nabla_{\hat{\mathbf{x}}} \mathbf{D}(\hat{\mathbf{x}})\|_2 - 1)^2] + \\ & \mathbb{E}_{\mathbf{x} \sim p_{T_1}} [\mathbf{C}_{T_2}(\mathbf{x})] + \alpha \|\mathbf{G}(\mathbf{x})\|_1 + \beta \|\mathbf{G}(\mathbf{x})\|_2. \end{aligned} \quad (4.3)$$

Here,  $\mathbf{C}_{T_2}$  denotes the classification decision function for  $T_2$ . The role of  $\mathbf{G}$  is now constrained to generate  $\tilde{\mathbf{x}}$  that is not only realistic as it belongs to  $p_{T_1}$  (due to the penalty from discriminator) but also classified as  $T_2$ , by incorporating a penalty from  $\mathbf{C}$ . In other words,  $\tilde{\mathbf{x}}$  belongs to  $T_{12}$ . The last two terms in (4.3) are weighted regularization functions of  $\mathbf{G}(\mathbf{x})$ , controlling sparsity through the  $l_1$  norm and the amplitude of perturbations through the  $l_2$  norm. The hyperparameters  $\alpha$  and  $\beta$  are constants that require fine-tuning for optimal performance.

## 4.2.2 Evaluation metrics

We evaluate the performance of the model using the following metrics. The generator aims to synthesize images belonging to  $T_{12}$  class.

### 4.2.2.1 Prediction Gain

This metric quantifies the disparity between the classifier’s prediction on  $\tilde{x}$  which is generated by perturbing  $T_1$  class and the input image  $x$  as shown in (4.4).

$$\mathbb{E}[\mathbf{C}_{T_2}(\tilde{\mathbf{x}}) - \mathbf{C}_{T_2}(\mathbf{x})] \quad (4.4)$$

The score is bounded between 0 and 1. A higher score signifies an increased probability that the counterfactual is classified as  $T_2$  instead of  $T_1$ .

#### 4.2.2.2 Realism

This is a metric for gauging whether the synthesized counterfactual image,  $\tilde{x}$ , belongs to the distribution of  $T_1$  class. This ensures that the synthetic samples still belong to the class of ISAR images that could be generated from actual physics-based phenomena such as occlusion. The ‘‘Realism’’ score is a *distributional conformity* check for counterfactuals  $\tilde{x}$  with respect to the  $T_1$  image statistics used in this work; it does not constitute a physical guarantee. A common approach in generative AI literature for gauging a synthetic sample’s realism is using denoising autoencoders [116]. Therefore, we train a denoising autoencoder  $\mathbf{DAE}(\cdot)$  on  $p_{T_1}$ . Then we provide  $\tilde{x}$  as input to  $\mathbf{DAE}(\cdot)$  and obtain the output  $\tilde{\tilde{x}}$ . The error between the input and output of the autoencoder, as shown in (4.5), indicates whether or not the input was a realistic sample - with a low score indicating high realism.

$$\|\tilde{x} - \tilde{\tilde{x}}\|_2^2. \quad (4.5)$$

### 4.2.2.3 Actionability/Physics-Based Insights

To analyze how occlusions of specific vehicular parts of the  $T_1$  vehicle affect classification outcomes, we employed a systematic approach utilizing image similarity metrics. We first generate ISAR images of the  $T_1$  class through electromagnetic simulations while intentionally occluding specific vehicular parts. Then we resize both the simulated and generated images to uniform dimensions to ensure consistency in feature extraction. We use the EfficientNetB0 model for extracting low-dimensional feature vectors from both sets of images. This model is selected for its state-of-the-art performance in feature extraction for image classification applications.

We consider two distance metrics - the cosine similarity and Euclidean distance - to quantify the similarity between the both sets of images. The cosine similarity [117] is defined as:  $\frac{\mathbf{F}_1 \cdot \mathbf{F}_2}{\|\mathbf{F}_1\| \|\mathbf{F}_2\|}$ . where  $\mathbf{F}_1$  and  $\mathbf{F}_2$  are the features extracted from simulated and generated images respectively. Cosine similarity focuses on the angle between two vectors in a multi-dimensional space, gauging their orientation rather than their magnitudes. This makes the metric robust against variations in scale, which is valuable for identifying structural similarities. The Euclidean distance [118] is defined as:  $\sqrt{\sum_{i=1}^n (F_{1i} - F_{2i})^2}$ , where  $F_{1i}$  and  $F_{2i}$  are the  $i^{th}$  components of  $\mathbf{F}_1$  and  $\mathbf{F}_2$ , respectively, and  $n$  is the dimensionality of the vectors. Euclidean distance quantifies the absolute differences between feature vectors and captures the overall variability in feature space.

For each generated counterfactual image, we compute both metrics against all

stored images of the  $T_1$  class obtained from simulations. Then, we select the top three simulated images with greatest similarity to the counterfactual image and infer from the similarity scores the specific occlusions that caused the specific  $T_1$  class image to be misclassified as  $T_2$ .

By integrating these metrics, we achieve a holistic view of the similarities and differences, where cosine similarity reveals structural alignments, and Euclidean distance quantifies the magnitude of deviations. This ultimately provides counterfactual insights that enhance our understanding of misclassification. We compare counterfactuals to *electromagnetically simulated* partial-occlusion cases to infer which occlusions are decision-relevant; a curated, measured partial-occlusion dataset was not available in this study.

Subsequent sections will comprehensively discuss how the specific body parts are subjected to intentional occlusion and their collective impact on the misclassification.

### **4.3 Experimental Setup**

This section briefly overviews how ISAR images of automotive targets are simulated and subsequently utilized for synthesizing counterfactual explanations. The radar signal model utilized in this chapter is based on the detailed formulation presented in chapter 2, section [2.1.2](#).

The automotive radar parameters used for simulating ISAR images are detailed in Table [2.1](#). These parameters are chosen to align with commonly used

commercial mm-wave automotive radar sensors closely. The complete details of the simulation methodology can be found in 2.1 and are briefly recapitulated here.

For the purposes of simulations, we use computer-aided design models of three different target vehicles: a mid-size car with dimensions of  $4.4\text{m} \times 1.7\text{m} \times 1.5\text{m}$ , an autorickshaw with dimensions of  $3.5\text{m} \times 1.7\text{m} \times 1.8\text{m}$  and a bicycle with dimensions of  $1.5\text{m} \times 0.5\text{m} \times 1\text{m}$ . The chassis of each vehicle is rendered with multiple metallic triangular faces - car: 6905, autorickshaw: 6949, and bicycle: 3919, as illustrated in Fig.4.3. The RCS of a triangular facet is characterized analytically as a function of radar aspect angle, wavelength, and size. The aspect angle is computed from the dot product of the incident vector and the normal vector of the plate, causing the RCS of each plate to fluctuate along the target trajectory. We incorporate a Bernoulli random variable into the scattering coefficient to account for the fluctuations in the visibility of each scatterer center [60]. This low-complexity approach avoids the computational expense of ray-tracing algorithms, especially for large dynamic objects at mm-wave frequencies. The vehicles are animated along 16 distinct turning trajectories from the perspective of a north facing ego radar: right turns (south to east, west to south, north to west, east to north), left turns (east to south etc.), U-turns (south to south etc.) and straight paths (north to south etc.). The duration of the motion of the target along each trajectory is 5 seconds resulting in approximately 48 images per vehicle per trajectory. Utilizing the radar signal model and processing methodologies outlined earlier, we generate ISAR images for each of these 16

trajectories for a signal-to-noise ratio of 10dB. This results in approximately 750 images per target vehicle. Note that none of the vehicle’s parts are occluded from the radar in these images. These images collectively form the  $p_{T_1}$ , and  $p_{T_2}$  distributions for the deep learning architectures described in the previous section.

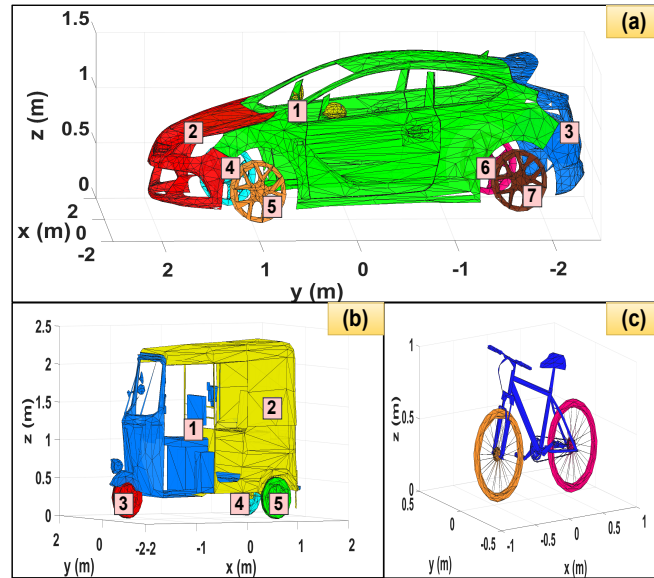


Figure 4.3: Computer-aided models of (a) car, (b) autorickshaw, and (c) bicycle, rendered with multiple triangular facets. The facets of the car are divided into seven parts, the autorickshaw into five parts, and the bicycle into three, as indicated by the different colors. Numeric labels (1–7 for the car; 1–5 for the autorickshaw) appear in (a) and (b). The bicycle in (c) is color-coded without numeric labels because it is not included in the physics-based occlusion dataset (see Table 4.1).

Next, we outline three scenarios for generating images with occluded body parts of targets. We divide the chassis of each vehicle, including all the facets, into distinct parts. The car is divided into 7 body parts consisting of each of the four wheels, the front, the rear, and the remaining body. Each part is shown with a different color in Fig.4.3. The autorickshaw is divided into 5 body parts - the three wheels, the front, and the rear. *The bicycle is excluded from the occlusion simulations in this chapter; it is shown in Fig. 4.3(c) for completeness and participates only as a class in the misclassification (counterfactual) experiments in section.4.4.* The main reason to focus on these components is to separate

parts of the vehicle into rigid components - the different sides of the vehicles - and the non-rigid components - the four wheels. The rigid components act like flat plate scatterers while the non-rigid components give rise to the additional micro-Doppler tracks that are characteristic of automotive targets.

First, we consider the occlusion of a single part of the vehicle by setting the scattering coefficients of the facets constituting the specific body part to zero. For the car, we simulate the occlusion of each of its seven body parts, resulting in 7 possible scenarios. This produces a total of  $7 \times 748$  occluded images. The same is repeated for the autorickshaw. Next, we consider the occlusion of any two body parts simultaneously. For the car, this yields  $\binom{7}{2} = 21$  combinations, and for the autorickshaw, it is  $\binom{5}{2} = 10$ . Finally, the exercise is repeated for a scenario where three body parts of a vehicle are occluded simultaneously for each vehicle (except for the bicycle). In summary, the total number of occluded images for the car and autorickshaw is determined by summing the results from all three scenarios as shown in Table.4.1. The images obtained through this

Table 4.1: Total Number of Occluded Images for Car and autorickshaw

Scenario	Total Occluded Images	
	Car	autorickshaw
Single-Part Occlusion	$7 \times 748$	$5 \times 748$
Double-Part Occlusion	$21 \times 748$	$10 \times 748$
Triple-Part Occlusion	$35 \times 748$	$10 \times 748$
<b>Total Classes <math>\times</math> Total Images</b>	$63 \times 47124$	$25 \times 18700$

exercise form the  $y$  class of data described in the previous section.

Finally, we use the deep learning classifier,  $C$ , trained on non-occluded data of car, autorickshaw, and bicycle classes to classify occluded images of  $y$ . While

several images are correctly labeled, we group those misclassified as  $T_2$  to  $p_{T_2}$  and use those for the RWGAN.

#### 4.4 Network Design and Hyperparameter Tuning

In this section, we discuss the design of the Generator (**G**) and Discriminator (**D**), along with two supporting networks: the classifier (**C**) and the denoising autoencoder (**DAE**). We focus on three distinct misclassification scenarios—(1) autorickshaw ( $T_1$ ) images misclassified as bicycle ( $T_2$ ); (2) misclassification of car ( $T_1$ ) as autorickshaw ( $T_2$ ); and (3) misclassification of car ( $T_1$ ) as bicycle ( $T_2$ ). For each scenario, we use the same architecture for **G** and **D**, with hyperparameter adjustments to optimize performance for each transformation. The architectures of **G** and **D** are shown in Figs. 4.4(a) and (b), respectively. During training, **G** and **D**, are trained to function adversarially, with the former aiming to generate counterfactual images indistinguishable from real images to the latter. The **G** and **D** networks consist of nineteen and seven layers, respectively. The input layer of the generator receives ISAR images of  $T_1$  with no occlusions of any part and pixel dimensions of  $168 \times 168$ . A cascade of convolutional layers (Conv2D) follows, with each layer comprising 64 nodes using a  $3 \times 3$  kernel size and  $2 \times 2$  stride. These convolutional layers play a crucial role in feature extraction and abstraction, gradually transforming the input data into higher-level representations. Batch normalization layers are applied after each convolutional layer to improve training stability and accelerate convergence, followed by Leaky ReLU activation functions to introduce non-linearity and enhance the model’s

representational capacity. Additionally, dropout layers are incorporated to prevent overfitting and improve the model's generalization ability. The architecture culminates with dense layers, which bridge the convolutional layers and the final output layer, and transposed convolutional layers that upsample the feature maps to generate the final output image.

The neural network,  $C$ , is designed as a three-class convolutional neural network to independently classify input images as  $T_1$ ,  $T_2$ , or  $T_3$  classes as depicted in Figs. 4.5. Here, the  $T_3$  label indicates images that do not belong to either  $T_1$  or  $T_2$ . The initial convolutional layer of  $C$  consists of 32 filters with a  $2 \times 2$  kernel size and uses the ReLU activation function. Subsequent convolutional layers follow a similar configuration, progressively reducing the spatial dimensions of the feature maps while increasing the depth to capture more abstract representations. Max-pooling operations are applied after each convolutional layer to downsample the feature maps, retain key information, reduce computational complexity, and prevent overfitting. After the convolutional layers, the output is flattened into a one-dimensional vector and passed through dense layers. The first dense layer has 256 nodes and uses the ReLU activation function. The output layer has three nodes corresponding to the three target classes ( $T_1$ ,  $T_2$ , and  $T_3$ ) and uses the softmax activation function. During training, the model parameters are iteratively refined using the Adam optimizer and the cross-entropy loss function, which measures the classification errors.

The objective of the denoising autoencoder (DAE) is to restore clean images from noisy or corrupted inputs. This network has been used in image classi-

fication [116] to identify if a test image belongs to a specific distribution of a class of images by computing the realism score of a generated synthetic image as described in Section. II.B. As shown in Figs. 4.6 the first layer accepts noisy or corrupted counterfactual images generated from the RWGAN. Then DAE proceeds with encoding through convolutional layers. Initially comprising 32 filters with  $3 \times 3$  kernel size and ReLU activation, these layers progressively reduce spatial dimensions while capturing increasingly abstract features. Max-pooling operations are used to downsample feature maps and dropout layers to mitigate overfitting. Following encoding, a compact latent space representation is formed. The subsequent decoder layers, consisting of transpose convolutional layers, aim to reconstruct clean images. Here, the layers upsample latent features to their original dimensions. The output layer, a single convolutional layer with a  $3 \times 3$  kernel and linear activation, completes the reconstruction process and yields noise-free images. During training, DAE is optimized using the mean squared error (MSE) loss function between pairs of noisy and clean images to measure reconstruction accuracy. The Nesterov-accelerated Adaptive Moment Estimation (NADAM) optimization algorithm [119], which combines the Adam optimization method and the concept of Nesterov Accelerated Gradient [120], iteratively adjusts model parameters to minimize the MSE loss, facilitating accurate image restoration. The performance metrics for different combinations of hyper-parameters are presented in 4.4.5, and 4.4.4.

#### 4.4.1 Autorickshaw to Bicycle Transformation

In this scenario, the objective is to minimally perturb the ISAR images of autorickshaws to generate realistic images that are misclassified as belonging to the bicycle class. The most effective architectures of **G** and **D**, for this scenario, are depicted in Figs. 4.4(a) and (b), respectively.

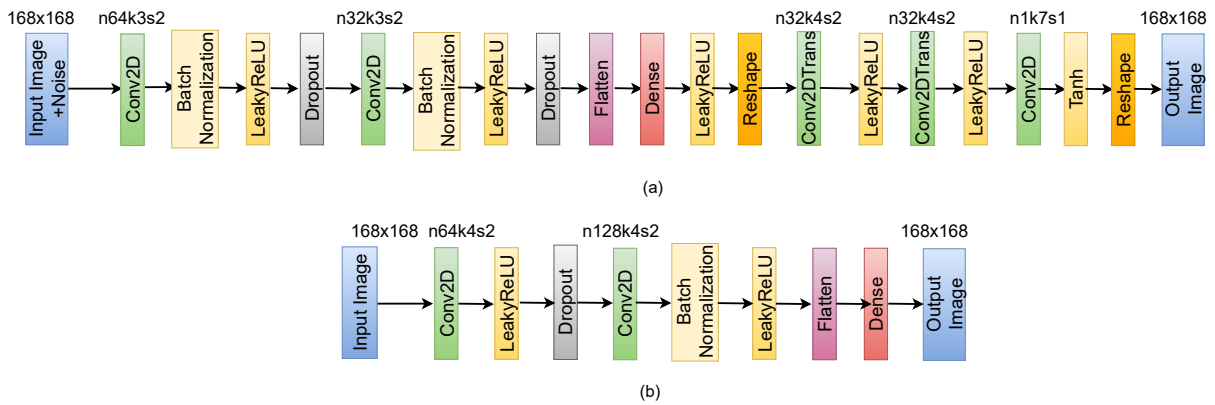


Figure 4.4: Architecture of the (a) generator and (b) discriminator network. Here,  $n$ ,  $k$ ,  $s$  are the number of filters, the kernel size, and the stride, respectively.

Through hyperparameter tuning and experimentation, the parameters of both networks are optimized to balance performance metrics such as image quality, computational time, and memory utilization. The  $l_1$  and  $l_2$  regularization terms ( $\alpha$  and  $\beta$ ) were fine-tuned to 0.0001 and 0, respectively. Convergence of the RWGAN was achieved within 1000 iterations, using a learning rate of 0.0002 and a gradient penalty of 15




Figure 4.5: Architecture of the classifier **C**. Here,  $n$ ,  $k$ ,  $s$  are the number of filters, the kernel size, and the stride, respectively.




Figure 4.6: Architecture of the (DAE). In this design,  $n$ ,  $k$ , and  $s$  denote the number of filters, the kernel size, and the stride, respectively.

#### 4.4.2 Car to Autorickshaw Transformation

The architectures of  $G$  and  $D$  for transforming ISAR images from the car to the autorickshaw class are depicted in Figs. 4.4(a) and (b), respectively. The fundamental structure with nineteen layers for  $G$  and seven layers for  $D$  remains intact, consistent with the previous case (autorickshaw to bicycle). Notably, hyperparameters have been adjusted to optimize performance for this transformation. The regularization parameters were calibrated to  $\alpha = 0.0001$  and  $\beta = 0.001$  in this scenario. The RWGAN successfully converged within 1000 iterations, with a learning rate of 0.0004 and a gradient penalty of 10.

#### 4.4.3 Car to Bicycle Transformation

Figures 4.4(a) and (b) depict the  $G$  and  $D$  networks used to transform ISAR images from car to bicycle. As with the previous transformation, the architectures of  $G$  and  $D$  remain unchanged, featuring the same nineteen-layer structure for  $G$  and seven-layer structure for  $D$ . Hyperparameters have been modified to enhance performance for this specific scenario. The regularization terms were also adjusted to  $\alpha = 0.0001$  and  $\beta = 0.001$ . The RWGAN achieved convergence within 1000 iterations, utilizing a learning rate of 0.0004 and a gradient penalty

set at 15.

#### 4.4.4 Classifier training

The C is trained on a dataset that includes images from three classes:  $T_1$ ,  $T_2$ , and  $T_3$  that comprises images that do not belong to either  $T_1$  or  $T_2$ . For training, we divided the dataset into 80% for training and 20% for a held-out test set to provide a robust evaluation of the classifier’s performance. We trained a CNN-based classifier for the different combinations of hidden layers and various numbers of filters used at each convolution layer. The classifier’s performance is reported using both training and test accuracy; training accuracy and loss (measured during training) verify convergence, while model selection is based on the test accuracy. In Table 4.2, we present the performance of each classifier combination. We selected the best-performing classifier with minimum complexity based on Table 4.2.

*Note on metrics:* Test accuracy is computed on the held-out test set in evaluation mode, whereas training accuracy is measured during training; consequently, once classes are well separated, the test accuracy can be slightly higher than the training accuracy.

#### 4.4.5 Auto-encoder training

We trained DAE on input class  $T_1$  data. We trained for different combinations of hidden layers, different numbers of filters used at each convolutions layer

Table 4.2: Classifier training results

Filter Layers	Batch size	Epoch	Accuracy(%)		Loss	
			Train	Test	Train	Test
Convolutional 64, 32, 16  Dense 256, 3	32	1	87.8	100	0.25	0.02
	32	3	99.3	100	0.02	0.004
	32	5	99.6	100	0.01	0.001
	32	10	99.9	100	0.01	0.001
	16	1	85.7	100	0.35	0.025
	16	3	99.6	100	0.01	0.001
	16	5	99.6	100	0.01	0.001
	16	10	99.8	100	0.01	0.001
Convolutional 128, 64, 32  Dense 256, 3	32	1	86.4	100	0.33	0.01
Convolutional 128, 64, 32 Dense 128, 3	32	1	86.4	99.4	0.42	0.56

for the encoder and decoder, and varying the number of epochs, as shown in Table 4.3. We evaluated the performance of the autoencoder based on the reconstruction error. We selected the best-performing autoencoder with minimum complexity based on Table 4.3.

Table 4.3: Autoencoder training

Filter #		Batch size	Epoch	Training Error	Test Error
Encoder	Decoder				
16, 8, 8	8, 8, 16	32	5	20±2.65	104±0.02
16, 8, 8	8, 8, 16	32	10	20±2.64	104±0.012
16, 8, 8	8, 8, 16	32	15	20±2.64	104±0.02
32, 16, 16	16, 16, 32	32	10	20±2.64	103±0.01
64, 32, 16	16, 32, 64	32	10	20±2.62	104±0.02
Number of hidden layers: 4					
32, 16, 8, 8	8, 8, 16, 32	32	10	20±2.62	104±0.01

## 4.5 Experimental Results and Discussion

This section discusses the experimental results explaining the misclassification of the ISAR images. We examine three scenarios: (1) autorickshaw to bicycle,

(2) car to autorickshaw, and (3) car to bicycle misclassification.

To show that the neural networks in the RWGAN are trained properly for all three scenarios, we present the change in prediction gain of the counterfactual images over training iterations in Fig.4.7a. This metric measures the extent to which the counterfactual images, generated by introducing perturbations to  $T_1$  class images, are classified as  $T_2$  class by the trained model. A higher prediction gain indicates a stronger likelihood that  $\tilde{x}$  is classified as  $T_2$  rather than  $T_1$ , thereby highlighting the model’s capacity to generate counterfactuals that resemble  $T_2$  class features. The result shows that as the training iterations increase, the prediction gain also increases. This means that the  $G$  is being trained such that some of the output of  $G$  will be classified as  $T_2$  rather than  $T_1$ . After 1000 iterations, we observe that transforming autorickshaw images into bicycles achieves the highest prediction gain of 0.56. Transforming car images into autorickshaws results in a prediction gain of 0.38, while transforming them into bicycles gives a prediction gain of only 0.18. These findings highlight that autorickshaw ISAR images are more likely to be misclassified as bicycles, even with minimal perturbation. However, achieving a similar level of misclassification of car images as bicycles is difficult, even with considerable perturbation. This is consistent with previous results in [91] where autorickshaw and car were more likely to be confused with each other than the car and bicycle.

Figure 4.7b shows the changing realism scores of the generated counterfactual images over the training iterations for all three scenarios. A lower realism metric suggests that the counterfactuals look realistic and fit well within the

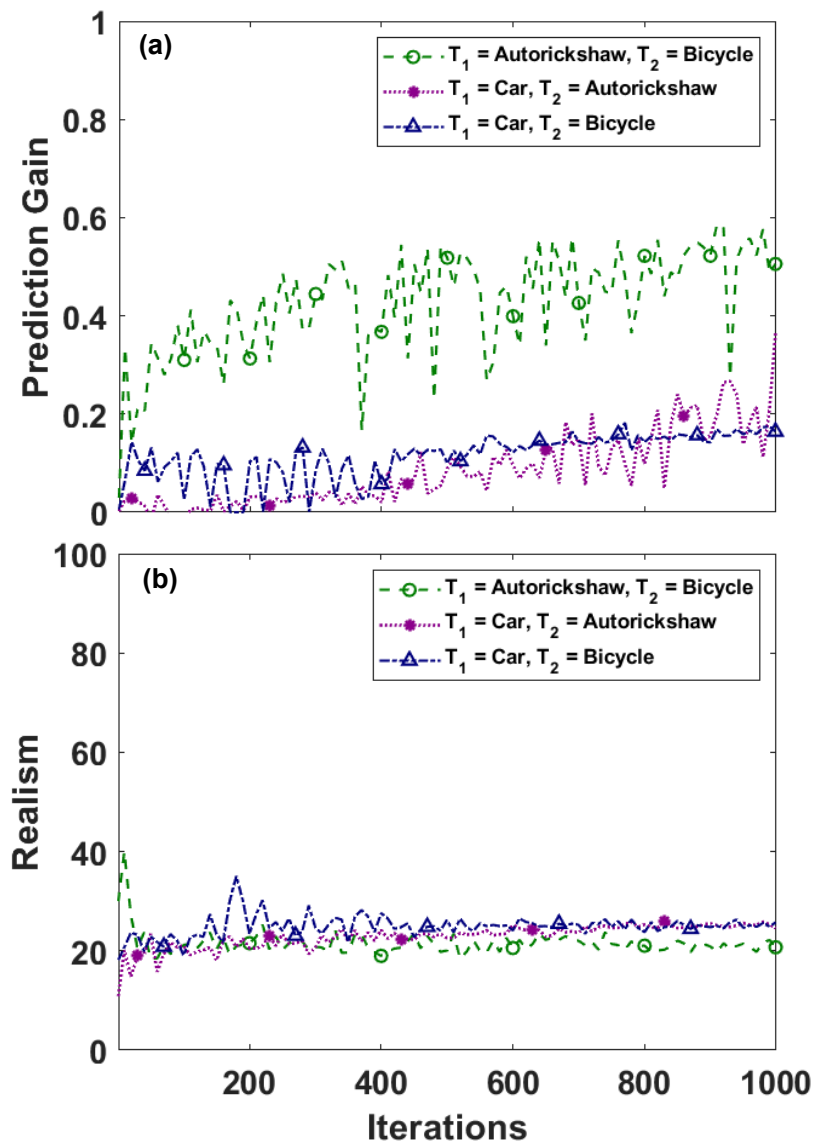


Figure 4.7: (a) Prediction gain, and (b) Realism score plots for three different cases, input class  $T_1$  to  $G$  and targeted (generated) class  $T_2$  for 1000 iterations.

distribution of the  $T_1$  class. In other words, the results from both these figures show that the generated counterfactual images belong to  $T_1$  distribution but are classified as belonging to the  $T_2$  class even after the perturbations are introduced. Throughout the 1000 training iterations for each scenario, the realism values consistently remain low. For instance, in the autorickshaw-to-bicycle scenario, the realism metric reaches around 21 after 1000 iterations. Similarly, in both the car-to-autorickshaw and car-to-bicycle scenarios, it stabilizes at approximately 25. These low values indicate that the counterfactuals effectively retain the characteristics of the original  $T_1$  class for all three scenarios.

#### 4.5.1 Misclassification of autorickshaw ISAR images as bicycle

We show 16 ISAR images (i-xvi) outlined in red of autorickshaw ( $T_1$ ) at different orientations in Fig.4.8(a).

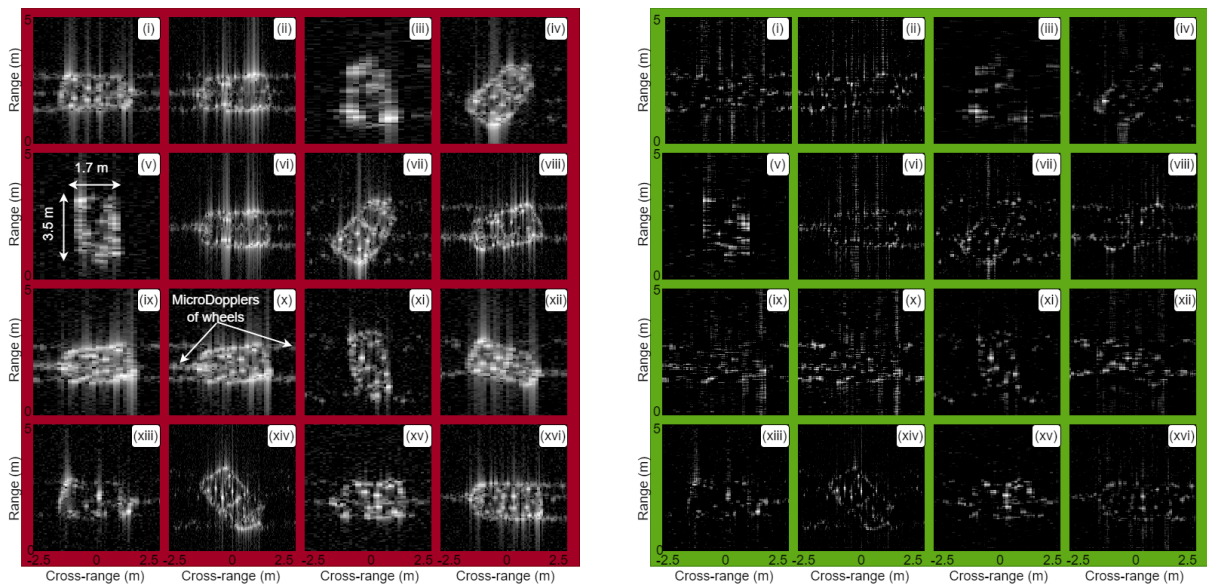


Figure 4.8: (a) Input ISAR images (i-xvi) represent ISAR images of  $T_1$  class vehicles (autorickshaw), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as  $T_2$  class vehicles (bicycle), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB.

These are generated from radar data at an SNR of 10dB. These images show the approximate shape, size, and orientation of the vehicle. Specifically, the ISAR images indicate that the autorickshaw appears as a triangle from the top view with dimensions of  $3.5\text{m} \times 1.7\text{m}$  and that each of the three wheels gives rise to a micro-Doppler track along the cross-range. These images are fed as input to the  $G$ . Figure 4.8 (b) shows the corresponding counterfactuals (i-xvi) outlined in green with minimal perturbation introduced by the RWGAN network. These images are misclassified as a bicycle ( $T_2$ ). Note that these images are distinct from a larger set of generated images that do not get classified as  $T_2$ . That set includes mostly images classified as  $T_1$  and fewer classified as  $T_3$  (neither  $T_1$  nor  $T_2$ ). The output counterfactual images differ from the input images since they indicate the additional perturbations introduced to each pixel of the input images. For improved visualization, Fig. 4.9 presents enlarged views of two representative input-counterfactual pairs from Fig. 4.8, where the subtle changes in the spatial extent of the chassis and in the number of micro-Doppler tracks can be seen more clearly.

In Section 4.3, we described generating 18700 ISAR images from the autorickshaw with parts of the chassis occluded from the radar. Specifically, we considered cases where each of the five different vehicular parts was occluded individually or in conjunction with other parts, resulting in 25 total distinct combinations, as described in Table 4.1.

To analyze the causes of misclassification, we utilize EfficientNetB0 to extract low-dimensional feature vectors of these simulated ISAR images and the

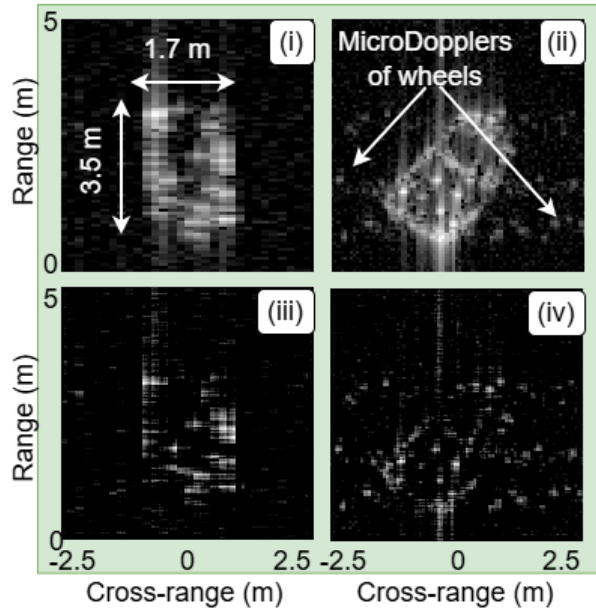


Figure 4.9: Enlarged view of representative input-counterfactual pairs for the autorickshaw-to-bicycle misclassification case. Images (i) and (ii) show input ISAR images of  $T_1$  class vehicles (autorickshaw) at two different orientations. Images (iii) and (iv) show the corresponding counterfactual ISAR images generated by the RWGAN and classified as  $T_2$  class vehicles (bicycle). The range and cross-range span are 5 m with a normalized dynamic range of 60 dB.

counterfactual images generated by RWGAN. Table 4.4 shows the highest three similarity scores for each counterfactual image (represented by *Image #* in the first column of the table) with corresponding three simulated ISAR images. For each of the identified simulated images, we list the occluded vehicular parts. For instance, in Fig.4.8 (b) (i), the misclassification of the autorickshaw ISAR image as a bicycle image can be attributed to the occlusion of the front chassis, rear chassis, and front wheel. In Fig.4.8 (b) (ii), the misclassification is primarily due to the occlusion of the right back wheel and left back wheel or due to the occlusion of the front chassis, front wheel, and right back wheel. Cosine similarity is particularly effective as it ranges from -1 to 1, where higher values indicate that the features of the counterfactual image share a similar direction to the simulated image in the high-dimensional space. This property makes cosine similarity

robust against variations in image scale, allowing us to focus on the structural relationships. In contrast, Euclidean distance quantifies the absolute differences between the feature vectors. This metric ranges from 0 to  $\infty$ , where lower scores near zero indicate similarity between counterfactual and simulated images.

From this table, we infer that the likeliest causes of misclassification can be summarized as follows: *When one or two wheels are occluded, the absence of the corresponding micro-Doppler tracks in the ISAR images results in misclassification to the bicycle since the latter has one less wheel than the autorickshaw. Similarly, when a part of the chassis is occluded, the smaller spatial extent of the overall ISAR image results in misclassification to the bicycle class.*

#### **4.5.2 Misclassification of car ISAR images as autorickshaw**

Next, we will discuss the car-to-autorickshaw misclassification scenario. Figure 4.10(a) shows 16 input ISAR images (i-xvi) of the car at various orientations, outlined in red.

The images show the rectangular shape of the top view of the car at different orientations with respect to the radar. Also, the dimensions of the length and width of the car can be estimated from the image. We also observe four distinct micro-Doppler tracks from the four wheels. The corresponding counterfactual ISAR images generated from RWGAN that are misclassified as the autorickshaw are shown in Fig.4.10(b), outlined in green. The results show that heavier occlusions are required for misclassifying car to autorickshaw when compared

Table 4.4: Counterfactual explanations for misclassification of autorickshaw as a bicycle.

Image #	Cosine Similarity	Euclidean Distance	Occluded parts
(i)	0.88	5.11	Front chassis, rear chassis, and front wheel
	0.87	5.12	Front wheel, right back wheel, and left back wheel
	0.86	5.23	Right back wheel and left back wheel
(ii)	0.87	5.35	Right back wheel and left back wheel
	0.87	5.57	Front chassis, front wheel, and right back wheel
	0.86	5.64	Rear chassis, right back wheel, and left back wheel
(iii)	0.93	3.95	Front wheel and left back wheel
	0.92	4.09	Front chassis, front wheel, and right back wheel
	0.92	4.15	Right back wheel and left back wheel
(iv)	0.92	4.43	Rear chassis, right back wheel, and left back wheel
	0.91	4.65	Front chassis, rear chassis, and right back wheel
	0.90	4.68	Front chassis and right back wheel
(v)	0.84	6.14	Rear chassis, right back wheel, and left back wheel
	0.83	6.36	Front chassis, right back wheel, and left back wheel
	0.82	6.48	Rear chassis
(vi)	0.82	7.05	Front chassis, front wheel, and right back wheel
	0.81	7.08	Rear chassis, front wheel, and right back wheel
	0.80	7.18	Front wheel, right back wheel, and left back wheel
(vii)	0.90	4.96	Rear chassis, right back wheel, and left back wheel
	0.89	4.98	Front chassis, front wheel, and right back wheel
	0.88	5.15	Rear chassis and left back wheel
(viii)	0.87	5.71	Front chassis
	0.87	5.76	Rear chassis, right back wheel, and left back wheel
	0.86	5.88	Rear chassis and right back wheel
(ix)	0.92	3.98	Front wheel and left back wheel
	0.91	4.28	Front chassis, right back wheel, and left back wheel
	0.90	4.38	Front chassis and right back wheel
(x)	0.88	4.91	Front chassis, front wheel, and right back wheel
	0.88	5.01	Right back wheel and left back wheel
	0.87	5.17	Front wheel and right back wheel
(xi)	0.91	4.70	Rear chassis, front wheel, and left back wheel
	0.90	4.72	Rear chassis, right back wheel, and left back wheel
	0.90	4.77	Rear chassis and front wheel
(xii)	0.90	4.41	Front chassis, front wheel, and right back wheel
	0.89	4.46	Front chassis and left back wheel
	0.89	4.57	Front wheel and left back wheel
(xiii)	0.88	5.44	Front chassis and left back wheel
	0.88	5.48	Front chassis, right back wheel, and left back wheel
	0.87	5.59	Front chassis, rear chassis, and right back wheel
(xiv)	0.91	5.40	Front chassis
	0.90	5.57	Front chassis and front wheel
	0.89	5.62	Rear chassis and front wheel
(xv)	0.87	5.10	Rear chassis, front wheel, and right back wheel
	0.86	5.52	Rear chassis and right back wheel
	0.85	5.70	Rear chassis, right back wheel, and left back wheel
(xvi)	0.87	5.35	Front chassis, front wheel, and right back wheel
	0.87	5.43	Front wheel and left back wheel
	0.86	5.53	Front wheel, right back wheel, and left back wheel

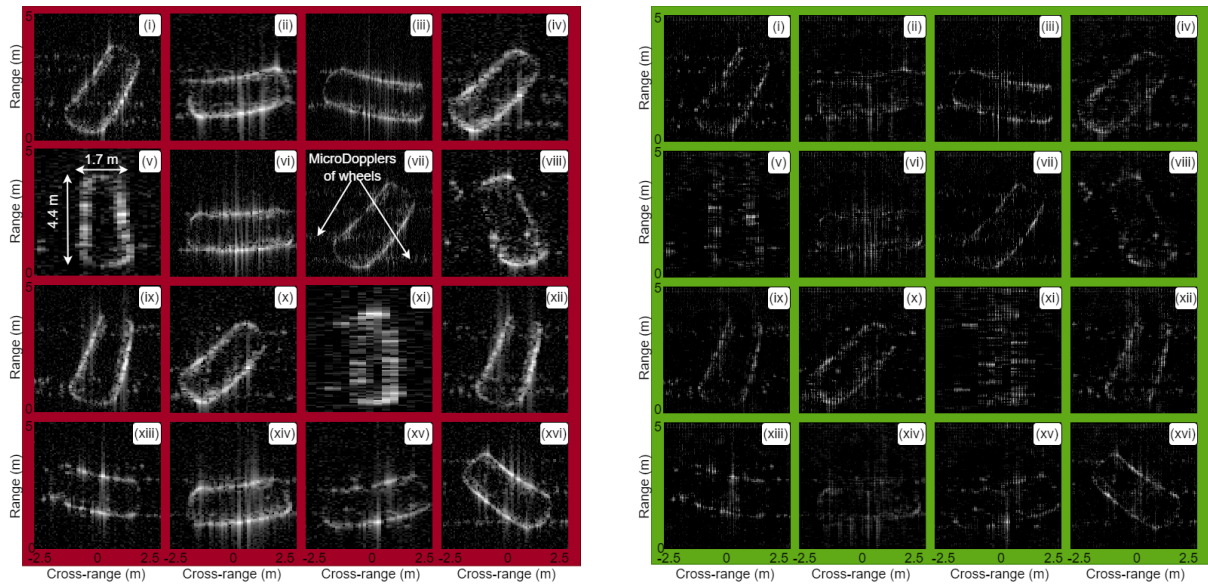


Figure 4.10: (a) Input ISAR images (i-xvi) represent ISAR images of  $T_1$  class vehicles (car), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as  $T_2$  class vehicles (autorickshaw), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB.

to the previous scenario of autorickshaw to bicycle.

Table 4.5 presents the counterfactual explanations for the misclassification of each image. The counterfactual image of Fig.4.10(b)(i), is most similar to the ISAR images with occlusion of the front chassis, right front wheel, and left back wheel or occlusions of the rear chassis, right front wheel. To better illustrate these occlusion patterns, Fig. 4.11 shows two representative car-to-autorickshaw examples at a larger scale, where images (i)–(ii) depict input ISAR images of the car  $T_1$ , and panels (iii)–(iv) show the corresponding counterfactual images misclassified as autorickshaw  $T_2$ . This zoomed view makes the localized shadowing of the chassis and wheels, which drives the misclassification, easier to discern.

This trend persists across all examples, as illustrated in Table 4.5. *Overall the similarity scores are poorer for this misclassification scenario than the*

Table 4.5: Counterfactual explanations for misclassification of a car as an autorickshaw.

Image #	Cosine Similarity	Euclidean Distance	Occluded parts
(i)	0.83	5.85	Front chassis, Right front wheel, and left back wheel
	0.82	5.91	Rear Chassis, right front wheel, and right back wheel
	0.86	6.01	Front chassis and left back wheel
(ii)	0.82	6.04	Front chassis, rear Chassis, and left back wheel
	0.82	6.09	Right front wheel, right back wheel, and left back wheel
	0.81	6.20	Right front wheel, right back wheel, and left back wheel
(iii)	0.78	7.71	Right front wheel and left back wheel
	0.78	7.88	Right back wheel, left front wheel, and left back wheel
	0.77	7.93	Front chassis and right back wheel
(iv)	0.81	5.78	Left front wheel
	0.81	5.79	Right back wheel
	0.80	5.83	Right front wheel and right back wheel
(v)	0.84	5.42	Front chassis, rear Chassis and left back wheel
	0.83	5.59	Rear Chassis, right back wheel and left back wheel
	0.82	5.65	Front chassis, rear Chassis and right back wheel
(vi)	0.79	6.85	Rear Chassis, left front wheel, and left back wheel
	0.78	6.90	Rear Chassis, right back wheel, and left back wheel
	0.78	6.93	Right front wheel, left front wheel, and left back wheel
(vii)	0.73	7.40	Rear Chassis
	0.72	7.42	Rear Chassis and left back wheel
	0.72	7.43	Rear Chassis and right back wheel
(viii)	0.85	5.25	Front chassis, right front wheel, and left back wheel
	0.83	5.51	Front chassis
	0.83	5.52	Front chassis and left front wheel
(ix)	0.82	6.18	Rear Chassis and right back wheel
	0.82	6.24	Rear Chassis, left front wheel, and left back wheel
	0.81	6.26	Rear Chassis, right front wheel, and left front wheel
(x)	0.84	5.39	Front chassis, right front wheel, and left back wheel
	0.84	5.46	Right front wheel, left front wheel, and left back wheel
	0.83	5.49	Front chassis, left front wheel, and left back wheel
(xi)	0.79	6.01	Front chassis, rear Chassis, and left back wheel
	0.79	6.10	Front chassis, rear Chassis, and right back wheel
	0.78	6.11	Front chassis and rear Chassis
(xii)	0.83	5.49	Right front wheel and left back wheel
	0.83	5.56	Front chassis and left front wheel
	0.82	5.59	Rear Chassis
(xiii)	0.84	6.18	Rear Chassis, right back wheel, and left front wheel
	0.82	6.30	Right front wheel, left front wheel, and left back wheel
	0.82	6.34	Rear Chassis, and right back wheel
(xiv)	0.81	6.22	Rear Chassis
	0.80	6.30	Right front wheel and left back wheel
	0.80	6.33	Rear Chassis, left front wheel, and left back wheel
(xv)	0.86	5.19	Front chassis, right front wheel, and left back wheel
	0.85	5.27	Right front wheel, right back wheel, and left front wheel
	0.84	5.35	Right back wheel, and left front wheel
(xvi)	0.82	6.28	Rear Chassis, right front wheel, and left front wheel
	0.81	6.40	Rear Chassis, left front wheel, and left back wheel
	0.81	6.41	Rear Chassis and right back wheel




Figure 4.11: Enlarged view of representative input-counterfactual pairs for the car-to-autorickshaw misclassification case. Images (i) and (ii) show input ISAR images of  $T_1$  class vehicles (car) at two different orientations. Images (iii) and (iv) show the corresponding counterfactual ISAR images generated by the RWGAN and classified as  $T_2$  class vehicles (autorickshaw). The range and cross-range span are 5 m with a normalized dynamic range of 60 dB.

*previous indicating that it is more likely that an autorickshaw is misclassified as a bicycle than a car to be misclassified as an autorickshaw. The results also show the shadowing of the front or rear chassis, or one or more wheels cause the misclassification.*

### 4.5.3 Misclassification of car as bicycle

Finally, we consider the car-to-bicycle scenario.

Figure 4.12(a) presents six input ISAR images of the car at various orientations, outlined in red, while Fig.4.12(b) shows the corresponding counterfactuals using RWGAN that are misclassified as bicycle class, outlined in green. There are some key differences in this scenario compared to the two previous scenarios: First, we are able to generate *much fewer* counterfactual images of cars that

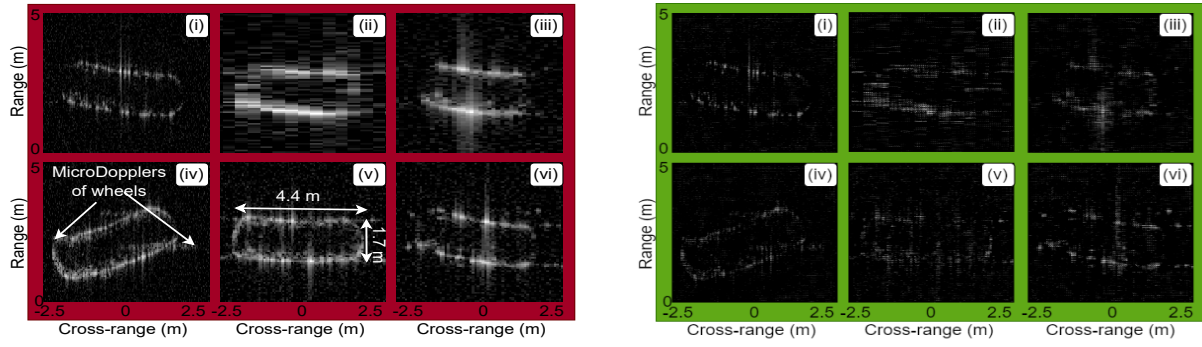


Figure 4.12: (a) Input ISAR images (i-xvi) represent ISAR images of  $T_1$  class vehicles (car), outlined in red. (b) Corresponding counterfactual images (i-xvi) depict ISAR images classified as  $T_2$  class vehicles (bicycle), outlined in green. The range and cross-range span are 5m with a normalized dynamic range of 60 dB.

are misclassified as bicycles. Second, considerable perturbation or shadowing is required for a car's ISAR images to be misclassified as a bicycle. This is consistent with our prior work, which reported more ambiguity between car and autorickshaw ISAR images than between car and bicycle due to the greater mismatch of the spatial dimensions and number of wheels between a car and a bicycle [91]. Table 4.6 lists the cases with the closest similarity between

Table 4.6: Counterfactual explanations for misclassification of a car as a bicycle.

Image #	Cosine Similarity	Euclidean Distance	Occluded parts
(i)	0.82	6.62	Front chassis, right back wheel, and left front wheel
	0.82	6.67	Rear Chassis, right front wheel, and left back wheel
	0.81	6.72	Front chassis, right front wheel, and left back wheel
(ii)	0.77	6.22	Front chassis, rear Chassis, and left back wheel
	0.77	6.31	Front chassis, and rear Chassis
	0.76	6.35	Rear Chassis, left front wheel, and left back wheel
(iii)	0.86	5.20	Rear Chassis, right back wheel, and left back wheel
	0.85	5.24	Rear Chassis, left front wheel, and left back wheel
	0.85	5.27	Right front wheel, left front wheel, and left back wheel
(iv)	0.84	5.91	Rear Chassis, right front wheel, and left front wheel
	0.83	6.01	Rear Chassis, left front wheel, and right back wheel
	0.83	6.10	Rear Chassis, right back wheel, and left back wheel
(v)	0.81	6.15	Rear Chassis, right back wheel, and left back wheel
	0.81	6.16	Rear Chassis, left front wheel, and left back wheel
	0.80	6.28	Rear Chassis, right front wheel, and left front wheel
(vi)	0.81	6.60	Front chassis and left front wheel
	0.80	6.77	Rear Chassis, left front wheel, and right back wheel
	0.79	6.79	Rear Chassis, right front wheel, and right back wheel

the counterfactual images and the corresponding simulated images. Here, we observe that *the car image is misclassified as a bicycle only when two or more vehicular parts and a wheel are occluded*. For instance, in image Fig.4.12(b) (iii), the rear chassis, right back wheel, and left back wheel are occluded.

#### 4.6 Summary of findings

- ISAR images are effective features for automotive target classification with deep neural networks. However, the classification performance deteriorates when some vehicle parts are occluded from the radar. This work *explains* how occluding parts of a target automotive vehicle can misclassify the corresponding ISAR image as another class, even under high SNR conditions. We address the problem using the deep learning-based RWGAN architecture. *The ISAR images of an un-occluded vehicle of  $T_1$  class are provided as input to the WGAN to generate counterfactual ISAR images that are realistic enough to belong to  $T_1$  class but are misclassified as  $T_2$  class. (Classifier:  $\sim 99$ – $100\%$  test accuracy on a held-out split; see Table 4.2.)*
- Based on the analysis of the counterfactual images, we infer that occluding any one or more wheels of a four-wheeled car or a part of its chassis can cause it to be misclassified as a three-wheeled autorickshaw. Similarly, the occlusion of any one or more wheels of an autorickshaw can cause it to be misclassified as a bicycle. However, two or more wheels and a significant part of the chassis must be occluded for a car to be mistaken for a bicycle.

- This is because the micro-Dopplers of the wheels of vehicles give rise to distinctive long cross-range tracks in the ISAR images, and the number of wheels can sometimes be inferred based on the number of such cross-tracks in the image (*car: four tracks; autorickshaw: three; bicycle: two*). Similarly, the rectangular spatial extent of a car in the top view ISAR image is greater than the triangular spatial extent of an autorickshaw and the almost point target-like spatial extent of a bicycle.
- Our approach demonstrates interpretability in radar-based classification by generating actionable counterfactual explanations that align occluded vehicular parts with potential misclassifications. This highlights how specific occlusion patterns influence classification outcomes, providing an interpretable framework for automotive radar systems. Although this study focused on three target classes, the insights can likely be extended to other automotive targets with varying structures, such as trucks, bikes, and pedestrians, where different occlusion patterns might affect classification similarly.
- In future work, we aim to explore how this framework can generalize to other target types by incorporating more varied vehicle structures and testing the influence of occlusions on additional automotive classes. Further research will also investigate whether this counterfactual explanation approach can address other phenomena impacting radar classification accuracy, such as clutter, multipath effects, and broader radar data representations, including joint time-frequency micro-Doppler signatures and radar point cloud data.

## Chapter 5

# Applications of ISAR Images of Automotive Targets

*In this chapter, we explore two advanced applications of ISAR imaging of automotive targets (1) The development of automated parking tests using ISAR images; and the (2) Implementation of around-the-corner sensing capabilities for enhanced situational awareness in urban environments.*

*We discuss both of these applications in greater detail in the following two sections.*

### **5.1 Automated parking test using ISAR images**

This section explores the development of automated parking tests using ISAR images generated by a single-channel externally mounted automotive radar system. The primary objective is to propose a novel approach leveraging ISAR images

for two distinct parking scenarios: angle parking and perpendicular parking. We advanced the field by introducing an automated parking test algorithm that utilized ISAR images to assess and verify the accuracy of parking maneuvers. By integrating radar-based imaging into parking assessments, this research enhanced parking efficiency and reliability in urban environments.

A key component towards improving road safety is to have well-trained and tested automobile drivers on the road. Recently, there has been research and development of automated driving license tests in order to address the costs, man-hours, and other types of inadequacies associated with manual driving tests, such as the tedium and corruption of the driving test inspectors. Driving tests typically involve testing the drivers in several key components such as lane changing, turning, following road rules, and parking. Currently, cameras have been the sensors of choice for automotive license tests [121]. In this work, we propose the use of mm-wave radars specifically for conducting automated parking tests. The main advantage of radar over other sensors, such as cameras, is that they can be used in low light/visibility conditions. Frequency-modulated continuous wave radars are currently heavily used in automotive systems for pedestrian detection, blind spot detection, object classification, and collision avoidance [11]. In our work, we propose that we use high-resolution radar images to test if the driver is following the designated test trajectory to the final parking spot.

Concurrent to automatic license testing development are the parking assistance systems on automobiles using cameras, lidars, and radars. These systems have been used to detect occupancy of parking slots [122, 123] and in aiding the

drivers in the correct positioning of the vehicle, especially while backing in [124]. The main distinction between the system that we propose to the existing parking assistance systems is that the radar is mounted *outside* the vehicle. Thus the objective is to *test/train* the driver as opposed to *assisting* the driver to complete the parking. Second, conventional automotive radars provide point cloud information of the target scatterers in the radar field of view. Each scatterer is marked with associated features of range, Doppler velocity, and signal strength. High-end systems consisting of multiple receiver channels provide additional azimuth and elevation information as well. In the proposed system, instead of generating point cloud data of the scatterers on the vehicle, we consider a low-cost and complexity, single-channel wideband radar that generates high-resolution range-crossrange images using ISAR imaging techniques. In chapter 2 we have shown that ISAR images provide rich information of size, orientation, and trajectory of dynamic targets [91, 109]

The main challenges of ISAR are twofold: First, the translational motion parameters of the automotive targets must be properly estimated in order to calibrate/compensate for their translational motion; second, the angular turning velocity must be accurately measured in order to obtain accurate cross-range estimates. In the parking test that we propose, we specifically address these challenges by determining if the parking is correct based on the estimated two-dimensional trajectory followed by the vehicle under test (VUT) using the ISAR images. We have experimentally validated our proposed parking test using measurement radar data gathered from the 77 GHz Texas Instruments AWR-

1843 automotive radar. Wide-band single-channel radar data gathered along the target trajectory are suitably processed to obtain ISAR images from which the trajectory of the target is estimated.

### 5.1.1 Methodology

The radar signal model used in this experiment builds upon the detailed formulation discussed in chapter 2, section 2.1.2. Here, we adapt the model to the specific conditions of this study: We assume that the VUT is moving on the  $XY$  ground plane with its height along the  $Z$  axis. The radar, mounted in a fixed position at the origin, operates at a short range in a line-of-sight environment relative to the VUT. The transmitted radar signal is a FMCW signal, as described in (2.5), with the same parameters for carrier frequency, chirp rate, and pulse repetition interval.

In this setup, the VUT is treated as an extended target with multiple dynamic scatterers. The radar return is influenced by both the translational and rotational motion of these scatterers, leading to Doppler shifts.

To process the received radar signals, we perform translational motion compensation within each  $T_{CPI}$  based on [40, 83]. The compensated signals are then transformed using a two-dimensional Fourier transform to produce range-Doppler plots. These plots are converted to range-cross-range plots, utilizing an angular velocity estimate derived from the yaw angle variations of the VUT. There are many ways to estimate the angular velocity,  $\omega$ , of a cooperative target

such as the use of additional sensors such as gyrometers and accelerometers. In this work, we estimate the  $\omega$  for every  $T_{CPI}$  from the change in the yaw,  $\theta$ , over consecutive intervals, where  $\theta[m]$  is

$$\theta[m] = \arctan \left( \frac{y_C[m] - y_C[m - 1]}{x_C[m] - x_C[m - 1]} \right). \quad (5.1)$$

Here  $(x_C, y_C)$  are the coordinates of the centroid of the VUT chassis with respect to the radar. We compute these coordinates based on the initial location of the VUT, the predefined trajectory, and the duration that the VUT takes to complete the trajectory.

**Rotation-rate estimation and smoothing.** The yaw sequence  $\theta[m]$  is computed from the centroid track  $(x_C[m], y_C[m])$  and smoothed with a quadratic (degree-2) polynomial fit to reduce numerical jitter on short slow-time windows. The CPI-wise angular velocity is then estimated as

$$\omega[p] = \frac{\theta[p] - \theta[p - 1]}{T_{CPI}}. \quad (5.2)$$

A quadratic fit was chosen as the minimum-capacity model adequate for the gentle curvature of the parking manoeuvres considered; higher-order fits are unnecessary here and may overfit noise.

**Parking Test Algorithm:** *Algorithm overview.* Algorithm 2 summarises the parking-test pipeline; the detailed description follows.

In the parking test, we infer whether the VUT has been appropriately parked based on the estimated trajectory of the VUT. In the *training algorithm*, we

---

**Algorithm 2:** Parking Test: training and test pipeline from ISAR images

---

**1. Training (for each trajectory, correct and incorrect):**

- (a) For each  $l = 1, \dots, L$ : convert the  $l^{\text{th}}$  image to grayscale; slide a fixed 2D bounding box; compute summed energy at each position.
- (b) Set  $(x_c^l, y_c^l) = \arg \max$  (summed energy) for each image.
- (c) Fit a 2D polynomial to  $\{(x_c^l, y_c^l)\}_{l=1}^L$ ; store the polynomial  $p_\tau$ .

**2. Testing (VUT sequence):**

- (a) Repeat the training steps on the test sequence to obtain  $\{(x_c^l, y_c^l)\}_{l=1}^L$  and fit  $p_{\text{test}}$ .
  - (b) Compare  $p_{\text{test}}$  to stored  $\{p_\tau\}$ ; decide *correct* / *incorrect* by minimum NMSE.
- 

perform the following steps: First, we collect  $L$  ISAR images of the car following a correct trajectory into the designated parking slot. Then we collect a similar set of ISAR images of the car following several incorrect trajectories resulting in parking outside of the designated parking slot. Then for each trajectory (both correct and incorrect), we convert the images to gray-scale. Then we choose a 2D bounding box comparable to the length and width of the car and slide the center of the box across the pixels of each  $l^{\text{th}}$  image while keeping the dimensions of the box fixed. For each sliding position of the bounding box, we compute the sum of the energy in the corresponding pixels bounded by the box. Then, we identify the 2D-pixel position,  $x_c^l, y_c^l$ , which has the maximum energy as the dominant scattering center position for that  $l^{\text{th}}$  image. We repeat these steps across all  $L$  ISAR images corresponding to the VUT motion. Then we curve fit a 2D polynomial across the dominant scatterer position to estimate the trajectory of the VUT across all the  $L$  images. We hypothesize that this two-dimensional polynomial function will correspond approximately to the trajectory of the center of the car. We repeat the exercise for the motion of the car along the incorrect

trajectories. These polynomial functions are stored and used while testing.

During the *test*, similar polynomials are estimated for the VUT motion and compared to the polynomials generated from the training. Then the VUT parking motion is deemed to be either correct or incorrect based on the closest fit of the test polynomial to the training polynomials. The advantages of this simple parking test are that the speed of VUT need not be identical to those of the car used for training. Hence, auxiliary sensors for estimating the translational motion characteristics of the test vehicles are not required. Further, the size of the VUT can differ a little from that used during training by adjusting the size of the bounding box.

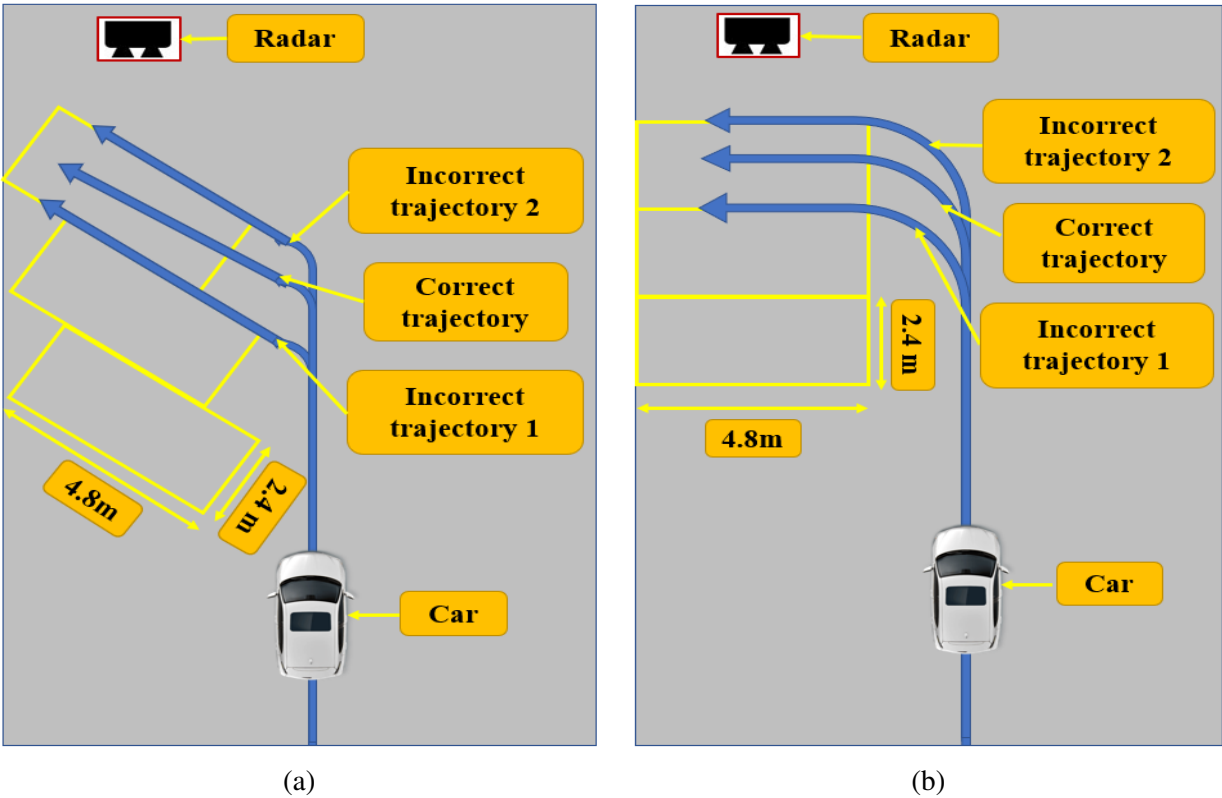


Figure 5.1: Three trajectories for the parking test for (a) 45° angle parking and (b) perpendicular parking.




Figure 5.2: The experiment setup for the parking test radar is situated at the location (0,0,0.5)m, and the car is parking at perpendicular parking.

### 5.1.2 Experimental Setup

We use the Texas Instruments AWR 1843, a 77GHz mm-wave radar, for experimental data collection. We configure the radar to transmit a frequency-modulated continuous wave signal with the parameters listed in Table 2.1 to operate as a short-range radar. The transmitted radar signal's pulse repetition interval ( $T_{PRI}$ ) is set to be 400  $\mu s$ , with a 133  $\mu s$  idle time and a 267  $\mu s$  ramp time. The chirp factor, or slope, is set to 7.532 MHz/ $\mu s$ , giving it a bandwidth of 2GHz, resulting in 7.5 cm range resolution. The sampling frequency is selected as 5MHz, resulting in 1328 fast-time samples in each  $T_{PRI}$ . A single coherent processing interval,  $T_{CPI}$ , of 0.1s duration is formed from slow time data of 250 chirps.

We have performed our experiments with two commonly used parking scenarios - the 45° angle parking and perpendicular parking as shown in Fig.5.1. The experimental setup for the parking test is shown in Fig.5.2, in which the

mm-wave radar is situated at the origin. The dimensions of the trajectory and parking lot for conducting the parking test are selected as per the standard defined by government agencies for parking [125]. The road along the parking trajectory is 7m in width. In the angle parking case, the car must first take a straight path, and then it should be parked at  $45^\circ$  from the road in a parking slot that is  $2.4 \times 4.8$  meters.

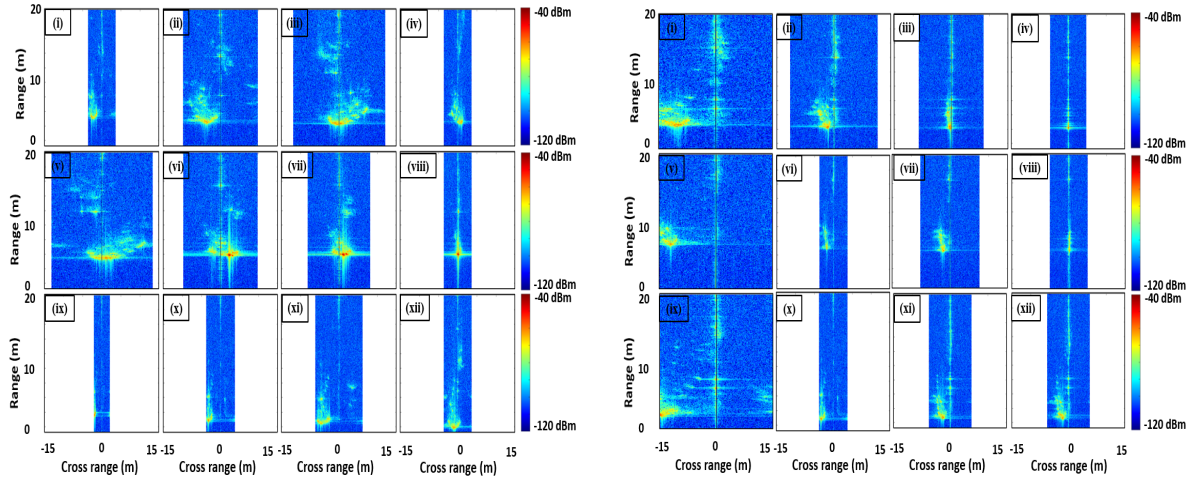


Figure 5.3: ISAR images of Ford Figo carrying out (a) perpendicular parking and (b) angle parking at 3.5, 4.0, 4.5, 5s. (1-iv) Top, (v-viii) middle, and (ix-xii) bottom rows in both images are generated for car following correct trajectory, incorrect trajectory-1, and incorrect trajectory-2 respectively.

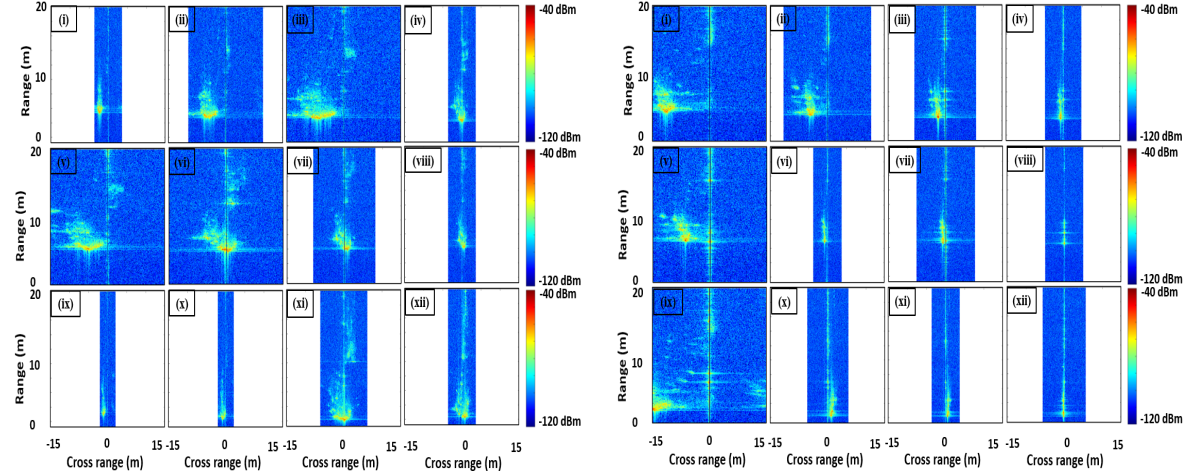


Figure 5.4: Bottom row shows ISAR images of Honda Brio carrying out (a) perpendicular parking and (b) angle parking, at 3.5, 4.0, 4.5, 5s. (i-iv) Top, (v-viii) middle, and (ix-xii) bottom rows in both images are generated for car following correct trajectory, incorrect trajectory-1, and incorrect trajectory-2 respectively.

In perpendicular parking, the car must take a straight path and then parked

at  $90^\circ$  from the road into a parking slot that is also of the same dimensions as the previous case. For our experimental data collection, we have chosen three trajectories for each type of parking test - one correct parking trajectory and two incorrect parking trajectories. We use two small-size cars for our experiment. The first car is a Ford Figo of  $3.9 \times 1.7 \times 2.5$  meters in size, and the second car is a Honda Brio of a comparable size of  $3.6 \times 1.7 \times 1.5$  meters.

### 5.1.3 Results

In this section, we will first discuss the ISAR images obtained while a car is being parked in the designated parking slot. We captured measurement data for 5 seconds for each parking test. We generate 50 ISAR images for each parking test using a  $T_{CPI}$  of 0.1 seconds. For evaluation, we use six test sequences from Honda Brio (three perpendicular and three angle) and six training sequences from Ford Figo (three perpendicular and three angle), each sequence being a 5 s run yielding  $L=50$  ISAR frames. In Fig 5.3a, we show the ISAR images for the perpendicular parking test for the Ford Figo. In each row, four images are shown corresponding to  $T_{CPI}$  equal to 3.5, 4.0, 4.5, and 5.0 s. The images from the straight-line motion of the car are not shown in these results. The first row corresponds to the ISAR images generated when the car is taking the correct trajectory and parked correctly. The second and third rows show the ISAR images for the case when the car is taking the two incorrect trajectories shown in Fig5.1(b) and (c); from these ISAR images, we observe the car's dimensions along with the range and cross-range. Also, in the first few view graphs (i, ii,v,

and vi), we observe that the car is oriented such that the longer dimension is along with the range while the shorter is along the cross-range. Then the car undergoes a turn such that the longer dimension is along the cross-range. Along the cross-range, we also observe the micro-Doppler tracks due to the wheels of the car. In all the figures, the cross-range axis varies because this axis depends on the target's rotational velocity  $\omega$ , which changes in every  $T_{CPI}$ . Our hypothesis is that based on the intensity of the ISAR image pixels corresponding to the car, we will be able to determine if the car followed the designated trajectory into the correct parking slot. Next, we show the results for the angle parking for the Ford Figo in Fig.5.3b for the same four-time instants. Again, we show the results for the correct parking (top row), and incorrect parking due to motion along the wrong trajectories (middle and bottom rows). Then, we repeated the measurements for the Honda Brio, and show the ISAR images for perpendicular parking in Fig.5.4a and angle parking in Fig.5.4b. The top rows for both figures correspond to the case when the car executed the correct parking, while the remaining two rows show the ISAR Images when the car performs incorrect parking. Just as in the previous case, the ISAR images give information about the car's size, position and its orientation along the range and cross-range axes.

In Table 5.1 we report the results of the parking test. The normalized mean square error (NMSE) between a test trajectory and each training trajectory is computed as the mean squared 2D coordinate error over all  $L=50$  frames, normalized by the test-trajectory energy (lower is better). We use the Ford Figo data for training and Honda Brio data for tests. The first three trajectories are

Table 5.1: Result of parking test algorithm using Ford Figo data for training and Honda Brio data for test.

True Trajectory	Predicted Trajectory					
	1	2	3	4	5	6
1	<b>0.332</b>	0.387	0.347	-	-	-
2	0.361	<b>0.347</b>	0.365	-	-	-
3	0.467	0.481	<b>0.466</b>	-	-	-
4	-	-	-	<b>0.512</b>	0.534	0.615
5	-	-	-	0.558	<b>0.512</b>	0.528
6	-	-	-	0.509	0.540	<b>0.411</b>

for the perpendicular parking case, while the remaining three are for the angle parking. Trajectories 1 and 4 correspond to the correct parking case, while the remaining correspond to incorrect parking. Table 5.1 shows the normalized mean square error (NMSE) between the estimated trajectory of the test case with each of the three training cases. The results indicate that the predicted trajectory from the algorithm is matched correctly to the ground truth in all cases based on the minimum NMSE. Thus the VUT has been correctly deemed to have either passed the parking test (cases 1 and 4) or failed the test (cases 2, 3, 5 and 6). **Note on model order.** Results reported in this section use the degree-2 smoothing for  $\theta[m]$  described in section 5.1.1. Within the manoeuvre range tested, this choice stabilises  $\omega$  without altering the qualitative outcomes; a broader sensitivity analysis to higher-order fits is outside the present scope.

## 5.2 Disaggregation and Detection of Multiple Automotive Targets for mm-wave Around-the-Corner Radar

**Overview.** We experimentally demonstrate that ACR detection of automotive targets with mm-wave (77GHz) short-range radars can be leveraged for look-

ahead sensing into NLOS road turns for collision avoidance. However, there is significant overlap in the range-Doppler returns from multiple extended targets in NLOS scenarios due to multipath-based ghost targets along the range and large micro-Doppler spreads from the vehicular wheels and arms and legs of pedestrians. As a result, it is challenging to resolve the returns from multiple targets before target detection and recognition. In this work, we propose the use of sparsity-based dictionary learning techniques to separate and reconstruct the radar signatures of each target from the aggregated radar returns of multiple targets. Our experimental results, with a 77GHz automotive radar with three target classes - car, pedestrian, and bike - show a mean successful detection of 88.3% after disaggregation.

**Background and related work.** In this chapter, we propose a novel application for short-range automotive radars for *look ahead sensing* in NLOS (blind) turns in ACR contexts for collision avoidance. Such scenarios may be encountered in urban environments where tall buildings occlude the driver's view (or the ego radar's line-of-sight) of the road into which an ego vehicle will turn. Note that detection based on measurements from other automotive sensors, such as cameras and lidars, is not viable in this scenario due to the NLOS conditions. Usually, only microwave radars have been used in NLOS modes, such as through-wall surveillance [126] and for ground and foliage penetration [127]. mm-wave radars, especially at W bands, have been restricted to LOS environments since the electromagnetic signal attenuates heavily through material mediums at these high frequencies. In this work, we demonstrate the viability of NLOS sensing

of targets in around-the-corner radar contexts using mm-wave radars. Over the last decade, *ACR* sensing has been researched primarily for detecting and localizing humans for surveillance purposes in [128–130] and with the help of reconfigurable intelligent surfaces in [131]. Experimental measurements have been conducted at microwave frequencies [130, 132] and at Ka bands [133] for primarily point targets or humans. Simulation studies have shown that three different types of propagation phenomenology occur when the radar is operated in the *ACR* scenario [134]. These are (1) direct through-wall propagation, (2) diffraction around the edges of the walls, and (3) multipath scattering off the lateral walls. At mm-wave frequencies, the dominant phenomenology at long distances is multipath scattering. The radar signatures most commonly extracted in *ACR* literature are the *high range resolution profiles* generated from wideband radar data from static targets from which higher-order multipath scattered returns are identified. Several types of detectors have been proposed for identifying a target by exploiting the multipath returns [133]. Alternatively, targets have been detected based on their micro-Doppler signatures gathered in *ACR* scenarios [134]. In addition to identifying these three mechanisms, [134] provides a frequency-dependent analysis of *ACR* phenomenology using FDTD simulations and measurements at 2.4, 5, 10, and 26 GHz. The study shows that *ACR* detectability and signature distortion are strongly dependent on carrier frequency and wall conductivity: at lower frequencies, the through-wall component dominates and the micro-Doppler signatures remain close to their free-space counterparts, whereas at higher frequencies and with lossy walls, the

NLOS signatures become weaker, more distorted by multipath, and observable only close to the corner. Based on these observations, [134] concludes that ACR detection strategies relying on multipath scattering should employ high carrier frequencies and must explicitly account for the nature of the wall. The automotive ACR experiment in this thesis extends these ideas to vehicular targets and to a higher carrier frequency of 77 GHz using a short-range mm-wave radar, with results reported for one representative wall configuration.

**Challenges.** The primary focus of our work is the disaggregation and subsequent detection of *multiple* dynamic automotive targets at W bands (77GHz) on the basis of their *range-Doppler ambiguity plots* in ACR scenarios. These plots are obtained by processing the two-dimensional fast time-slow time radar data gathered from the reflections of wideband linear frequency modulated waveforms transmitted by automotive radars [79]. The range-Doppler plots carry rich information on the size, shape, and motion trajectories of automotive targets and form excellent features for target identification [91, 109]. However, in short-range ACR scenarios, the detection and recognition of multiple extended targets are very challenging due to the significant overlap in their ambiguity plots. This overlap arises due to the ghost targets from multipath scattering along the range dimension and micro-Doppler spreads from the vehicular wheels and the arms and legs of the pedestrians along the Doppler dimension [135]. Separation of target features using traditional image processing techniques based on bounding boxes are largely ineffective since the range-Doppler spread of a stronger or nearer extended target often overwhelms a weaker target. The problem of

single-channel source separation, discussed above, has been studied in several application domains, including speech processing [136] and apportionment of electricity bills based on their energy consumption patterns [137]. Traditionally, source disaggregation was carried out by representing multiple sources using data-independent basis functions such as Fourier and wavelets. However, customized bases fine-tuned to the underlying data characteristics can often lead to sparser representations than data-independent transforms, eventually allowing more effective disaggregation. For example, in [138, 139], sparsity-based dictionary learning techniques were exploited to separate the Doppler spectrogram returns generated from narrowband time-domain radar data from multiple targets.

**Proposed approach.** In this work, we apply a similar formulation to the two-dimensional range-Doppler ambiguity maps generated from wideband FMCW radar data. Once the radar range-Doppler ambiguity maps are disaggregated, they are used for the detection and identification of each automotive target.

**Experimental validation and contributions.** In this work, we experimentally validate our proposed radar application with ACR measurement data collected from the TI AWR1843 sensor operating at 77GHz for three automotive targets - a small-size car, a bike, and a pedestrian. We collect multiple one-target and two-target class radar signatures along diverse target motion trajectories and generate their corresponding range-Doppler ambiguity plots. Then, we apply dictionary learning to achieve two objectives: first, to disaggregate returns from multiple targets; and second, to detect the presence of a target from the disaggregated returns of the corresponding target class. Note that this problem formulation

is fundamentally different from a *classification problem* where the objective is to classify data from a single target class based on trained classifiers. We use the single-target ambiguity maps to train dictionaries for each target class and use these dictionaries to disaggregate and identify these targets in multi-target scenarios. Our results show that the single target range-Doppler ambiguity maps are reconstructed with an error of 0.24 from the multiple target data and are correctly identified with an average accuracy of 88.3%. Our contributions can, therefore, be summarized as follows - First, we experimentally demonstrate a novel application of automotive radars for look-ahead sensing of dynamic targets in ACR scenarios for collision avoidance; Second, we show that W band mm-wave radars can be used in NLOS / ACR scenarios; Third, we present and experimentally validate a method for disaggregating and then detecting the radar returns from multiple target range-Doppler ambiguity signatures.

### 5.2.1 Methodology

The radar signal model used in this study is based on the detailed radar signal processing framework presented in chapter 2 section 2.1.2. We transmit an FMCW signal as described in (2.5). We consider an extended target with  $P$  point scatterers, each located at range  $r_p$  and moving with radial velocity  $v_p$ . The received signal, including delays and Doppler shifts due to target motion given by (2.6), we follow the same processing approach as described in section 2.1.2 to get range-Doppler ambiguity plot (2.14).

*Supervised Learning Technique for Target Disaggregation:* During training,

we derive data-driven dictionaries,  $\mathbf{D}_i \in \mathbb{C}^{MN \times Q}$ , from the radar range-Doppler ambiguity plots for each radar target class  $i$ . First, we vectorize each  $N \times M$  plot into a  $\mathbf{y}_i^{tr} \in \mathbb{C}^{NM \times 1}$  column vector corresponding to one CPI. We collect  $L$  such measurements to form a training matrix,  $\mathbf{Y}_i^{tr} \in \mathbb{C}^{NM \times L}$ , for each  $i^{th}$  target class. The objective is to represent the radar data using dictionaries,  $\mathbf{Y}_i^{tr} = \mathbf{D}_i \mathbf{Z}_i^{tr}$ , while enforcing a sparsity constraint on the corresponding coefficient matrix,  $\mathbf{Z}_i^{tr}$ . Hence the objective function is

$$\min_{\mathbf{D}_i, \mathbf{Z}_i^{tr}} \|\mathbf{Y}_i^{tr} - \mathbf{D}_i \mathbf{Z}_i^{tr}\|_2^2 \text{ s.t. } \|\mathbf{Z}_i^{tr}\|_0 \leq p_0. \quad (5.3)$$

Since the above formulation is NP-hard [140], we use compressed sensing principles to relax the  $l_0$  constraint to its  $l_1$ -norm counterpart as shown in

$$\min_{\mathbf{D}_i, \mathbf{Z}_i^{tr}} \|\mathbf{Y}_i^{tr} - \mathbf{D}_i \mathbf{Z}_i^{tr}\|_2^2 + \lambda \|\mathbf{Z}_i^{tr}\|_1, \quad (5.4)$$

where  $\lambda$  is the regularization parameter that trades off between the representation error and the sparsity constraint. The above problem can be solved using an iterative approach as discussed in [141]. We offer a brief overview of the algorithm here. In the first iteration, we initialize  $\mathbf{D}_i$  with random values from the training data. Then, we estimate the coefficient matrix using the ISTA algorithm [142] based on

$$\min_{\mathbf{Z}_i^{tr}} \|\mathbf{Y}_i^{tr} - \mathbf{D}_i \mathbf{Z}_i^{tr}\|_2^2 + \lambda \|\mathbf{Z}_i^{tr}\|_1. \quad (5.5)$$

Then  $\mathbf{D}_i$  is obtained from the new  $\mathbf{Z}_i$  using least squares on  $\|\mathbf{Y}_i^{tr} - \mathbf{D}_i \mathbf{Z}_i\|_2^2$ . The last two steps are reiterated till the objective function in (5.4) either asymptotically converges or falls below a previously stipulated threshold. These steps are repeated for all  $I$  radar target classes.

During the test, the dictionaries of all the classes are stacked together to form an aggregate dictionary  $\mathbf{D}_{agg} = [\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_I]$ . The radar ambiguity data during the test,  $\mathbf{y}^{test}$ , may consist of returns from one or more targets. The coefficient vector, corresponding to the test data is retrieved using the ISTA algorithm on

$$\min_{\mathbf{z}^{test}} \|\mathbf{y}_i^{test} - \mathbf{D}_{agg} \mathbf{z}^{test}\|_2^2 + \lambda \|\mathbf{z}^{test}\|_1. \quad (5.6)$$

Note that since  $\mathbf{D}_{agg} \in \mathbb{C}^{MN \times IQ}$ , is a collection of dictionaries from all target classes, the recovered coefficient vector is  $\mathbf{z}^{test} \in \mathbb{C}^{IQ \times 1}$ . Therefore, the coefficient vector corresponding to each  $i^{th}$  target class,  $\tilde{\mathbf{z}}_i$ , is extracted from  $\mathbf{z}^{test}$ . For example  $\tilde{\mathbf{z}}_1 = \mathbf{z}^{test}[1 \dots Q]$  and  $\tilde{\mathbf{z}}_2 = \mathbf{z}^{test}[Q + 1, \dots, 2Q]$  and so on. Then, the target returns from each target class are reconstructed using  $\tilde{\mathbf{y}}_i = \mathbf{D}_i \tilde{\mathbf{z}}_i$ . Based on the strength of the reconstructed signal, we determine whether or not the corresponding target exists in the channel. The algorithm described here is a su-

pervised learning technique where the possible target classes in the channel have been identified aprior. The presence of other target classes, besides, the possible list of targets, cannot be identified using this method. Also, the technique only identifies the presence or absence of a particular target class. It cannot be used to count the number of targets of a particular class.

### 5.2.2 Experimental Setup

We use TI AWR1843, a highly versatile mm-wave sensor, for collecting measurement data in ACR scenarios. We configure the sensor as a short-range radar that transmits frequency-modulated continuous wave signals at a center frequency of 77GHz. We set the pulse repetition interval of the signal to be  $400\mu s$  with an idle time of  $133\mu s$  and a ramp time of  $267\mu s$ . The chirp factor or slope of the signal is set as  $7.532\text{ MHz}/\mu s$  resulting in a bandwidth of 2GHz. The sampling frequency of the radar data is selected as 5MHz resulting in 1328 fast time samples. Slow-time data of 250 chirps are gathered to form a single CPI of 0.1s duration. The radar data over fast and slow time are processed with a two-dimensional (2D) Fourier transform to obtain range-Doppler ambiguity plots. The maximum unambiguous range is taken as 99.6m with a range resolution of 0.075m. Also, the maximum unambiguous Doppler is 2500 Hz with a Doppler resolution of 10 Hz. We deploy this radar in an ACR scenario, as shown in Fig.5.5. Two roads form a T junction along the  $XY$  plane, as shown in Fig.5.5. The main road (aligned along the  $X$  axis) along which the targets move on is 8.2m wide, while the second road (aligned along the  $Y$  axis), on which the radar is placed,



car which is  $4m \times 1.6m \times 1.5m$  (Fig.5.6c). Note that all three targets are lower than the height of the walls and hence not visible to the radar along the elevation. In our measurement scenario, the walls nearer to side A-B are more uneven than those alongside C-D. We collect experimental data of single target motions (Fig.5.6a and b) as well as two target motions (Fig.5.6c). In the case of the two target motions, the trajectories are chosen such that one target does not shadow the other from the radar. A total of 1200 single target measurement data samples are collected, of which 1050 are used for training and the remaining 150 are used for validation. 300 measurement data samples are collected from the two target scenarios, and all of them are used for validation.

### 5.2.3 Measurement Results

In this section, we present the results of disaggregation and detection of automotive targets in ACR scenarios.

*Scope of the reported rate.* The mean successful detection rate of 88.3% reported in this section pertains to the specific site, wall configuration, and trajectories described in Section 5.2.2. No material sweep or systematic study of wall thickness and surface roughness is performed here, and the numerical value should therefore be interpreted as a proof-of-concept for this configuration rather than as a universal performance guarantee for all urban ACR scenarios.

*Range-Doppler ambiguity plots of automotive targets in LOS/ACR scenarios:* The top three rows in Fig.5.7 show the range-Doppler plots for a small size

car, bike and human as each moves along the trajectory (iii) shown in Fig.5.5 with 12 kmph, 9kmph and 8kmph speed respectively. It is evident from these results that radar spotting of targets is possible in NLOS ACR conditions due to multipath scattering of the mm-waves off the lateral walls. Further, their apparent range peaks are displaced from the actual range. For example, the car is located between A and B in Fig.5.5. The actual range extent of the car is from 15.2m to 18.5m. However, the apparent ranges in the plots are from 18m to 21.5m. Multiple peaks are observed along the range dimension corresponding to second-order reflections off the front of the car and reflections from different parts of the car. The bike and pedestrian were each located between A and B. Since they are spatially narrow targets, they appear as point scatterers along the range. The actual ranges for both are 11.7m and 12.2m, respectively, while they appear at 15m and 18m due to multipath off the lateral walls. In all of the results, we observe strong returns at DC Doppler due to static clutter. We also observe micro-Doppler modulations, along the Doppler dimension, arising due to the rotational motion of the wheels in the case of the car and bike and the swinging motion of the arms and legs of the human. The signal returns from the pedestrian are much weaker than those of the car since the pedestrian is characterized by a much weaker radar cross-section than the car [60]. The range-Doppler ambiguity plots for the two-target scenarios in both LOS and ACR conditions are shown in the last three rows of Fig.5.7. The fourth row shows the plots for a car moving along the trajectory (ii) while a bike moves along the trajectory (iv). In the first two columns of the fourth row, both car and bike are in the LOS; hence

the returns are strong. In the last two columns of the fourth row, car and bike start moving away from the radar and come into the ACR region. Thus, we can observe the Dopplers for both the targets are increasing. The fifth row has plots of the car along the trajectory (iv) and human along the trajectory (i). In the first two columns of the fifth row, a car is in the LOS, and a human is in the ACR region, while in the last two columns of the fifth row, a car is in the ACR, and a human is in the LOS. The last row shows the range-Doppler plots where the human and bike are moving along trajectories (ii) and (iii), respectively. In the first two columns of the last row, the human is in the LOS, while the bike is in the ACR, and in the last two columns of the bottom row, the human moves to the ACR, and the bike comes in the LOS. It is evident in the figure that due to the multipath and micro-Dopplers from both targets, there is considerable overlap in the radar returns along both the range and Doppler dimensions.

***Disaggregation of multiple target radar data:*** We disaggregate the radar returns from multiple targets and reconstruct the range-Doppler plots of every constituent target. Then we validate the numerical accuracy of the reconstruction process by computing the normalized mean square error (NMSE) between the reconstructed plots and ground truth single target plots. In this sub-section, we synthesize the two target data by adding data from two single-target cases so as to have ground-truth data for validation purposes since it is not possible to exactly replicate single target motions again in two-target scenarios. First, we consider a car and a bike following the red and blue trajectories, respectively,

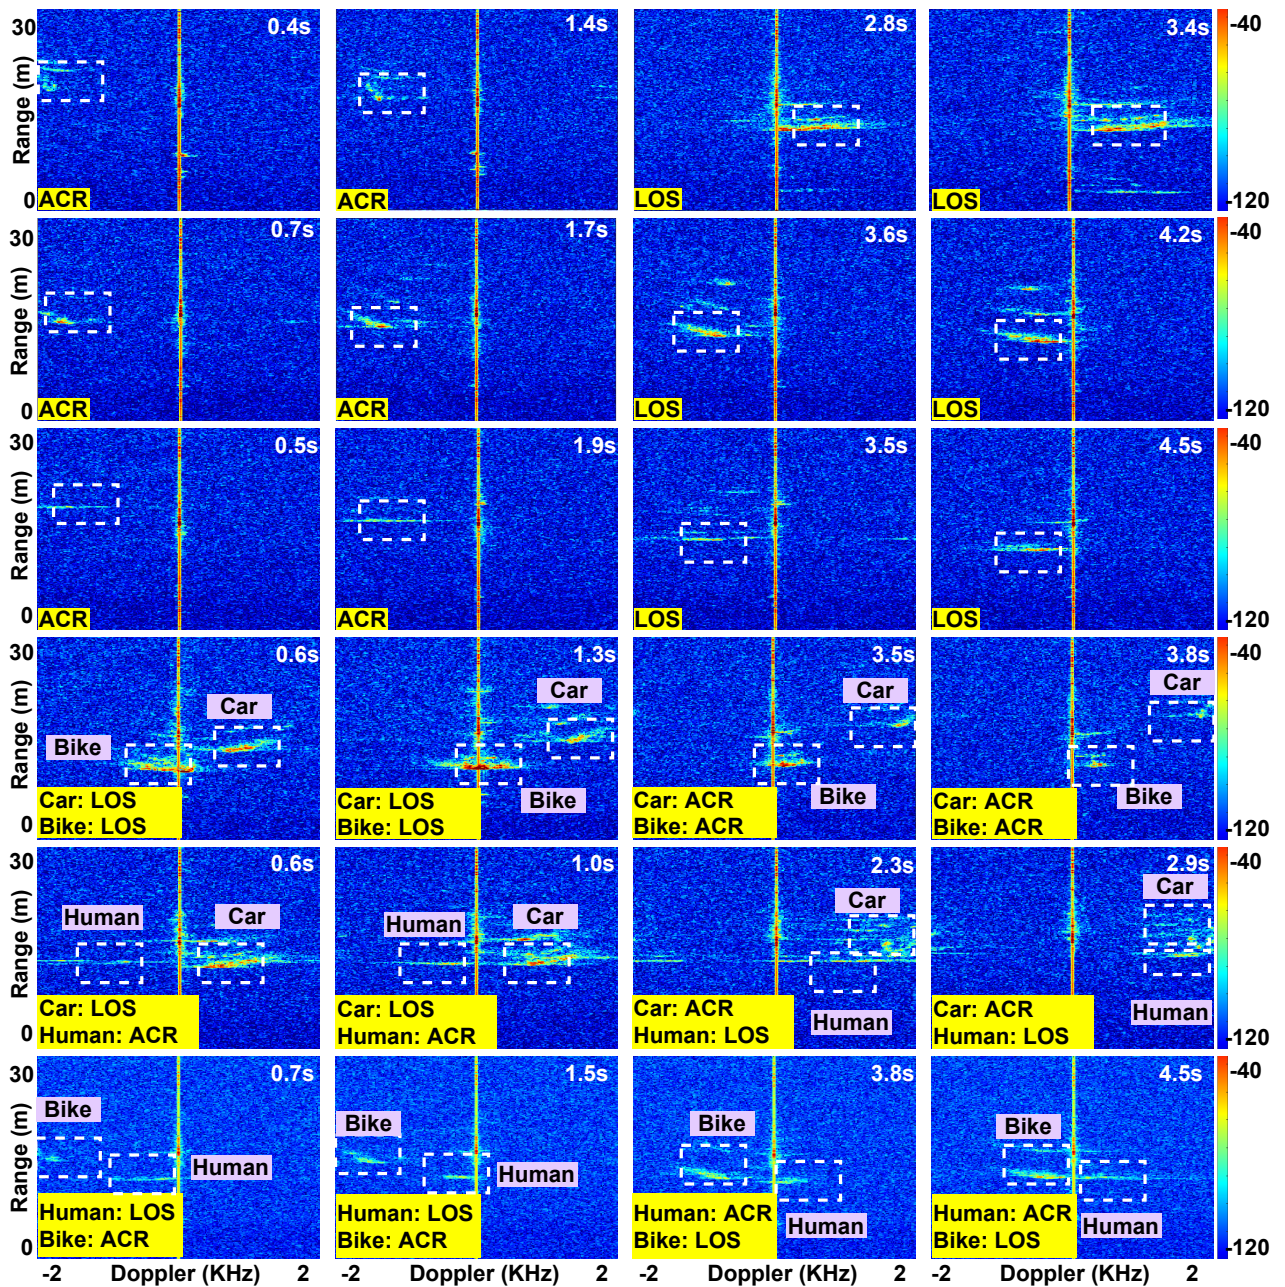


Figure 5.7: Range-Doppler ambiguity plots of automotive targets generated with a 77GHz automotive radar in ACR and LOS conditions. Top row: small size car; second row: bike; third row: human; fourth row: car and bike; fifth row: car and human and sixth row: car and human. The white squares in the figures depict the position of the targets.

as shown in Fig.5.8 a(i). The range-Doppler ambiguity plot corresponding to a single frame or CPI is presented in Fig.5.8 a(ii). Due to considerable overlap in the radar returns, separation using bounding box principles (as in computer vision) is not possible. Instead, we use the sparsity-based disaggregation algorithm presented in section 5.2.1. The ground truth and reconstructed plots for the car are shown in Figs.5.8 a(iii) and a(iv), respectively; while those for the bike are shown in Figs.5.8 a(v) and a(vi). Qualitatively, the reconstructed plots of the two targets closely resemble each other. Next, a similar analysis has been done for the case where two targets are human and car Fig.5.8 b(i)-(vi). Then, we consider the third case where the two targets are the bike and the human moving in opposite directions, as shown in Fig.5.8 c(i). Since their speeds are different, we observe some Doppler and range displacement in the two targets' range-Doppler ambiguity plot shown in Fig.5.8 c(ii). As mentioned before, both these targets are spatially narrow along the range dimension. However, we see multiple peaks along the range dimension due to multipath scattering. Also, there is a significant micro-Doppler spread from the wheel rotation in the case of the bike and the motion of legs for the human. Due to the close similarity between these two targets, we see that in some instances, the disaggregation is not carried out as well as in the previous two cases with the car. For example, the reconstructed plot for the bike, in Fig.5.8 c(iv), shows the returns from the bike (dotted arrow) along with returns from the human (solid arrow). Hence, the reconstructed plot does not visually resemble the ground truth plot shown in Fig.5.8 c(iii). The reconstruction of the human was more effective as the

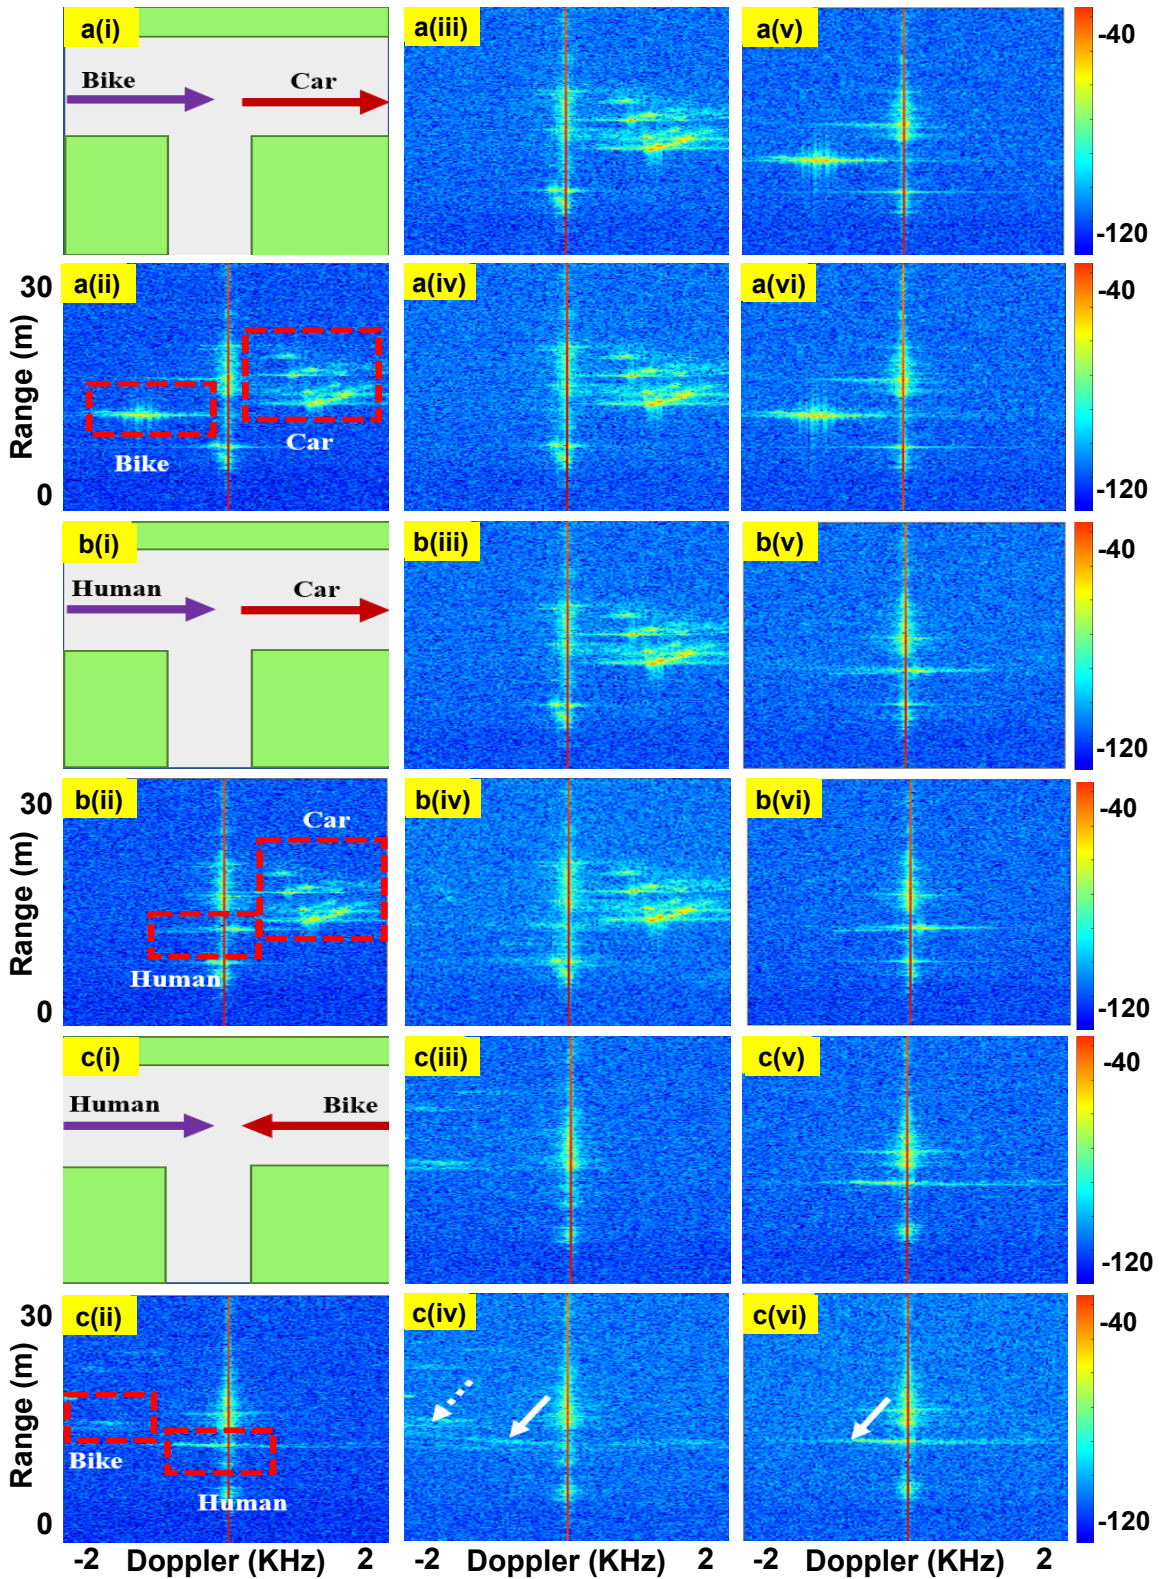


Figure 5.8: Trajectories followed by (a) car and bike, (b) car and human, (c) human and bike shown in a(i), b(i) and c(i) and their respective range-Doppler ambiguity plots shown in a(ii), b(ii) and c(ii). Ground truth plots of *car* for cases (a) and (b) shown in a(iii) and b(iii); *bike* for cases (a) and (c) shown in a(v) and c(iii); *human* for cases (b) and (c) shown in b(v) and c(v). Reconstructed plots of *car*, *bike* and *human* for the respective cases shown in a(vi)-c(vi).

reconstructed plot in Fig.5.8 c(vi) closely resembles the ground truth plot shown in Fig.5.8 c(v). Note that all the results presented for qualitative comparison are for a single frame. We perform a quantitative comparison by measuring the NMSE between the reconstructed and ground truth plots for all frames of the two target cases. The NMSE are 0.26, 0.30, and 0.19 for car, bike, and human respectively. The reconstruction for car and human is slightly more effective than the bike since the latter shares some common features with both the other targets. The bike is spatially narrow like a human but has a high micro-Doppler spread from multiple wheels like the car.

***Detection of multiple targets:*** In this section, we use the disaggregation algorithm to detect and identify the presence of multiple targets. The dictionaries for each of the target classes are trained using the single-target data, while we consider the *real* two target measurement data for the test. Since it is a supervised learning algorithm, we disaggregate the test radar data into the three target classes - car, bike, and human. Then, we check if the strength of the reconstructed signal for each class is above a pre-defined threshold to see whether the target class is present. Hence, there are three possible outcomes for each target class- a correct detection, a missed detection, and a wrong detection/false alarm. The thresholds for each target class are uniquely chosen since each target class is characterized by a different radar cross-section. In order to select an appropriate threshold, we first create a histogram of the strength of the target signals from the single target data for each class. Then we select the threshold such that the

probability of detection is maximized for a fixed value of the probability of false alarms, which is chosen to be 0.26 in this work. Based on these calculations, the threshold for the car, bike, and human are 0.2, 0.09, and 0.065 volts, respectively. We consider data from the three two-target scenarios and present the detection results in Table 5.2. The three scenarios are - 50 samples each of (1) car and

Table 5.2: Number of detections and false alarms for two target scenario.

Cases	Car	Bike	Human
Car-bike (100)	90	75	9
Car-human (100)	91	13	95
Human-bike (100)	10	79	93

bike, (2) car and human, and (3) human and bike. In the first case, some of the bike returns are missed by the algorithm, while there are some false alarms of humans. In the second case of car and human, almost all of the human cases are correctly identified, but some of the car returns are missed. Again, there are few false alarms from the bike. In the third case, when a human and bike are present, we obtain very few false alarms of the car. Based on the above results, we observe that the dictionaries of the bike are somewhat similar to those of the human due to their similarity in the range-Doppler ambiguity plot. As a result, the algorithm is at times confused between these two targets. The problem is likely to be alleviated with a greater amount of training data for the target classes under diverse conditions. Note that these results have been obtained when the training data are generated in single-target conditions while the test data are gathered in similar but not identical two target conditions. Finally, we note that using these algorithms, we were able to successfully detect and identify two targets in NLOS conditions using ACR range-Doppler ambiguity plots.

### 5.3 Summary of findings

In this chapter, we explored two applications of mm-wave automotive radar technology, demonstrating its versatility in enhancing automotive safety and functionality. Firstly, We have demonstrated the use of an externally mounted automotive radar to conduct automatic parking tests based on ISAR radar images. We have developed an automated parking test algorithm to classify the parking maneuvers of a VUT as either correct or incorrect based on the estimated trajectory of the VUT from the ISAR images. We have experimentally validated our proposed parking test algorithm using the data gathered using the TI-AWR 1843 mm-wave sensor. Secondly, We have proposed a novel application of mm-wave automotive radars for NLOS look-ahead sensing of multiple automotive targets for collision avoidance. We experimentally demonstrate the viability of this application using a 77GHz automotive radar to generate range-Doppler ambiguity plots of cars, bikes, and humans. The resulting plots show multipath scattered returns from the target along the range and micro-Doppler spread from the rotational motion of the wheels or the swinging motion of the arms and legs of the pedestrian. When multiple such targets are present together, the overlap in their range-Doppler ambiguity plots can result in challenges for correct detection and identification of the targets. We have used a sparsity-based disaggregation algorithm to separate the target returns and then correctly identify each target. Our results show a detection accuracy of 88.3% for real measured data.

*Wall interaction and skin depth at 77 GHz.* At W-band, the electromagnetic

skin depth in typical concrete, brick, and stone walls is on the order of only a few millimetres, so the field interacts mainly with a thin surface layer and deep through-wall propagation is strongly attenuated. In the reported 77 GHz ACR configuration, the dominant NLOS mechanism is therefore surface and edge scattering along the outer walls and corner, rather than volumetric propagation inside the wall. In this thesis, these effects are captured through an effective additional interaction loss and scattering term in the NLOS path; more detailed material- and structure-specific wall modelling is left as an avenue for future work.

Together, these applications underscore the potential of advanced radar technologies in improving automotive systems. The automated parking test algorithm helps to automate the test ability to parking test maneuver in driver license test, while the NLOS sensing application contributes to effective collision avoidance in complex environments. Future work will build on these findings to refine the algorithms and expand their applicability across diverse driving scenarios.

## **Chapter 6**

# **Conclusion and Future Work**

The contributions of this dissertation are as follows:

### **6.1 High-Fidelity ISAR Image Simulation Framework for Automotive Targets at mm-wave Radar Frequency**

This thesis presents a high-fidelity point-scatterer based simulation framework of low computational complexity that effectively models the radar scattering phenomenology of common road vehicles. The simulation methodology was validated through benchmarking with measurement data collection of automotive targets with the TI AWR1843 radar. Through the simulation framework, we successfully generated a large volume of high-resolution ISAR radar images of dynamic automotive targets at mm-wave frequencies characterized by fine range and cross-range resolution. This contribution advances radar imaging in automotive applications, providing a valuable tool for generating high-quality radar images in large quantities, thereby supporting a variety of applications,

including (1) providing training data for classifiers, (2) rapid testing of new sensors and algorithms, (3) for understanding underlying physics of electromagnetic phenomenology.

## **6.2 Modeling Clutter in Automotive Radar Scenarios**

Characterizing road clutter is essential for optimizing radar detectors and generating realistic simulation models. This thesis introduces detailed models for various types of clutter based on measurement data from 77 GHz automotive radar, specifically under Indian road conditions. The clutter model considers road surfaces with varying roughness across diverse environmental conditions, including temperature, wind speeds, and surface moisture. The research developed mathematical models for static and dynamic clutters encountered on roads, demonstrating their impact on radar image quality and classification performance. Notably, the Gaussian Mixture Model (GMM) distribution was identified as the best fit for capturing the diversity in clutter characteristics.

Road-surface backscatter (mean  $\sigma^0$ ): concrete  $\approx -25.0$  dB; asphalt  $\approx -23.6$  dB.

## **6.3 Public Database of Automotive Radar Data**

A key contribution of this thesis is creating and releasing a comprehensive public dataset designed to facilitate the development and benchmarking of new algorithms within the radar community. This dataset includes a curated collection of labeled ISAR radar images, both simulated and measured, featuring automotive

targets across diverse aspects, trajectories, and noise and clutter conditions. *Data access:* Public dataset link — <https://tinyurl.com/msu6aj7y>.

## 6.4 Classification of Automotive Targets Using ISAR Images

This thesis demonstrates the effectiveness of ISAR images as features for object classification in automotive radar systems. The study employed off-the-shelf classification algorithms, including two ML-based approaches (SVM and RF) and two deep learning algorithms based on transfer learning (AlexNet and GoogLeNet). Initially, the classification algorithms were trained and tested using simulated ISAR images to analyze the impact of noise, clutter, and test and training data volume on classification performance. An important related contribution to this work is that this is the first time, to the best of our knowledge, that classifiers trained on simulation data were tested on measurement data. This showcases the fidelity of the simulation framework and the robustness of ISAR features for classification. On simulated data, transfer learning baselines reached up to 99.9% accuracy under range–Doppler clutter (AlexNet); robust performance across noise/clutter settings (Table 2.4). Training on simulation and testing on measured data achieved  $F_1=94.2\%$  (AlexNet) vs. 75.3% (SVM). (Table 2.6).

## 6.5 Explainable AI in Radar Image Classification

This thesis explores explainable AI techniques to improve trust and transparency in the decision-making process of automotive radar systems, paving the way toward broader commercial adoption and ensuring compliance with regulatory standards. Specifically, the focus is to study the causes of the misclassification of ISAR images with high SNR. A generative AI framework is introduced, which realistically perturbs ISAR images belonging to one class of automotive targets to be now misclassified as a second class. This generative approach introduces perturbations that mirror the occlusions of some parts of the vehicle and shows that this leads to misclassification under high SNR conditions. For example, occluding one or more wheels of a four-wheeled car could cause it to be misclassified as a three-wheeled autorickshaw. Similarly, occluding one wheel of an autorickshaw could result in its misclassification as a bicycle.

## 6.6 Applications

**Automated Parking Test Using ISAR Images.** This thesis also introduces a novel approach to automated parking tests using ISAR images. Based on these images, a method was developed to guide and validate parking maneuvers, and an algorithm was created to test for accurate parking in both angled and perpendicular scenarios.

**Around-the-Corner Radar Sensing.** Finally, the thesis contributes to de-

veloping and validating techniques for around-the-corner radar sensing, which significantly enhances situational awareness in look-ahead scenarios for automotive applications. A method was proposed to detect and classify targets in non-line-of-sight conditions using ISAR images, and experiments were conducted to demonstrate the feasibility of this approach. The findings from this research have important implications for improving safety and awareness in complex urban environments where traditional line-of-sight sensors may be insufficient. Mean successful detection after disaggregation: **88.3%**.

Overall, the advancements presented in this thesis significantly improve automotive radar technologies, particularly in enhancing imaging, classification, and sensing capabilities, while also promoting transparency and trust in AI-based radar systems. We believe that these developments will influence the future of automotive safety and automation, creating more intelligent and reliable vehicles.

## **6.7 Future Research Avenues**

Based on the findings and contributions of this thesis, several avenues for future research are proposed:

### **6.7.1 Real-World Testing and Validation**

Comprehensive real-world testing and validation of automotive ISAR imaging are essential to ensure the practical applicability of the developed models and algorithms. Future research should involve conducting large-scale field tests to

assess the effectiveness of these techniques across various automotive scenarios. Collaborating with industry partners will be essential to integrating these technologies into commercial systems, validating their performance, and ensuring their reliability in real-world applications. Concretely, priority data collection should cover multiple physical instances per class (cars—sedan/hatchback/SUV; autorickshaws; bikes; pedestrians) across urban ACR corners and open-lot parking scenes, with right/left turns, U-turns, and angled/perpendicular parking trajectories sampled at slow/medium/fast speeds, including curated two-target compositions (e.g., car+bike, car+pedestrian) to stress disaggregation.

### **6.7.2 ISAR Classification and Data Verification**

Future work will address the limitations of current data collection methods by focusing on acquiring multiple instances of each target class. While the existing models are based on single-target instances, real-world applications often involve greater diversity within target classes. Exploring unsupervised learning algorithms to handle new and unseen targets will be crucial. This research will move beyond supervised learning constraints, improving the system's ability to generalize to novel radar data and enhance classification accuracy.

### **6.7.3 Advanced Explainable AI Techniques**

To further advance explainable AI in radar applications, future research should focus on extending current techniques to handle a broader range of radar phenomena. For example, a study investigating how explainable AI methods can be

applied to clutter and multipath effects and how they can be adapted to different types of automotive radar signatures, such as joint time-frequency micro-Doppler signatures and radar point cloud representations. These advancements will enhance our understanding of AI model decisions and improve the transparency and trustworthiness of radar-based systems.

#### **6.7.4 Radar-Camera Sensor Fusion**

Camera and radar offer complementary sensing paradigms. While camera images provide crucial information regarding the target shape, color, texture, and high-resolution front-view images, radar provides top-view ISAR images and Doppler velocity information. Modern intelligent vehicles are also likely to be equipped with both sensors. Therefore, combining radar and camera data through sensor fusion presents a critical avenue for future research. Developing methods to integrate ISAR and camera images to provide three-dimensional target images can significantly enhance target classification and tracking capabilities across diverse and challenging scenarios.

#### **6.7.5 Enhanced Diversity in Radar Systems and Algorithms**

Future research should expand the diversity of radar systems from monostatic to bistatic and multistatic frameworks. Similarly, investigation into new signal processing algorithms and the metrics used for their evaluation in these frameworks would be valuable for expanding the performance scope. Developing standardized evaluation metrics will help assess the effectiveness of automotive

radar systems in detecting and classifying targets under different conditions. Such efforts will enable a more comprehensive comparison and enhancement of radar technologies, driving forward the field of automotive radar. Future work includes extending the ISAR/ACR analyses to **bistatic** (and, where feasible, multistatic) configurations.

## References

- [1] “WHO Global health observatory data,” <https://www.who.int/data/gho>, accessed: 2024-08-20.
- [2] “WHO Global status report on road safety 2023,” <https://www.who.int/publications/i/item/9789240086517>, accessed: 2024-08-20.
- [3] “WHO Risk factors for road traffic injuries,” <https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries>, accessed: 2024-08-20.
- [4] “Federal motor vehicle safety standards and regulations ,” <https://www.nhtsa.gov/sites/nhtsa.gov/files/fmvss-quickrefguide-hs811439.pdf>, accessed: 2021-08-03.
- [5] “Vehicle safety 2016 European Commission EUROPA ,” [https://ec.europa.eu/transport/road\\_safety/sites/default/files/ersosynthesis2016-vehiclesafety15\\_en.pdf](https://ec.europa.eu/transport/road_safety/sites/default/files/ersosynthesis2016-vehiclesafety15_en.pdf), accessed: 2021-08-03.
- [6] “Vehicular safety standards and regulations ,” <https://www.siam.in/technical-regulation.aspx?mpgid=31&pgidtrail=34>, accessed: 2021-08-03.

- [7] “SAE International releases updated visual chart for its levels of driving automation standard for self-driving vehicles,” <https://tinyurl.com/3ekycpn4>, accessed: 2024-08-22.
- [8] K. Bengler, K. Dietmayer, B. Farber, M. Maurer, C. Stiller, and H. Winner, “Three decades of driver assistance systems: Review and future perspectives,” *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 4, pp. 6–22, winter 2014.
- [9] “Advanced driver assistance systems - European Commission,” [https://ec.europa.eu/transport/road\\_safety/sites/default/files/pdf/ersosynthesis2018-adas-summary.pdf](https://ec.europa.eu/transport/road_safety/sites/default/files/pdf/ersosynthesis2018-adas-summary.pdf), accessed: 2021-08-03.
- [10] S. Campbell, N. O’Mahony, L. Krpalcova, D. Riordan, J. Walsh, A. Murphy, and C. Ryan, “Sensor technology in autonomous vehicles: A review,” in *2018 29th Irish Signals and Systems Conference (ISSC)*. IEEE, 2018, pp. 1–4.
- [11] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt, “Millimeter-wave technology for automotive radar sensors in the 77 ghz frequency band,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 60, no. 3, pp. 845–860, 2012.
- [12] R. H. Rasshofer and K. Gresser, “Automotive radar and lidar systems for next generation driver assistance functions.” *Advances in Radio Science*, vol. 3, 2005.

- [13] M. Schneider, “Automotive radar-status and trends,” in *German microwave conference*, 2005, pp. 144–147.
- [14] T. Zhou, M. Yang, K. Jiang, H. Wong, and D. Yang, “Mmw radar-based technologies in autonomous driving: A review,” *Sensors*, vol. 20, no. 24, p. 7283, 2020.
- [15] “The global automotive radar market size ,” <https://www.fortunebusinessinsights.com/industry-reports/automotive-radar-market-101688>, accessed: 2021-08-03.
- [16] A. Sume, M. Gustafsson, M. Herberthson, A. Janis, S. Nilsson, J. Rahm, and A. Orbom, “Radar detection of moving targets behind corners,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 6, pp. 2259–2267, 2011.
- [17] B. Fleming, “Recent advancement in automotive radar systems [automotive electronics],” *IEEE Vehicular Technology Magazine*, vol. 7, no. 1, pp. 4–9, 2012.
- [18] Y. Cui, R. Chen, W. Chu, L. Chen, D. Tian, Y. Li, and D. Cao, “Deep learning for image and point cloud fusion in autonomous driving: A review,” *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–18, 2021.
- [19] I. Jouny, F. Garber, and S. Ahalt, “Classification of radar targets using synthetic neural networks,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 29, no. 2, pp. 336–344, 1993.

- [20] H.-C. Chiang, R. L. Moses, and L. C. Potter, "Model-based classification of radar images," *IEEE Transactions on Information Theory*, vol. 46, no. 5, pp. 1842–1854, 2000.
- [21] I. Bilik, J. Tabrikian, and A. Cohen, "Gmm-based target classification for ground surveillance doppler radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 1, pp. 267–278, 2006.
- [22] I. Urazghildiiev, R. Ragnarsson, P. Ridderstrom, A. Rydberg, E. Ojefors, K. Wallin, P. Enochsson, M. Ericson, and G. Lofqvist, "Vehicle classification based on the radar measurement of height profiles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 2, pp. 245–253, 2007.
- [23] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, "Tracking of extended objects with high-resolution doppler radar," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 5, pp. 1341–1353, 2016.
- [24] S. Kuutti, R. Bowden, Y. Jin, P. Barber, and S. Fallah, "A survey of deep learning applications to autonomous vehicle control," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [25] X. Cai, M. Giallorenzo, and K. Sarabandi, "Machine learning-based target classification for mmw radar in autonomous driving," *IEEE Transactions on Intelligent Vehicles*, pp. 1–1, 2021.

- [26] Y. Kim and H. Ling, “Human activity classification based on micro-doppler signatures using a support vector machine,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 5, pp. 1328–1337, 2009.
- [27] Y. Kim and T. Moon, “Human detection and activity classification based on micro-doppler signatures using deep convolutional neural networks,” *IEEE Geoscience and Remote Sensing Letters*, vol. 13, no. 1, pp. 8–12, 2015.
- [28] S. Vishwakarma and S. S. Ram, “Dictionary learning with low computational complexity for classification of human micro-dopplers across multiple carrier frequencies,” *IEEE Access*, vol. 6, pp. 29 793–29 805, 2018.
- [29] M. S. Seyfioğlu, A. M. Özbayoğlu, and S. Z. Gürbüz, “Deep convolutional autoencoder for radar-based classification of similar aided and unaided human activities,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1709–1723, 2018.
- [30] F. Fioranelli, M. Ritchie, and H. Griffiths, “Classification of un-armed/armed personnel using the netrad multistatic radar for micro-doppler and singular value decomposition features,” *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 9, pp. 1933–1937, 2015.
- [31] R. Prophet, M. Hoffmann, A. Ossowska, W. Malik, C. Sturm, and M. Vossiek, “Pedestrian classification for 79 ghz automotive radar sys-

- tems,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1265–1270.
- [32] P. Khomchuk, I. Stainvas, and I. Bilik, “Pedestrian motion direction estimation using simulated automotive mimo radar,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 3, pp. 1132–1145, 2016.
- [33] M. Ritchie, F. Fioranelli, H. Borrión, and H. Griffiths, “Multistatic micro-doppler radar feature extraction for classification of unloaded/loaded micro-drones,” *IET Radar, Sonar & Navigation*, vol. 11, no. 1, pp. 116–124, 2016.
- [34] P. Molchanov, K. Egiazarian, J. Astola, R. Harmanny, and J. De Wit, “Classification of small uavs and birds by micro-doppler signatures,” in *2013 European Radar Conference*. IEEE, 2013, pp. 172–175.
- [35] I. Bilik and J. Tabrikian, “Radar target classification using doppler signatures of human locomotion models,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1510–1522, 2007.
- [36] S. Chen, H. Wang, F. Xu, and Y.-Q. Jin, “Target classification using the deep convolutional networks for sar images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 8, pp. 4806–4817, 2016.
- [37] S. A. Wagner, “Sar atr by a combination of convolutional neural network and support vector machines,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 6, pp. 2861–2872, 2016.

- [38] Z. Lin, K. Ji, M. Kang, X. Leng, and H. Zou, “Deep convolutional highway unit network for sar target classification with limited labeled training data,” *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 7, pp. 1091–1095, 2017.
- [39] S. S. Ram and A. Majumdar, “High-resolution radar imaging of moving humans using doppler processing and compressed sensing,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 2, pp. 1279–1287, 2015.
- [40] V. C. Chen, *Inverse Synthetic Aperture Radar Imaging; Principles*. Institution of Engineering and Technology, 2014.
- [41] S. S. Ram, “Fusion of inverse synthetic aperture radar and camera images for automotive target tracking,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [42] M. Martorella, E. Giusti, A. Capria, F. Berizzi, and B. Bates, “Automatic target recognition by means of polarimetric isar images and neural networks,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 11, pp. 3786–3794, 2009.
- [43] M. Vespe, C. Baker, and H. Griffiths, “Automatic target recognition using multi-diversity radar,” *IET Radar, Sonar & Navigation*, vol. 1, no. 6, pp. 470–478, 2007.

- [44] S.-H. Park, M.-G. Joo, and K.-T. Kim, "Construction of isar training database for automatic target recognition," *Journal of Electromagnetic Waves and Applications*, vol. 25, no. 11-12, pp. 1493–1503, 2011.
- [45] J. S. Kulpa, M. Malanowski, D. Gromek, P. Samczyński, K. Kulpa, and A. Gromek, "Experimental results of high-resolution isar imaging of ground-moving vehicles with a stationary fmcw radar," *International Journal of Electronics and Telecommunications*, vol. 59, no. 3, pp. 293–299, 2013.
- [46] C. J. Li and H. Ling, "Wide-angle isar imaging of vehicles," in *2015 9th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 2015, pp. 1–2.
- [47] C. Li and H. Ling, "Wide-angle, ultra-wideband isar imaging of vehicles and drones," *Sensors*, vol. 18, no. 10, p. 3311, 2018.
- [48] A. A. Danylov, T. M. Goyette, J. Waldman, M. J. Coulombe, A. J. Gatesman, R. H. Giles, X. Qian, N. Chandrayan, S. Vangala, K. Termkoa *et al.*, "Terahertz inverse synthetic aperture radar (isar) imaging with a quantum cascade laser transmitter," *Optics express*, vol. 18, no. 15, pp. 16 264–16 272, 2010.
- [49] M. Boshra and B. Bhanu, "Bounding sar atr performance based on model similarity," in *Algorithms for Synthetic Aperture Radar Imagery VI*, vol. 3721. International Society for Optics and Photonics, 1999, pp. 716–729.

- [50] H. Essen, M. Hagelen, W. Johannes, R. Sommer, A. Wahlen, M. Schlechtweg, and A. Tessmann, “High resolution millimetre wave measurement radars for ground based sar and isar imaging,” in *2008 IEEE Radar Conference*. IEEE, 2008, pp. 1–5.
- [51] G. Rubin, E. V. Sager, and D. H. Berger, “Gpu acceleration of sar/isar imaging algorithms,” in *Antenna Measurement Techniques Association Symposium*, 2010, pp. 430–435.
- [52] D. H. Nguyen, J. H. Kay, B. J. Orchard, and R. H. Whiting, “Classification and tracking of moving ground vehicles,” *Lincoln Laboratory Journal*, vol. 13, no. 2, pp. 275–308, 2002.
- [53] A. R. Brenner and J. H. Ender, “Very wideband radar imaging with the airborne sar sensor pamir,” in *IGARSS 2003. 2003 IEEE International Geoscience and Remote Sensing Symposium. Proceedings (IEEE Cat. No. 03CH37477)*, vol. 1. IEEE, 2003, pp. 533–535.
- [54] D. A. Morgan, “Deep convolutional neural networks for atr from sar imagery,” in *Algorithms for Synthetic Aperture Radar Imagery XXII*, vol. 9475. International Society for Optics and Photonics, 2015, p. 94750F.
- [55] O. Karabayır, O. M. Yücedağ, M. Z. Kartal, and H. A. Serim, “Convolutional neural networks-based ship target recognition using high resolution range profiles,” in *2017 18th International Radar Symposium (IRS)*. IEEE, 2017, pp. 1–9.

- [56] B. Xue and N. Tong, "Diod: Fast and efficient weakly semi-supervised deep complex isar object detection," *IEEE Transactions on Cybernetics*, 2018.
- [57] Y. Deep, P. Held, S. S. Ram, D. Steinhauser, A. Gupta, F. Gruson, A. Koch, and A. Roy, "Radar cross-sections of pedestrians at automotive radar frequencies using ray tracing and point scatterer modelling," *IET Radar, Sonar & Navigation*, 2020.
- [58] S. S. Ram, C. Christianson, Y. Kim, and H. Ling, "Simulation and analysis of human micro-dopplers in through-wall environments," *IEEE Transactions on Geoscience and remote sensing*, vol. 48, no. 4, pp. 2015–2023, 2010.
- [59] G. Duggal, K. V. Mishra, and S. S. Ram, "Micro-doppler and micro-range detection via doppler-resilient 802.11ad-based vehicle-to-pedestrian radar," in *IEEE Radar Conference 2019*, 2019, pp. 1–6.
- [60] G. Duggal, S. Vishwakarma, K. V. Mishra, and S. S. Ram, "Doppler-resilient 802.11ad-based ultra-short range automotive joint radar-communications system," *IEEE Transactions on Aerospace and Electronic Systems*, 2020.
- [61] K. Sarabandi, E. S. Li, and A. Nashashibi, "Modeling and measurements of scattering from road surfaces at millimeter-wave frequencies," *IEEE Transactions on Antennas and Propagation*, vol. 45, no. 11, pp. 1679–1688, 1997.

- [62] M. Saidi, B. Hoeltzener, A. Toumi, A. Khecnhaf, and D. Aboutajdine, “Automatic recognition of isar images: Target shapes features extraction,” in *2008 3rd International Conference on Information and Communication Technologies: From Theory to Applications*. IEEE, 2008, pp. 1–6.
- [63] M. Giallorenzo, X. Cai, A. Nashashibi, and K. Sarabandi, “Radar backscatter measurements of road surfaces at 77 ghz,” in *2018 IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*. IEEE, 2018, pp. 2421–2422.
- [64] G. P. Kulemin, *Millimeter-wave radar targets and clutter*. Artech House, 2003.
- [65] A. Adadi and M. Berrada, “Peeking inside the black-box: a survey on explainable artificial intelligence (xai),” *IEEE access*, vol. 6, pp. 52 138–52 160, 2018.
- [66] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins *et al.*, “Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai,” *Information fusion*, vol. 58, pp. 82–115, 2020.
- [67] X.-H. Li, C. C. Cao, Y. Shi, W. Bai, H. Gao, L. Qiu, C. Wang, Y. Gao, S. Zhang, X. Xue *et al.*, “A survey of data-driven and knowledge-aware explainable ai,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 1, pp. 29–49, 2020.

- [68] A. Dhurandhar, P.-Y. Chen, R. Luss, C.-C. Tu, P. Ting, K. Shanmugam, and P. Das, “Explanations based on the missing: Towards contrastive explanations with pertinent negatives,” *Advances in neural information processing systems*, vol. 31, 2018.
- [69] P. Hase and M. Bansal, “Evaluating explainable ai: Which algorithmic explanations help users predict model behavior?” *arXiv preprint arXiv:2005.01831*, 2020.
- [70] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang, “A survey on neural network interpretability,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 5, no. 5, pp. 726–742, 2021.
- [71] S. Joshi, O. Koyejo, W. Vijitbenjaronk, B. Kim, and J. Ghosh, “Towards realistic individual recourse and actionable explanations in black-box decision making systems,” *arXiv preprint arXiv:1907.09615*, 2019.
- [72] S. Wachter, B. Mittelstadt, and C. Russell, “Counterfactual explanations without opening the black box: Automated decisions and the gdpr,” *Harv. JL & Tech.*, vol. 31, p. 841, 2017.
- [73] M. I. Skolnik *et al.*, *Introduction to radar systems*. McGraw-hill New York, 1980, vol. 3.
- [74] J. Lekner, *Theory of reflection of electromagnetic and particle waves*. Springer Science & Business Media, 1987, vol. 3.
- [75] J. R. Wait, *Electromagnetic waves in stratified media: Revised edition including supplemented material*. Elsevier, 2013, vol. 3.

- [76] C. Li, Z. Peng, T.-Y. Huang, T. Fan, F.-K. Wang, T.-S. Horng, J.-M. Munoz-Ferreras, R. Gomez-Garcia, L. Ran, and J. Lin, “A review on recent progress of portable short-range noncontact microwave radar systems,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 65, no. 5, pp. 1692–1706, 2017.
- [77] S. S. Ram, “Representation of radar micro-dopplers using customized dictionaries,” *Topics in Radar Signal Processing*, p. 197, 2018.
- [78] G. T. Ruck, D. E. Barrick, W. D. Stuart, and C. K. Krichbaum, *Radar cross section handbook*. Plenum press New York, 1970, vol. 1.
- [79] M. A. Richards, *Fundamentals of radar signal processing*. Tata McGraw-Hill Education, 2005.
- [80] F. Ulaby, M. C. Dobson, and J. L. Álvarez-Pérez, *Handbook of radar scattering statistics for terrain*. Artech House, 2019.
- [81] E. S. Li and K. Sarabandi, “Low grazing incidence millimeter-wave scattering models and measurements for various road surfaces,” *IEEE Transactions on Antennas and Propagation*, vol. 47, no. 5, pp. 851–861, 1999.
- [82] P. Pandey, D. Kumar, A. Prakash, J. Masih, M. Singh, S. Kumar, V. K. Jain, and K. Kumar, “A study of urban heat island and its association with particulate matter during winter months over delhi,” *Science of the Total Environment*, vol. 414, pp. 494–507, 2012.

- [83] B. Haywood and R. Evans, "Motion compensation for isar imaging," in *Proceedings of Australian Symposium on Signal processing and applications*, 1989, pp. 112–117.
- [84] K. D. Ward, S. Watts, and R. J. Tough, *Sea clutter: scattering, the K distribution and radar performance*. IET, 2006, vol. 20.
- [85] A. A. Neath and J. E. Cavanaugh, "The bayesian information criterion: background, derivation, and applications," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 4, no. 2, pp. 199–203, 2012.
- [86] S. Watanabe, "A widely applicable bayesian information criterion," *Journal of Machine Learning Research*, vol. 14, no. Mar, pp. 867–897, 2013.
- [87] P. A. Bosman and D. Thierens, "Negative log-likelihood and statistical hypothesis testing as the basis of model selection in ideas," in *Proceedings of the Tenth Dutch–Netherlands Conference on Machine Learning*. Tilburg University, 2000.
- [88] J. H. Friedman, *The elements of statistical learning: Data mining, inference, and prediction*. springer open, 2017.
- [89] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [90] S. Lee, Y.-J. Yoon, J. Yoon, H. Sim, and S.-C. Kim, "Periodic clutter suppression in iron road structures for automotive radar systems," *IET Radar, Sonar & Navigation*, vol. 12, no. 10, pp. 1146–1153, 2018.

- [91] N. Pandey and S. S. Ram, “Classification of automotive targets using inverse synthetic aperture radar images,” *IEEE Transactions on Intelligent Vehicles*, 2022.
- [92] S. S. Ram, G. Singh, and G. Ghatak, “Estimating radar detection coverage probability of targets in a cluttered environment using stochastic geometry,” in *2020 IEEE International Radar Conference (RADAR)*. IEEE, 2020, pp. 665–670.
- [93] —, “Optimization of radar parameters for maximum detection probability under generalized discrete clutter conditions using stochastic geometry,” *IEEE Open Journal of Signal Processing*, vol. 2, pp. 571–585, 2021.
- [94] J. B. Billingsley, *Low-angle radar land clutter: measurements and empirical models*. IET, 2002.
- [95] G. Kulemin, E. Goroshko, and E. Tarnavsky, “Land backscattering model for millimeter wave radar,” in *Proceedings of the International Conference Modern Problems of Radio Engineering, Telecommunications and Computer Science, 2004*. IEEE, 2004, pp. 138–141.
- [96] A. M. Raynal, D. L. Bickel, D. F. Dubbert, T. Verge, B. L. Burns, R. Dunkel, and A. W. Doerry, “Radar cross section statistics of cultural clutter at ku-band,” *Radar Sensor Technology XVI*, vol. 8361, 2012.
- [97] Y.-Z. Ma, C. Cui, B.-S. Kim, J.-M. Joo, S. H. Jeon, and S. Nam, “Road clutter spectrum of bsd fmcw automotive radar,” in *2015 European Radar Conference (EuRAD)*. IEEE, 2015, pp. 109–112.

- [98] S. Lee, B.-H. Lee, J.-E. Lee, and S.-C. Kim, “Statistical characteristic-based road structure recognition in automotive fmcw radar systems,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 7, pp. 2418–2429, 2018.
- [99] M. Skolnik, *Introduction to Radar Systems*, ser. Electrical engineering series. McGraw-Hill, 2001. [Online]. Available: <https://books.google.co.in/books?id=Y6-APwAACAAJ>
- [100] C. Li, J. Zhao, J. Yin, G. Zhang, and X. Shan, “Analysis of rcs characteristic of dihedral corner and triangular trihedral corner reflectors,” in *2010 5th International Conference on Computer Science & Education*. IEEE, 2010, pp. 40–43.
- [101] I. R. C. (IRC), “Guidelines for parking facilities in urban areas (irc:sp-12-2015),” *Ministry of Road Transport & Highways (MORTH, formerly MOST)*, 2015.
- [102] C. Forbes, M. Evans, N. Hastings, and B. Peacock, *Statistical distributions*. John Wiley & Sons, 2011.
- [103] S. M. Ross, *Introduction to probability models*. Academic press, 2014.
- [104] R. Dubé, M. Hahn, M. Schütz, J. Dickmann, and D. Gingras, “Detection of parked vehicles from a radar based occupancy grid,” in *2014 IEEE Intelligent Vehicles Symposium Proceedings*. IEEE, 2014, pp. 1415–1420.

- [105] K. Patel, K. Rambach, T. Visentin, D. Rusev, M. Pfeiffer, and B. Yang, “Deep learning-based object classification on automotive radar spectra,” in *2019 IEEE Radar Conference (RadarConf)*. IEEE, 2019, pp. 1–6.
- [106] L. Du, H. Liu, Z. Bao, and J. Zhang, “Radar automatic target recognition using complex high-resolution range profiles,” *IET Radar, Sonar & Navigation*, vol. 1, no. 1, pp. 18–26, 2007.
- [107] Z. Huang, C. O. Dumitru, Z. Pan, B. Lei, and M. Datcu, “Classification of large-scale high-resolution sar images with deep transfer learning,” *IEEE Geoscience and Remote Sensing Letters*, vol. 18, no. 1, pp. 107–111, 2020.
- [108] O. Schumann, J. Lombacher, M. Hahn, C. Wöhler, and J. Dickmann, “Scene understanding with automotive radar,” *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 2, pp. 188–203, 2019.
- [109] N. Pandey, G. Duggal, and S. S. Ram, “Database of simulated inverse synthetic aperture radar images for short range automotive radar,” in *2020 IEEE International Radar Conference (RADAR)*. IEEE, 2020, pp. 238–243.
- [110] V. Petsiuk, “Rise: Randomized input sampling for explanation of black-box models,” *arXiv preprint arXiv:1806.07421*, 2018.
- [111] V. Petsiuk, R. Jain, V. Manjunatha, V. I. Morariu, A. Mehra, V. Ordonez, and K. Saenko, “Black-box explanation of object detectors via saliency maps,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 443–11 452.

- [112] Z. Huang, Y. Liu, X. Yao, J. Ren, and J. Han, “Uncertainty exploration: toward explainable sar target detection,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–14, 2023.
- [113] M. Tan and Q. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [114] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [115] N. Pandey and S. S. Ram, “Explainable artificial intelligence based classification of automotive radar targets,” in *2023 IEEE Radar Conference (RadarConf23)*. IEEE, 2023, pp. 1–6.
- [116] A. V. Looveren and J. Klaise, “Interpretable counterfactual explanations guided by prototypes,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2021, pp. 650–665.
- [117] J. Han, J. Pei, and H. Tong, *Data mining: concepts and techniques*. Morgan kaufmann, 2022.
- [118] L. Wang, Y. Zhang, and J. Feng, “On the euclidean distance of images,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 27, no. 8, pp. 1334–1339, 2005.
- [119] T. Dozat, “Incorporating nesterov momentum into adam,” 2016.

- [120] A. Tato and R. Nkambou, “Improving adam optimizer,” 2018.
- [121] A. U. Nambi, I. Mehta, A. Ghosh, V. Lingam, and V. N. Padmanabhan, “Alt: towards automating driver license testing using smartphones,” in *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*, 2019, pp. 29–42.
- [122] T. Lin, H. Rivano, and F. Le Mouël, “A survey of smart parking solutions,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 12, pp. 3229–3253, 2017.
- [123] V. Paidi, H. Fleyeh, J. Håkansson, and R. G. Nyberg, “Smart parking sensors, technologies and applications for open parking lots: a review,” *IET Intelligent Transport Systems*, vol. 12, no. 8, pp. 735–741, 2018.
- [124] D. Fernández-Llorca, I. García-Daza, A. Martínez-Hellín, S. Álvarez-Pardo, and M. Á. Sotelo, “Parking assistance system for leaving perpendicular parking lots: Experiments in daytime\nighttime conditions,” *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 2, pp. 57–68, 2014.
- [125] L. T. Authority, “Handbook on vehicle parking provision in development proposals,” 2005.
- [126] M. G. Amin, *Through-the-wall radar imaging*. CRC press, 2017.
- [127] D. J. Daniels, “Surface-penetrating radar,” *Electronics & Communication Engineering Journal*, vol. 8, no. 4, pp. 165–182, 1996.

- [128] A. Sume and A. Örbom, “See-around-corners with coherent radar. x-band measurements at lilla gåra, august 2007,” *Swedish Def. Res. Agency, Linköping, Sweden, Tech. Rep. FOI*, 2007.
- [129] O. Rabaste, E. Colin-Koeniguer, D. Poullin *et al.*, “Around-the-corner radar: detection of a human being in non-line of sight,” *IET Radar, Sonar & Navigation*, vol. 9, no. 6, pp. 660–668, 2015.
- [130] M. Gustafsson, Å. Andersson, T. Johansson *et al.*, “Extraction of human micro-doppler signature in an urban environment using a “sensing-behind-the-corner” radar,” *IEEE Geoscience and Remote Sensing Letters*, vol. 13, no. 2, pp. 187–191, 2015.
- [131] A. Aubry, A. De Maio, and M. Rosamilia, “Reconfigurable intelligent surfaces for n-los radar surveillance,” *IEEE Transactions on Vehicular Technology*, vol. 70, no. 10, pp. 10 735–10 749, 2021.
- [132] G. Gennarelli and F. Soldovieri, “A linear inverse scattering algorithm for radar imaging in multipath environments,” *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 5, pp. 1085–1089, 2013.
- [133] O. Rabaste, J. Bosse, D. Poullin *et al.*, “Detection–localization algorithms in the around-the-corner radar problem,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 6, pp. 2658–2673, 2019.
- [134] S. Vishwakarma, A. Rafiq, and S. S. Ram, “Micro-doppler signatures of dynamic humans from around the corner radar,” in *2020 IEEE International Radar Conference (RADAR)*. IEEE, 2020, pp. 169–174.

- [135] V. C. Chen, F. Li, S.-S. Ho *et al.*, “Micro-doppler effect in radar: phenomenon, model, and simulation study,” *IEEE Transactions on Aerospace and electronic systems*, vol. 42, no. 1, pp. 2–21, 2006.
- [136] D. Barry, E. Coyle, D. Fitzgerald *et al.*, “Single channel source separation using short-time independent component analysis,” in *Audio Engineering Society Convention 119*. Audio Engineering Society, 2005.
- [137] J. Kolter, S. Batra, and A. Ng, “Energy disaggregation via discriminative sparse coding,” *Advances in neural information processing systems*, vol. 23, 2010.
- [138] S. Vishwakarma and S. S. Ram, “Detection of multiple movers based on single channel source separation of their micro-dopplers,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 1, pp. 159–169, 2017.
- [139] —, “Classification of multiple targets based on disaggregation of micro-doppler signatures,” in *2016 Asia-Pacific Microwave Conference (APMC)*. IEEE, 2016, pp. 1–4.
- [140] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM journal on computing*, vol. 24, no. 2, pp. 227–234, 1995.
- [141] S. Vishwakarma and S. S. Ram, “Mitigation of through-wall distortions of frontal radar images using denoising autoencoders,” *IEEE Transactions on Geoscience and Remote Sensing*, 2020.
- [142] I. W. Selesnick, “Sparse signal restoration,” *Connexions*, pp. 1–13, 2009.