

THE JIDOKA SYSTEM FOR MULTIPLE-SENSOR TERRAIN CLASSIFICATION

**Michael N. Huhns, Benito Mendoza,
Alicia Ruvinsky, and Dale McCants**

**Center for Information Technology
University of South Carolina**

Technical Report CSE TR-2006-13

July 2006

TABLE OF CONTENTS

1	INTRODUCTION	5
	PART ONE: SYSTEM DESCRIPTION	6
2	SYSTEM SENSORS	6
2.1	Radar	6
2.2	Camera	9
2.3	Accelerometer	12
3	SETUP	14
4	USER INTERFACE	15
4.1	Classification Panel	15
4.2	Camera Panel	17
4.3	Radar Panel	18
4.4	Accelerometer Panel	20
5	THE ROAD TO TERRAIN CLASSIFICATION	21
5.1	Threshold Analysis with Systematic Data Samples	21
5.1.1	Distinguishing Gravel from Other Terrains	23
5.1.2	Wet Gravel vs. Dry Gravel	24
5.1.3	Distinguishing Wet Concrete from Dry Concrete, Wet Sand or Dry Sand	25
5.1.4	Distinguishing Dry Sand from Dry Concrete and Wet Sand	26
5.1.5	Distinguishing Dry Concrete from Wet Sand	27
5.1.6	The algorithm	27
5.2	Threshold Analysis with Real-World Data Samples	28
5.2.1	Red Variance (Indoors vs. Outdoors)	29
5.2.2	Red Average (Indoors vs. Outdoors)	30
5.2.3	PointC (Indoors vs. Outdoors)	31
5.2.4	PointA (Indoors vs. Outdoors)	32
5.2.5	Algorithm Adjustments for Real-World Data	33
5.2.6	Adding Dirt	35
5.3	Road Surface Detection using an Accelerometer	40
5.3.1	Sampling	40
5.3.2	The 10th of a second analysis	44
5.3.3	Time domain	45

5.3.4	Spectrum analysis	46
5.3.5	The classifier procedure	48
6	CLASSIFICATION	51
6.1	Preprocessing: Extracting data of interest from the sensors	51
6.2	Learning: Classification Rules	52
6.2.1	Training	52
6.2.2	Testing	52
6.3	Real-Time Terrain Classification	53
	PART TWO: SYSTEM LEARNING	54
7	TRAINING AND TESTING	54
7.1	Saving Data	54
7.2	Invoking the C4.5 Application	56
7.3	Creating Rules	57
7.4	Implementing Dynamically Generated Rules	59
	PART THREE: SYSTEM ASSESSMENT	61
8	RESULTS	61
8.1	The rule with highest confidence first	61
8.2	Voting	62
8.3	Weighted Rules	64
8.4	Normalized Weight	65
9	CONCLUSIONS	67
9.1	Future Improvements to Terrain Detection	67

FINAL REPORT

Summary of JIDOKA Project:

The JIDOKA Integrated Development Environment (IDE)

1 Introduction

The JIDOKA project's mission is to create a computer architectural framework consisting of electronic sensing components with the purpose of providing vehicles with enhanced structural awareness. The functional intention of JIDOKA is both preventative and restorative. Prevention will be founded on calculations of the statistical likelihood of future mechanical failure. These calculations will be influenced by the current structural integrity of the vehicle as well as the impact of the exigent terrains to which it is exposed. The non-deterministic nature of real-world real-time systems motivates the implementation of restorative solutions to unexpected problems. This will be enabled by a sensor-based awareness of existing mechanical failures based on variations in the vehicle's performance over time.

The JIDOKA system will be used for either warning operators of possible or existing problems, or to automatically correct or adjust for such problems. The sensors comprising the JIDOKA framework will acquire real world data to be processed and characterized in order to generate intelligent decision making. These decisions will result in the more efficient utilization of the vehicle while enabling a more reliable assurance of the current, as well as future, safety of the vehicle.

In this research, we design a road surface detection system that utilizes sensors for detecting the characteristics of the terrain on which a vehicle is driving. Terrain detection is an essential aspect of the greater JIDOKA vision due to the significant effects of terrain on the durability and reliability of the car's overall functional integrity. For example, wet pavement is notorious for its effect of greatly increasing the likelihood of tires slipping or sliding, resulting in dangerous, uncontrollable spins. If the vehicle is aware that it is on wet pavement, it can take preventative measures, beginning with informing the driver of the wet terrain and the need for caution. These preventative measures depend on the degree of autonomy of the vehicle, but may span from a passive action such as simple driver communication in order to allow the driver to react to the information, to a more assertive action such as full axel management or speed control in order to decrease the likelihood of tire slip on wet pavement.

This document consists of a detailed description of the terrain detection component of the JIDOKA architecture. Part One gives a thorough description of the physical system. It begins with Chapter 2, which details the individual hardware sensing components of the system and their specifications. Chapter 3 outlines the setup procedure and subsequent use of the system, followed by a detailed description of how to interact with the system graphically in Chapter 4. The system description part is concluded with a glimpse in Chapter 5 into the directions that led to the final classification procedure, and a brief introduction to the training and execution of this final classification mechanism of the system. This is explained in more detail in Part Two, Chapters 6 and 7. Part Three consists of results obtained from system executions, as well as conclusions about the capabilities and limitations of the terrain-detection system.

PART ONE: System Description

2 System Sensors

The road surface detection system component of the JIDOKA project consists of three sensing components: an accelerometer, a digital camera, and a pulse radar. Each of these sensors independently and concurrently captures surface terrain data. The camera and the radar's data sets correspond to the impending terrain on which the vehicle will be traversing momentarily, while the accelerometer senses the vertical acceleration of the vehicle upon terrain traversal.

The sensors regularly transmit captured data via a USB interface to the JIDOKA software component running on a central computing device (currently a laptop computer). The data is then processed and the current terrain classification is made with respect to the sensor-supplied input.

In the following sections, we describe in greater detail the specification of each sensing component and a brief overview of their functionality contribution to terrain detection.

2.1 Radar

The pulse radar was manufactured by ABM Sensor Technology model ABM300/400 100R. The radar used for this project has been specially modified by ABM Sensor Technology to provide transmit power adjustability for pulse width and provide an analog echo profile for dynamic analysis. The radar is mounted to the front of the vehicle via a mounting bracket made specifically for the requirements of this project. The height of the tip of the radar antenna must be between 48-49 inches from the road surface in order to separate the transmission profile from the echo profile. The height is also based on the voltage (selected after extensive testing) in order to provide the best results with this particular radar. (Changes in voltage cause variance in the specific pulse width generated by the radar.) The radar must be mounted normal (perpendicular) to the surface for best results, in order to receive a strong echo profile. The analog signal from the radar is fed to the Analog Devices AD9223 analog-to-digital converter, which has a USB interface connection to the laptop computer.

Table 1. Radar specifications

RADAR GENERAL SPECIFICATIONS	
Conduit Entry	1/2" NPT
Enclosure	Aluminum or Stainless Steel
Enclosure Rating	NEMA 4 (IP65)
Antenna	Polypropylene Dielectric Rod Standard (Optional PTFE)
Ambient Operating Temperature	-40° to +140° F (-40° to 60° C)
Process Operating Temperature (PP Rod):	-40° to 190° F (-40° to 90° C)
Pressure	1-10 bar / 15-150 psi
Approvals	FCC Part 15—Low Power Communication Device
Accuracy	+/- 0.25% of maximum Target range (In Air)
Loss of Echo	Hold 30 seconds, 22 mA Output Time
Calibration	Push-button or programmable via optional communication port
Diagnostics	Via communication port (echo profile, echo stability, operation errors)
Power AC	AC units 115 VAC 60Hz or 230 VAC 50Hz, 1.7 VA
Power DC	DC units 12 to 30 VDC, 0.07 A max @ 24 VDC
Transmitter Power	50 uW Average
Output	4-20 mA, optional RS-232, RS-485, or Modbus RTU
4 to 20 mA Max. Loop Resistance	110 VAC @ 750 Ohms (isolated) 12 VDC @ 250 Ohms 24 VDC @ 750 Ohms
Frequency	6.3 GHz




Figure 1. Radar component identification image.




Figure 2. Radar and mount assembly.




Figure 3. Radar and mount assembly from side.

Analog-to-Digital Converter

Analog Devices manufactured the A/D converter used to interface the radar with the laptop computer. The A/D converter consists of two separate boards that connect together to perform the conversion and interface the computer. The Analog Devices model AD9223 monolithic converter interfaces the radar via coaxial cable. The AD9223 board also contains a specially built timing encoding oscillator, which operates at 1.2288 MHz. The timing encoder functions to tell the AD9223 the rate at which to cut up and process the incoming analog signal. The AD9223 connects directly to the Analog Devices High Speed ADC USB Evaluation board; model HSC-ADC-EVALA-SC FIFO 4 Rev 2A. The HSC-ADC-EVALA-SC interfaces with the laptop computer via a USB connection. Due to the large amount of detail for the specifications of these two boards, their specifications are not included in this document, but can be found on the Analog Devices Web site www.analog.com.

2.2 Camera

The digital camera used in the JIDOKA terrain detecting system was manufactured by Arcvision Technology Corp. The model number of the camera is ARC-12B49. The camera is mounted to the front of the vehicle on the same specially made mounting bracket as the radar. The height of the camera is approximately 60 inches from the road surface. The mounting position was chosen out of convenience, because the mounting bracket made for the radar provided a logical placement for the camera. The ARC-12B49 is a high quality color ccd camera that is weatherproof and has night vision (full camera specifications are shown in Table 2). The signal from the ARC-12B49 connects with the laptop computer via a PCI bus video card. The input from the camera is then dynamically read into the Jidoka Software to assist in road surface detection.

Table 2. Camera Specifications

ARC-12B49 CAMERA SPECIFICATIONS	
Model Number	ARC-12B49
Image Pick-up Device	Color SONY 1/3" Super HAD CCD Sensor
Effective Picture Elements	NTSC: 510 x 492, PAL: 500 x 582 (HxV)
Horizontal Resolution	380 TV Lines
Minimum Illumination	0.5 LUX @ F2.0
S/N Ratio	More than 48 dB
Auto Electronic Shutter	NTSC: 1/60s-1/100,000s, PAL: 1/50s-1/110,000s
Auto Gain Control	Yes
Auto White Balance	Yes (2500K-9500K)
Scanning System	Interlace 2:1
Gamma Characteristic	0.45
Lens Furnished	Board Lens 3.6mm / F2.0
Synchronous System	Internal, Negative Sync.
Back Light Compensation	Yes
Video Output	1 Vp-p / 75 Ohms
Power Consumption	1.5 W (CCD Camera only)
Infrared Luminary	12 pieces IR LED, Life: 6000 hrs.
Wave Length	850 nm
Beam Spread	35°
Projection Distance	Up to 10 meters
Auto Power ON/OFF	Cds auto control
Illumination Active Limit	Under 10 LUX
Power Consumption	1.3 W (IR-LED ON only)
Power Supply	DC 12V \pm 10% / 220 mA min. (Camera+IR-LED ON)
Waterproofing Criterion	Aluminum case, IP 67 waterproof design
Operating Temp	-10C – 60C (14F – 122F)



Figure 4. Image of camera ARC-12B49.




Figure 5. Dimensions of camera ARC-12B49.

2.3 Accelerometer

The accelerometer is manufactured by Techkor, with a model number of M9E-USB. This component is mounted to the lower control arm of the vehicle, on the passenger side of the vehicle, via a specially made bracket. This mounting location was chosen so that the only added damping effect between the accelerometer and the road is the tire. Tire pressure will have an effect on the magnitude of the sensed acceleration but only to a relatively small degree, assuming the tire is kept in its suggested pressure range. The M9E-USB is an accelerometer and data logger in a single package, which communicates to a host device through a standard USB connection. The unit contains a piezoelectric accelerometer, constant current power source, amplifier, filter, digital signal processor, data memory, and USB transceiver. See Figure 6 for a detailed specification of the accelerometer. The data collection parameters are configured from the Jidoka software.

Input Specifications

- Piezoelectric Sensor
- 35g full scale range
- 0.4 Hz to 12,000 Hz Frequency Response

DSP Specifications

- 125 Hz to 65K Hz Sampling Rate
- 16 bit resolution
- 128K byte Data Storage
- Programmable Data Form
 - Time Waveform
 - Enveloped Time Waveform
 - 50 - 1,000 Hz band
 - 500 - 10,000 Hz band

Communication Specifications

- USB 1.1 (low speed device, 3m cable)
- USB 2.0 (full speed device, 5m cable)
- MIL-C-26482 Connector (four pin)

General Specifications

- Power over USB
- Hermetically Sealed to IP
- Corrosion Resistant 316 Stainless
- 1/4"-28 Thread Mount in Base
- Case Isolated
- ESD Protection




Figure 6. Accelerometer Specification

The accelerometer is intended to distinguish between smooth terrains, such as pavement, and texturally rougher terrains, such as gravel. Due to the visual similarity of different surfaces, the camera sensor was found to be not accurate enough to distinguish between pavement and gravel for all forms of gravel. The texture dimension is consistently distinguishable for all forms of gravel and all forms of pavement. Yet, acceleration introduces a speed variable which is provided to the system via a wheel speed sensor. The JIDOKA classification system uses the accelerometer input coupled with the speed input to assist in minimizing overlapping characterizations of pavement and gravel.

The frequency range of the accelerometer was chosen to be able to detect gravel of nominal size 0.3 cm (~1/8") when the vehicle is traveling at 120 kph (33.3 meters/second). At this speed, the vibration, along a vertical axis, caused by gravel of this size would be approximately 10 KHz. To capture this would require a sensor with an input upper frequency response greater than 10 KHz, which the chosen accelerometer has. In order to characterize this signal as accurately as possible, it must be sampled at greater than the Nyquist frequency, requiring a DSP operating at greater than 20KHz. The chosen accelerometer satisfies this requirement as well.

3 Setup

Physical Connections

Analog Devices AD9223 Board:

- J3 to -5V (red wire)
- J6 to -12V and -5V (two black wires)
- J2 to +12V (green wire)
- J5 to -5V (green wire)
- J4 to +5V (orange wire)
- J1 (BNC cable) to radar unit

Analog Devices ADC-EVALA-SC FIFO 4 Rev 2A Board:

- J309 to 3.3V Regulated Power Supply
- USB Male A to laptop USB Male B

Power Connections to Radar:

- +24V labeled wire to CSI/SPECO Regulated Power Supply Red Terminal
- -24V labeled wire to CSI/SPECO Regulated Power Supply Black terminal
- +13V labeled wire to Tenma DC Regulated Power Supply
- -13V labeled wire to Tenma DC Regulated Power Supply

Connections to Camera:

- Camera signal to Video1 port on PCMCIA card of laptop
- Camera power to power supply connector

Connections to Accelerometer:

- Single USB connector to USB port on laptop

Power-Up Procedure

1. Plug connector into cigarette lighter receptacle
2. Turn on power inverter
3. Turn on 24V DC power supply
4. Turn on Tenma DC power supply
5. Set current on Tenma power supply to maximum (max. clockwise)
6. Set voltage on Tenma power supply to 13.0V
7. Boot laptop and run the Jidoka-classifier program


4 User Interface

A user interacts with the system via a graphical user interface consisting of individual panels for each sensing component, as well as a primary panel on which classification is presented visually. The panels are intended to illustrate the functional processes of each sensor as well as the classification system. This facilitates the understanding of what is going on in the system, as well as providing a monitor for observing and maintaining the proper execution of the software.


In this section, we describe the interface of each of the four panels comprising the JIDOKA interface. We also briefly detail what the user will observe upon proper execution of the system.

4.1 Classification Panel


The classification panel's main component is a graph on which the current terrain classification is represented. A LED located next to the classified terrain type illuminates in order to indicate the description of the specific terrain being traversed, as well as the kind of precipitation (if any) that is affecting the vehicle's interaction with the terrain surface. The terrain type graph also maintains a short history of recent terrain and precipitation classifications in order to provide a slightly larger contextual awareness to the operator. Figure 7 contains images depicting the classification of various terrains and precipitations.



(a) Snow on pavement



(b) Dirt and no precipitation



(c) Gravel and no precipitation

Figure 7 (a), (b), (c). Classification Panel GUI for all three terrain types with varying precipitations

4.2 Camera Panel


The camera panel displays an image of the impending terrain on which the vehicle will traverse momentarily. Though slight, the time between sensing the approaching terrain and the actual traversal will provide the vehicle with the opportunity to mechanically prepare for any potentially dangerous situations. For example, if the camera senses snow on the road, the car will inform the driver of the hazardous conditions as well as manage speed level in an attempt to prevent the vehicle from sliding. See Figure 8 for an example of the camera panel.




Figure 8. Camera Panel GUI capturing approaching gravel terrain.

4.3 Radar Panel


The radar panel contains a graph that exhibits the waveform received by the pulse radar. The pulse radar transmits a pulse towards the ground and then receives the echoed transmission resulting from the pulse reflecting from the terrain. Different terrain types and precipitation types will reflect the pulse in distinguishable ways. In Figure 9 below, the echoed waveforms from different terrains and precipitations are shown.



(a) Snow on pavement



(b) Gravel and no precipitation



(c) Dirt and no precipitation

Figure 9 (a), (b), (c). Radar Panel GUI for all three terrain types with varying precipitations

4.4 Accelerometer Panel

The accelerometer panel graphs the changes in the vertical velocity of the vehicle as observed by the accelerometer. The waveforms generated by the accelerometer are representative of the texture of the terrain on which the vehicle is being driven. For example, pavement will result in significantly less vertical displacement at 25 mph than gravel.

In order to adequately analyze the data generated by the accelerometer, the speed of the vehicle at the time of the accelerometer data capture is required. This is necessary because vertical displacement is amplified with speed. Without an awareness of the speed of the vehicle, confusion will arise between, for example, pavement at 60 mph and gravel at 15 mph. Though pavement is generally smooth, high speeds will generate vertical displacements similar to that of a slower speed on a rough terrain such as gravel. Figure 10 shows a number of accelerometer data captures from various terrains at different speeds.




Figure 10. The accelerometer panel shows the values sent by the accelerometer.

5 The Road to Terrain Classification

Many obstacles were overcome in attempting to design a provably satisfactory classification algorithm for the terrain and precipitation requirements specified for the JIDOKA project. Much may be learned from an analysis of the route we took in reaching our final classification procedure.

5.1 Threshold Analysis with Systematic Data Samples

Our initial approach used the radar and the camera attributes (Wave Width, Wave Amplitude, and the RGB Average and Variance) to compute the Euclidean distance between a new instance and the existing classes (using 2 or more attributes combined). Using this approach, the average accuracy was 82.85% as shown in Table 3. All the combinations were tried and the best are presented below. Table 3 depicts the classification results from using Euclidean distance on data acquisition from the radar and camera. Table 4 illustrates a confusion matrix expressing the satisfactory results generated from using the Euclidean distance of red average and red variance to discriminate gravel from other terrain types.

Table 3. Classification results from radar and camera data acquisition using Euclidean Distance.

Using Both, Camera (RGB scaled) and Radar							
	Wet Concrete	Dry Concrete	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Accuracy
Wet Concrete	71	9	0	0	25	0	67.61905
Dry Concrete	0	100	0	0	4	1	95.2381
Wet Gravel	0	0	79	26	0	0	75.2381
Dry Gravel	0	0	20	85	0	0	80.95238
Wet Sand	0	1	0	0	103	1	98.09524
Dry Sand	0	15	0	0	6	84	80
						Average Accuracy	82.85714

Gravel is the only terrain that can be classified perfectly, since the classifier does not misclassify gravel and nothing else is misclassified as gravel, as seen in Table 3. However, this can be attributed to the red values from the Camera. Table 4 shows how it is possible to get this perfect accuracy on the gravel classification using only red average and red variance. Decreasing the dimensionality of classification as we did here with gravel is useful because the combined use of many attributes can be self-defeating, since some attributes can reduce the efficiency of others.

Table 4. Satisfactory discrimination of gravel using only red average and red variance.

	Wet Concrete	Dry Concrete	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Accuracy
Wet Concrete	97	8	0	0	0	0	92.380
Dry Concrete	12	52	0	0	24	17	49.523
Wet Gravel	0	0	103	2	0	0	98.095
Dry Gravel	0	0	5	100	0	0	95.2381
Wet Sand	14	42	0	0	12	37	11.428
Dry Sand	0	24	0	0	17	64	60.952
						Average Accuracy	67.936

In these tests, the eight attributes used for classification are listed below:

1. wave width
2. wave amplitude
3. red average
4. red variance
5. green average
6. green variance
7. blue average
8. blue variance

The categories for which classification was defined are as follows:

1. wet concrete
2. dry concrete
3. wet gravel
4. dry gravel
5. wet sand
6. dry sand

Based on the confusion matrices in Table 3, the results were very low. Gravel was the only category classified satisfactorily using the red variance. It was obvious that Euclidean distance was not good enough, and that new features were needed. The next approach was to perform threshold analysis to find values for attributes that could be used to distinguish between categories. The new attributes were identified by eyeballing the graph in Figure 11. Three attributes, PointA, PointB, and PointC, were chosen based on visual assessment. and selecting points that provided the greatest distinction between categories. These were defined as follows:

- **PointA** is the point 160 of the wave form
- **PointB** is the point 240 of the wave form
- **PointC** is the point 360 of the wave form

The results in Figure 11 show the category separation provided by each point.

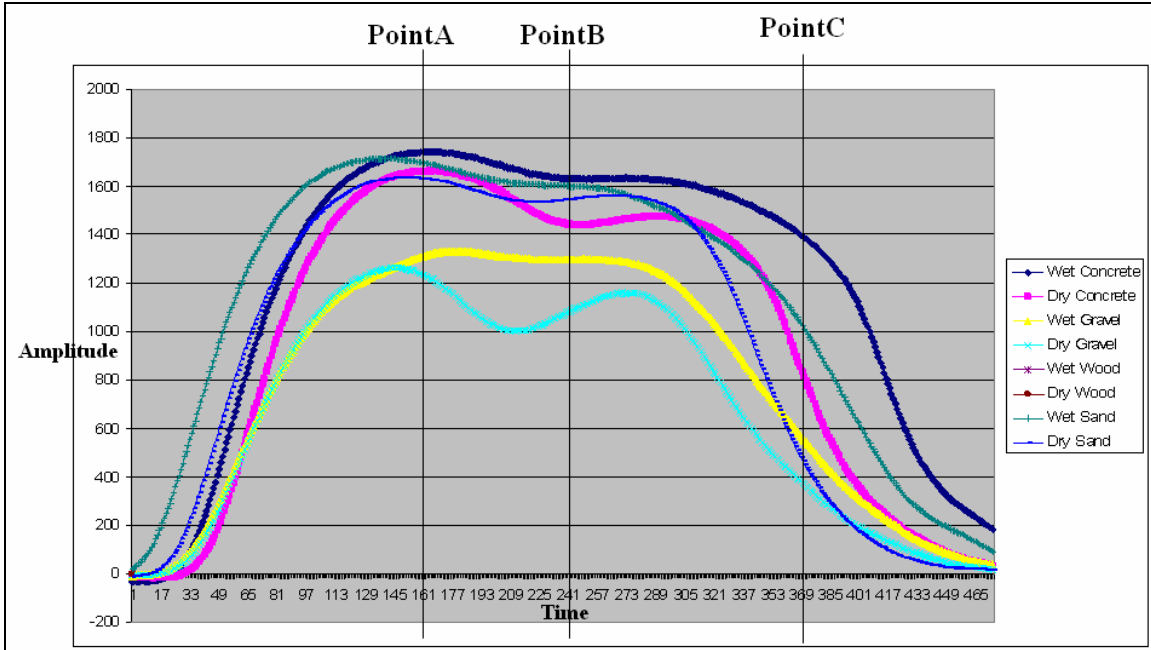


Figure 11. Three attributes, PointA, PointB, and PointC, were chosen based on visual assessment.

An analysis of the each individual attribute was carried out. The goal was to find the attribute that provides greatest discrimination of terrain and precipitation. The following sections detail our findings.

5.1.1 Distinguishing Gravel from Other Terrains

Distinguishing gravel from the other terrains is possible using the red variance of the camera pictures and a threshold of 3439.7183 (determined experimentally). Table 5 shows the confusion matrix and Figure 12 gives a better view of the difference among the terrains. The set of points circled in green corresponds to gravel (wet and dry). The set on the left is concrete and the one on the right is sand.

Table 5. Confusion matrix generated by using threshold analysis on red variance to classify concrete, gravel, and sand.

RedVariance				
	Concrete	Gravel	Sand	Accuracy
Concrete	112	0	98	53.33333
Gravel	0	210	0	100
Sand	18	0	192	91.42857




Figure 12. Results from using threshold analysis on red variance to classify gravel. Gravel values (dry and wet) are circled in green.

Red variance can discriminate gravel from the other terrains with 100 percent accuracy. The distinguishing threshold was derived by taking the average of the lowest red variance value for a gravel sample and the highest value for a sample in the both the concrete and sand categories. This provides enough of a buffer for adequate distinction between gravel and the other terrains. Thus, anything having a red variance bigger than 3439.7183, can be classified as gravel. However red variance cannot distinguish between wet gravel and dry gravel.

5.1.2 Wet Gravel vs. Dry Gravel

Red average is the best attribute to distinguish between wet gravel and dry gravel. However, it is not possible to reach 100 percent accuracy. Figure 13 shows the results of using threshold analysis of red average to distinguish wet and dry gravel. The wet gravel values are on the left side of the graph circled in green, while the values on the right represent dry gravel. The threshold marked by the red line is a red average value of 282.6675. Using red average, there is 98 percent accuracy for telling wet gravel from dry gravel.

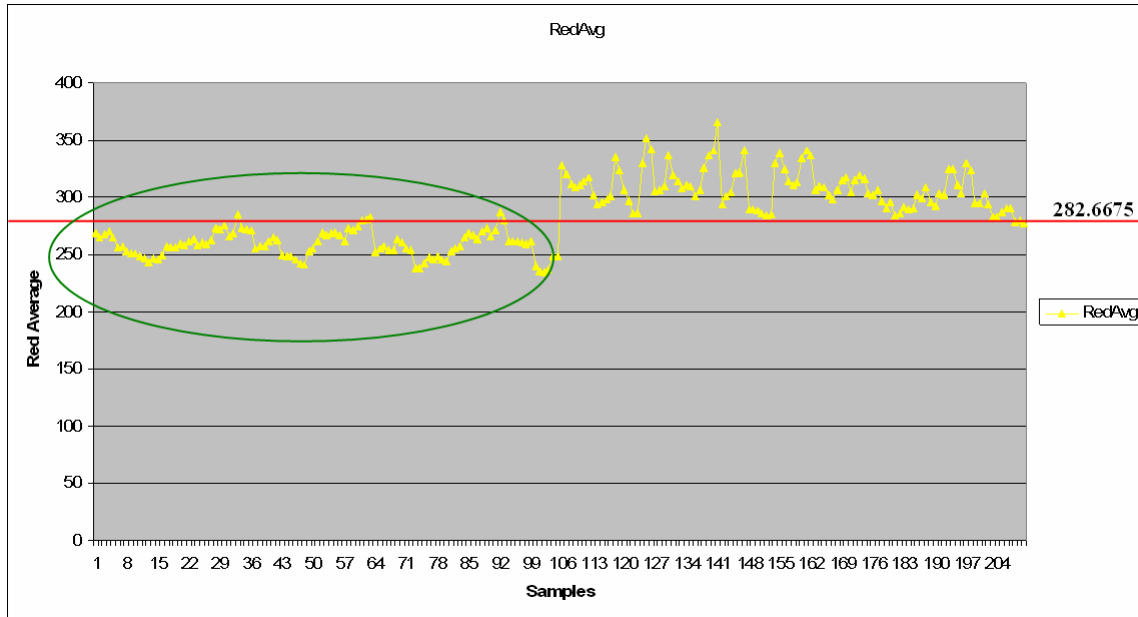


Figure 13. Results of using threshold analysis of red average to distinguish wet and dry gravel. The wet gravel values are on the left side of the graph circled in green. The values on the right represent dry gravel. The threshold marked by the red line is a red average value of 282.6675. This represents a 98% accuracy for telling wet gravel from dry gravel.

5.1.3 Distinguishing Wet Concrete from Dry Concrete, Wet Sand or Dry Sand

PointC is the attribute that best separates wet concrete from the other classes. Figure 14 shows the differences between wet concrete and the other classes based on PointC. In this figure, the wet concrete samples are circled. Using the discriminating threshold of 1575.233, wet concrete can be classified with 100 percent accuracy.




Figure 14. PointC is used to distinguish wet concrete from dry concrete or sand. Wet concrete values are circled in green. The remaining values represent dry concrete and sand (wet and dry). The threshold is 1575.233. The accuracy is 100% for wet concrete.

5.1.4 Distinguishing Dry Sand from Dry Concrete and Wet Sand

The best attribute to discriminate dry sand from the other terrains, dry concrete and wet sand, is PointA. All the samples of dry sand have a PointA value of less than 1600 as shown in Figure 15, such that PointA provides an accuracy reading for distinguishing dry sand of 100 percent.




Figure 15. PointA is used to distinguish dry sand from dry concrete or wet sand. Dry sand values are circled in green. The remaining values represent dry concrete and wet sand. The threshold is 1600. The accuracy is 100%.

5.1.5 Distinguishing Dry Concrete from Wet Sand

To distinguish between dry concrete and wet sand, the discriminating attribute is also PointA. Most of the dry concrete samples have a PointA value of below 1660.613. Figure 16 shows the difference between these dry concrete and wet sand. There is 98 percent accuracy for this distinction.




Figure 16. PointA is used to distinguish dry concrete from wet sand. Dry concrete values are circled in green. The remaining values represent wet sand. The threshold is 1660.613. The accuracy is 98%.

5.1.6 The algorithm

A classification procedure can be created based on the information obtained from the analysis of the attributes. The algorithm encompassing all the thresholds specified so far is described by the following flow chart:

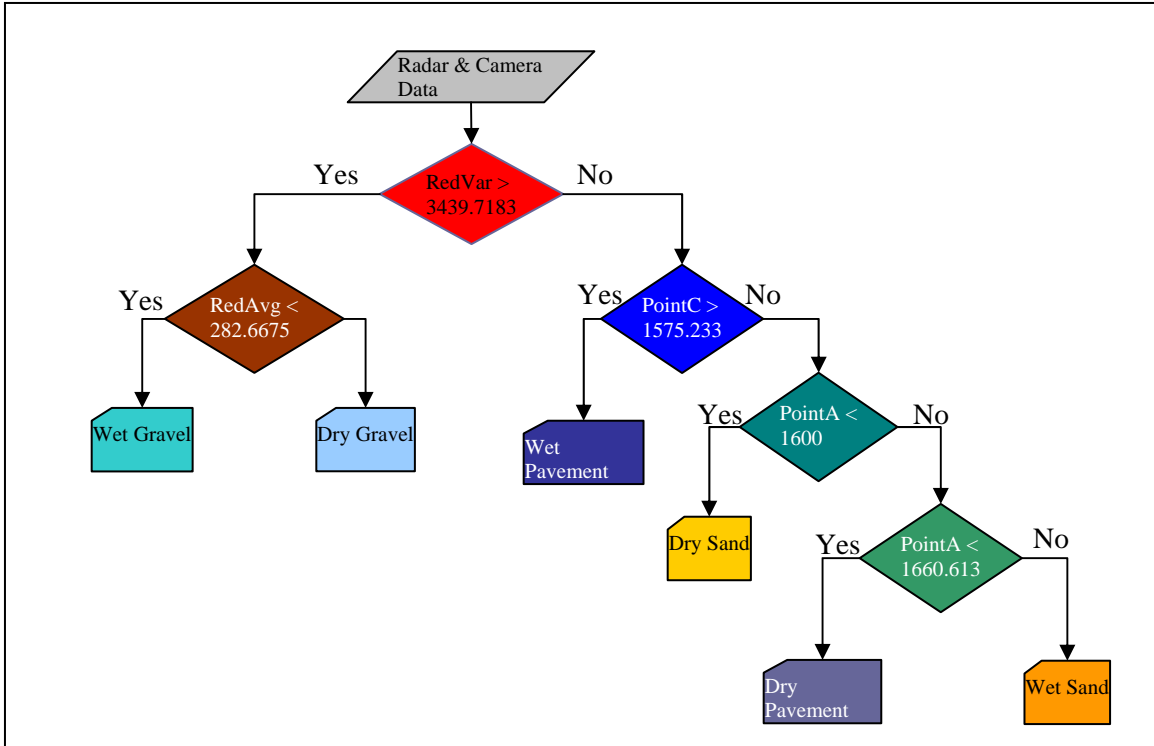


Figure 17. Algorithm describing the initial threshold discriminator, *Threshold_1*.

5.2 Threshold Analysis with Real-World Data Samples

Previously we presented an approach to classify pavement, gravel, and sand (either wet or dry). We will refer to that approach as *Threshold_1*. The samples or data used in generating *Threshold_1* as presented in the previous section were taken inside (indoors) of a lab (building) in an attempt to generate controlled, systematic data. The lab had artificial illumination and was surrounded by walls, thus creating significant discrepancies between this data and real-world samples. In this section, we present the results of using the *Threshold_1* approach with data that was taken outdoors. Since the results of using *Threshold_1* to classify outdoor data have lower accuracy than when used with indoor data, we present two modifications to increase the accuracy of classification with outdoor data.

Table 6. Results obtained from using *Threshold_1* with real-world outdoor data. There are 108 samples for each road surface/precipitation pairing.

	Wet Pavement	Dry Pavement	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Accuracy
Wet Pavement	0	8	0	57	6	37	0
Dry Pavement	0	0	0	0	0	108	0
Wet Gravel	0	0	0	108	0	0	0
Dry Gravel	0	0	0	104	0	4	96.2963
Wet Sand	0	18	0	0	3	87	2.77777
Dry Sand	0	38	0	0	8	62	57.4074
						Average Accuracy	26.0802

Table 7. Results obtained using *Threshold_1* with systematic indoor data. There are 105 samples of each surface/precipitation pairing.

	Wet Pavement	Dry Pavement	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Accuracy
Wet Pavement	105	0	0	0	0	0	100
Dry Pavement	0	103	0	0	2	0	98.09524
Wet Gravel	0	0	102	3	0	0	97.14286
Dry Gravel	0	0	3	102	0	0	97.14286
Wet Sand	0	2	0	0	103	0	98.09524
Dry Sand	0	0	0	0	0	105	100
						Average Accuracy	98.4127

Table 6 shows the results of using *Thresholds_1* to classify outdoor data. This may be compared to Table 7 which shows the results of using *Thresholds_1* to classify indoor data. The accuracy of this approach to classify real-world terrains is very low. Dry gravel is the only one that was recognized with a good accuracy (96.2963%). The resultant accuracy of the rest of the terrains is very low and some of them are 0 percent. The reason for those results is that the outdoor data is very different from the indoor data.

The following four subsections contain a brief analysis of the attributes used by the procedure with both indoor and outdoor data. Subsection 5.2.1 contains a brief analysis of the red variance attribute originally, used to distinguish gravel from the rest. Subsection 5.2.2 contains a brief analysis of the red average attribute, used previously to distinguish wet gravel from dry gravel. Subsection 5.2.3 discusses the PointC attribute, used previously to distinguish wet pavement from dry sand, wet sand and dry pavement. Subsection 5.2.4 is about PointA, the attribute used to distinguish dry sand from wet sand and dry pavement, and dry pavement from wet sand. Subsection 5.2.5 presents a modification to *Thresholds_1* to classify outdoor data; the modified procedure is called *Thresholds_Outdoors*. Subsection 5.2.6 presents a modification to classify outdoors data including dry and wet dirt.

5.2.1 Red Variance (Indoors vs. Outdoors)

Recall Figure 12. Results from using threshold analysis on red variance to classify gravel. Gravel values (dry and wet) are circled in green.), which shows the red variance of all the indoors terrains. When using the red variance with indoor data, it is very easy to distinguish the gravel samples from the others, since all the gravel samples have a red variance greater than 3439.7183. But the case is not the same with outdoors samples. Figure 18 shows the red variance of samples taken outdoors. Most of the outdoor samples of gravel are above 3439.7183, however some of the wet pavement samples are above that threshold too. In other words, in this case red variance cannot completely discriminate gravel from the other terrains since it sometimes gets confused with wet pavement.




Figure 18. Results for using red variance threshold value of 3439.7183 for identifying gravel samples. The circled set is gravel (wet and dry). The set to the left for the circle is pavement, and the one on the right is sand.

5.2.2 Red Average (Indoors vs. Outdoors)

Originally, the best attribute to distinguish between wet gravel and dry gravel was red average. Recall Figure 13 showing the discrimination capacity of red average with indoor data. Using the threshold 282.6675 it was possible to get 97.14286 percent accuracy for both wet gravel and dry gravel (indoors). However this attribute is not good with outdoors data. Figure 19 not only shows that both wet gravel and dry gravel (outdoors) are above that threshold, but it also shows that there is no good threshold to use.

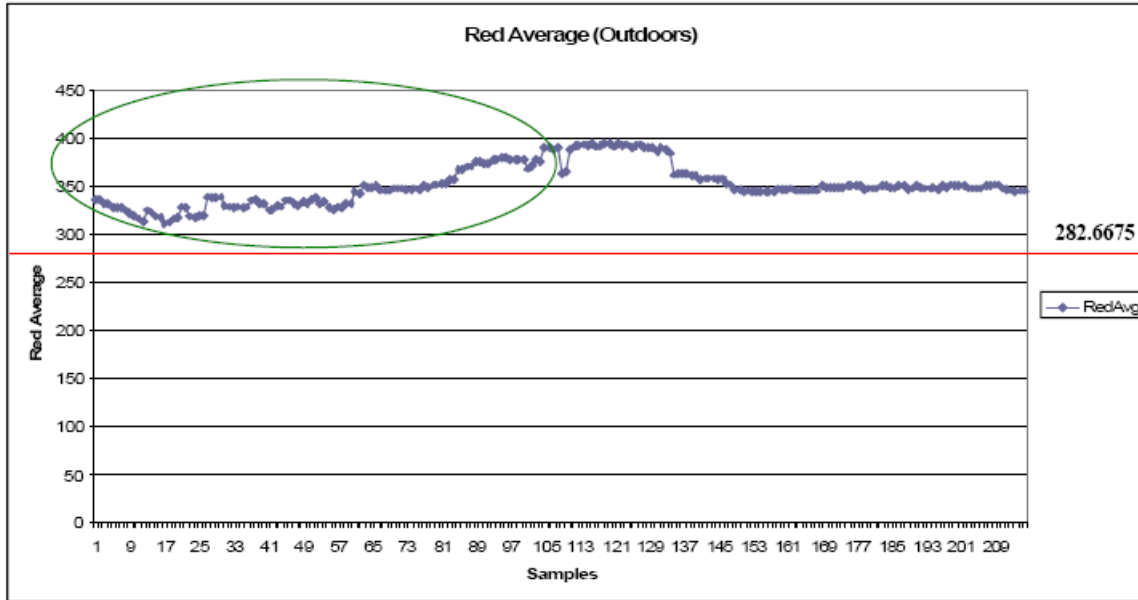


Figure 19. Results for using red average threshold value of 282.6675 for discriminating wet and dry gravel samples. The threshold of 282.6675 derived from indoor samples is below all of the outdoor samples. It is therefore not possible to distinguish wet gravel from dry gravel using this attribute. The circled set is wet gravel. The set to the right is dry gravel.

5.2.3 PointC (Indoors vs. Outdoors)

PointC is the attribute that best discriminates wet pavement from the other terrains. Recall Figure 14 which shows the differences among the wet pavement and the other terrains on PointC for indoor data. Everything with a PointC value greater than 1575.233 can be considered wet pavement. However, this attribute does not work very well for outdoor data. Figure 20 shows that all the samples are below the threshold, and there is no threshold that we can use.




Figure 20. PointC is not a good choice for discrimination when using outdoor data. The wet pavement set is the circled one. The sets on the right are the other terrains. All the samples are below the threshold (1575.233).

5.2.4 PointA (Indoors vs. Outdoors)

The best attribute to discriminate dry sand from the other terrains (except gravel, which we could distinguish using the red variance) is PointA. All the samples of dry sand have a PointA less than 1600. Again, this attribute with that threshold does not work with outdoor data. Figure 21 shows that for outdoor data, most of the dry pavement samples and some of the wet sand, and some of the dry sand samples are above it.




Figure 21. The PointA attribute with a threshold of 1600 is not good to distinguish dry sand from the others when using outdoor data. The dry sand set is the circled one. The samples on the left with a PointA greater than 1600 are wet pavement.

To distinguish between dry pavement and wet sand, the best discriminating attribute is PointA again for indoor data. Most of the dry pavement samples are below 1660.613.




Figure 22. The PointA and the 1660.613 threshold are not good to distinguish dry pavement from wet sand. The dry pavement set is the circled one. The set on the right is wet sand.

Figure 22 shows that for outdoor data both dry pavement and wet sand are below 1660.613. So this threshold is not good for discriminating purposes.

5.2.5 Algorithm Adjustments for Real-World Data

The *Threshold_1* procedure needs some modifications to do a good job with outdoor data. Figure 23 shows a modification to this procedure that does better with outdoor data. This new procedure is called *Threshold_Outdoors*.

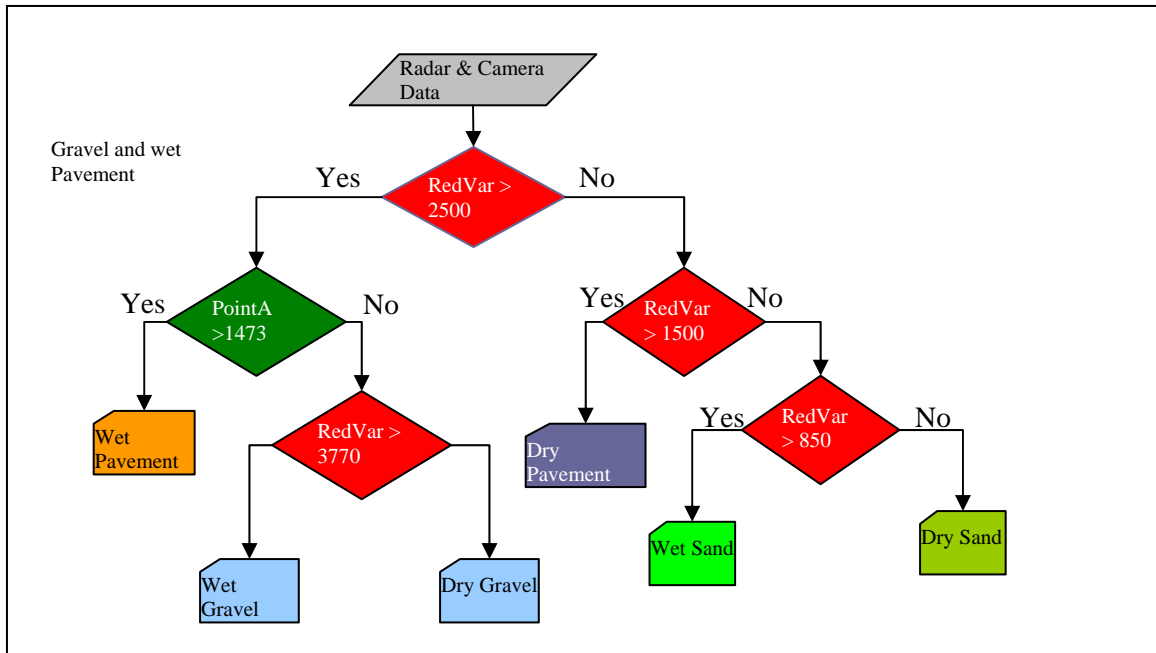


Figure 23. *Threshold_Outdoor*: Classification procedure for real-world data set.

Table 8 shows the results of using *Threshold_Outdoors* to classify outdoor data. This procedure works very well with outdoor data. The problem is that 10 percent of the wet gravel samples are classified as wet pavement and 26 percent of the dry gravel samples are classified as wet gravel. That occurs because some of those dry gravel samples have a red variance greater than 3770. Figure 24 illustrates this fact.

Table 8. results of using *Threshold_Outdoors* to classify outdoors data

	Wet Pavement	Dry Pavement	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Accuracy
Wet Pavement	108	0	0	0	0	0	100
Dry Pavement	0	108	0	0	0	0	100
Wet Gravel	10	0	98	0	0	0	90.74074
Dry Gravel	0	0	28	80	0	0	74.07407
Wet Sand	0	0	0	0	108	0	100
Dry Sand	0	0	0	0	0	108	100
						Average Accuracy	94.1358




Figure 24. Red Variance is a good attribute to distinguish wet gravel from dry gravel (outdoors). The circled samples are the dry gravel samples misclassified as wet gravel since they have a Red Variance greater than 3770. The set on the left is wet gravel and the one on the right is dry gravel.

5.2.6 Adding Dirt

A requirement of the JIDOKA system was to distinguish dirt from the other road surfaces. To distinguish dirt from the other terrains we made some modification to *Threshold_Outdoors*. The new procedure called *Threshold_Outdoors_D* is shown in Figure 25.

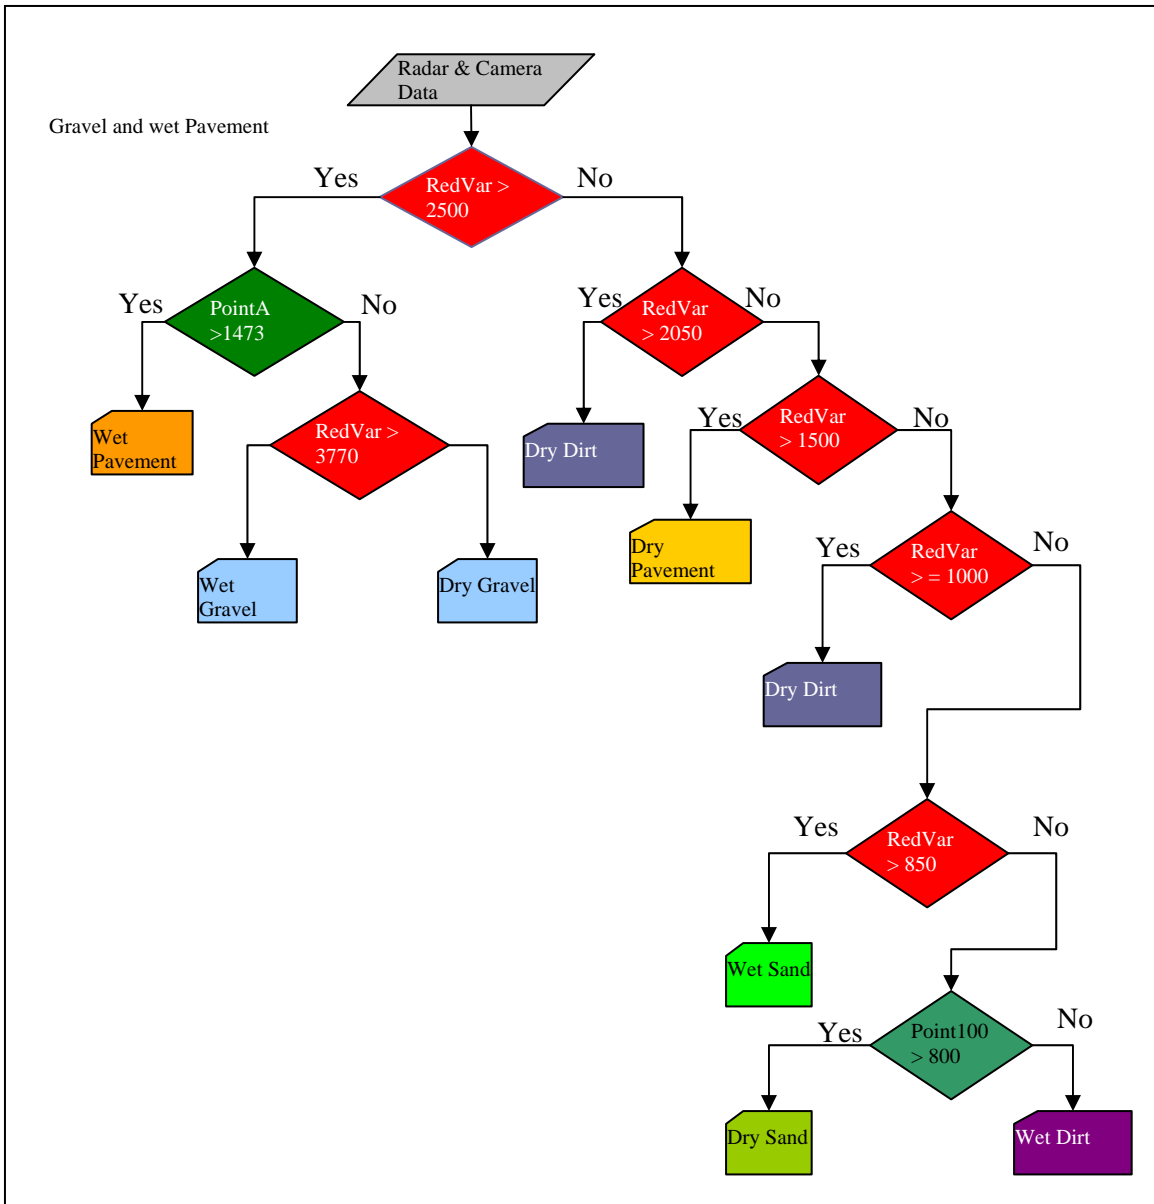


Figure 25. *Threshold_Outdoor_D*: Classification procedure for inclusion of Dirt samples.

In this new procedure, we use the red variance to distinguish dry dirt from sand and wet dirt and the Point100 attribute (a new attribute corresponding to the point 100 of the waveform from the radar) to distinguish dry sand from wet dirt. The results of this procedure with all the terrains are showed in Table 9. Results of using the *Threshold_Outdoors_D* procedure.) The results are not as good as when there were only three road surfaces. The problem is that for the radar, dirt is very similar to gravel, and for the camera, dirt is very similar to sand.

Table 9. Results of using the Threshold_Outdoors_D procedure.

	Wet Pave.	Dry Pave.	Wet Gravel	Dry Gravel	Wet Sand	Dry Sand	Wet Dirt	Dry Dirt	Accuracy
Wet Pave.	108	0	0	0	0	0	0	0	100
Dry Pave.	0	108	0	0	0	0	0	0	100
Wet Grav.	10	0	98	0	0	0	0	0	90.74074
Dry Grav.	0	0	28	80	0	0	0	0	74.07407
Wet Sand	0	0	0	0	99	0	0	9	91.66667
Dry Sand	0	0	0	0	0	101	7	0	93.51852
Wet Dirt	0	0	0	0	1	38	69	0	63.88889
Dry Dirt	6	1	22	4	2	0	0	73	67.59259
Average Accuracy									85.18519

Table 10 shows the results of this last procedure eliminating the sand category and samples. The results of eliminating the sand category are a little bit better, since all the wet dirt samples are classified correctly.

Table 10. Results of using the Threshold_Outdoors_D procedure without the Sand category and samples.

	Wet Pavement	Dry Pavement	Wet Gravel	Dry Gravel	Wet Dirt	Dry Dirt	Accuracy
Wet Pavement	108	0	0	0	0	0	100
Dry Pavement	0	108	0	0	0	0	100
Wet Gravel	10	0	98	0	0	0	90.74074
Dry Gravel	0	0	28	80	0	0	74.07407
Wet Dirt	0	0	0	0	108	0	100
Dry Dirt	4	1	24	4	0	75	69.444
Average Accuracy							89.04321

Figure 26 shows how different the red variance of dirt is with respect to the other terrains. Figure 27 shows the difference of all the samples for Point100. Figure 28 shows the difference of Point100 between dry sand and wet dirt.




Figure 26. Red Variance (outdoors). The circled set is dirt, the first half is wet (left) and the second dry. The sets on the left are sand, gravel and pavement. All the thresholds used by the Threshold_Outdoors_D procedure are highlighted.




Figure 27. The Point100 attribute is used to distinguish dry sand from wet dirt. Here are shown all the samples from all the terrains. The set on the square is sand, the one on the right is dirt, and the one on the left is gravel. The circled set is pavement.

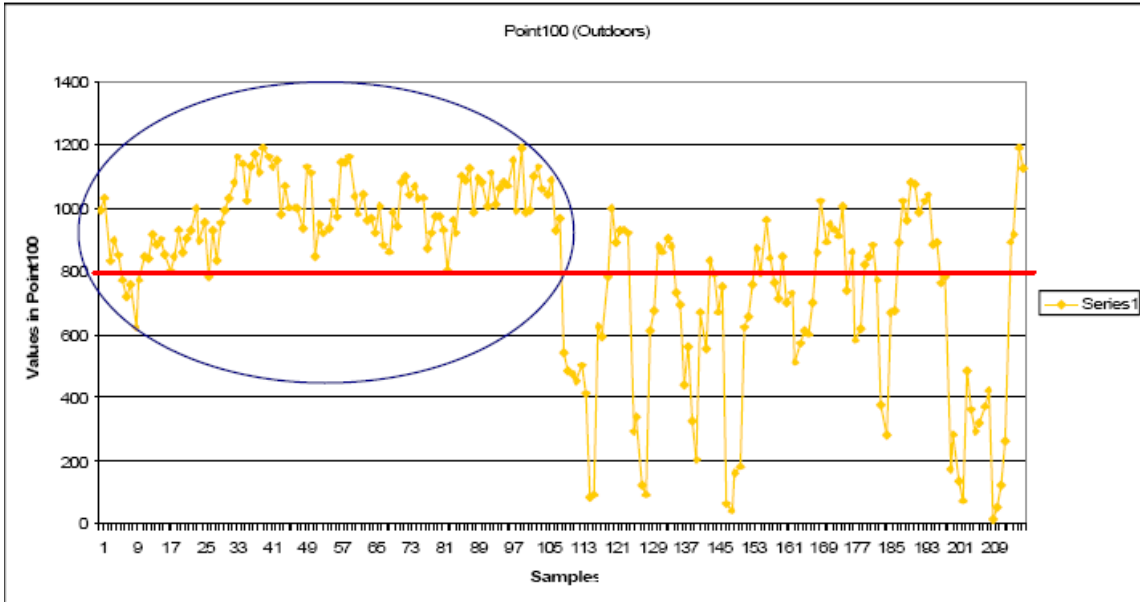


Figure 28. The Point100 attribute is used to distinguish dry sand from wet dirt. The circled set is dry sand and the other is wet dirt. The threshold used is 800. Although the classification is not perfect, Point100 was the best attribute for this discrimination.

Several other points of the waves from the radar were analyzed, but none was as good as the ones used in this algorithm. Figure 29 shows all the analyzed points.

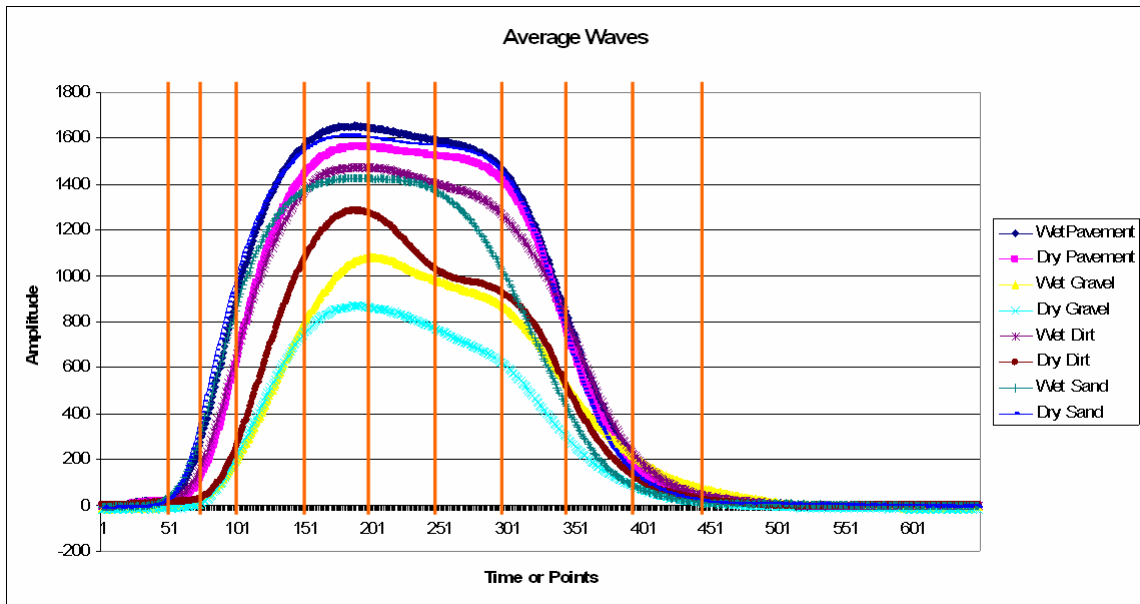


Figure 29. Several other points of the waves from the radar were analyzed, but only Point100 and PointA (160) were used in this last algorithm. The points shown here are 50, 70, 100, 150, 200, 250, 300, 350, 400, and 450.

5.3 Road Surface Detection Using an Accelerometer

In this section, we present the results of using an accelerometer to detect road surfaces. We can distinguish gravel from pavement with 98.94737 percent accuracy.

Table 11. Results obtained using accelerometer system to distinguish gravel from pavement. We classified 380 samples of each terrain at different speeds.

	Gravel	Pavement	Accuracy
Gravel	375	5	98.68421
Pavement	3	377	99.21053
		Accuracy	98.94737

The accelerometer model M9E-USB is used on a 1993 Chevy Astro Van. It is mounted to the lower control arm, on the passenger side, via a specially made bracket so that the only added damping effect between the accelerometer and the road is the tire. Tire pressure will have an effect on the magnitude of the sensed acceleration but only to a relatively small degree if the tire is kept in the proper suggested pressure range. The M9E-USB is an accelerometer and data logger in a single package, which communicates to a host device through a standard USB connection. The unit contains a piezoelectric accelerometer, constant current power source, amplifier, filter, digital signal processor, data memory, and USB transceiver. The unit collects and transmits vibration data via a standard USB link. The data collection parameters are configured from a laptop.

5.3.1 Sampling

The data sent by M9E-USB is available in g forces; it can be in the form of time trace or enveloped time trace. The sampling rate we use is 10000 Hz. So we receive 10000 g's in a second. We have tried our system with the car traveling to different speeds (10, 15, 20, 25 and 30 mph). Table 12. Wave form of examples of gravel and pavement. Each wave form represents a period of time of 7.6 sec. shows the waveforms on time domain of samples of pavement and gravel for a period of 7.6 seconds (76000 g's) at the speeds mentioned above. It is clear that for each speed the amplitude of the waveform of both road surfaces are different. However pavement at 30 mph is very similar to gravel at 10 mph in amplitude. Table 13 shows spectrum of the same samples. The concentration of the lower frequencies on the spectrum of the gravel samples is bigger for all the speed cases. Looking at that feature of the spectrum we could distinguish pavement from gravel since even pavement 30 mph is different than gravel 10 mph (something that in the time domain was not quite possible to distinguish).

Table 12. Wave form of examples of gravel and pavement. Each wave form represents a period of time of 7.6 sec.

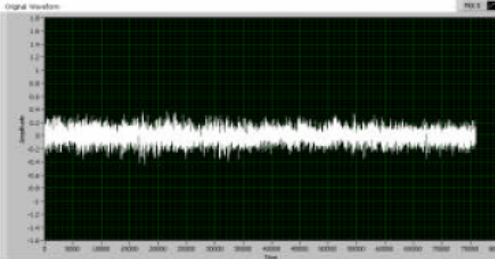
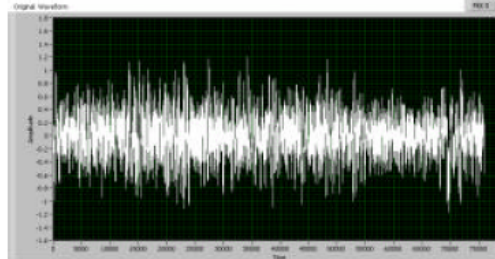
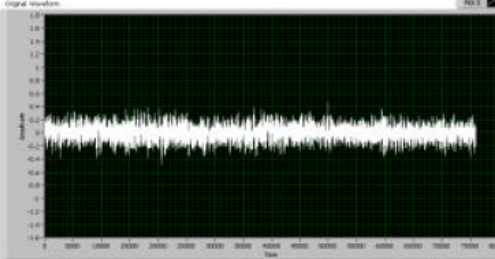
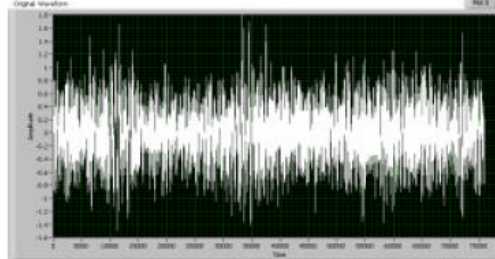
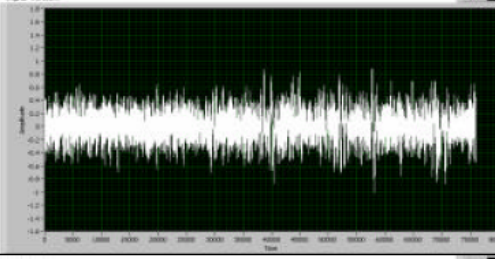
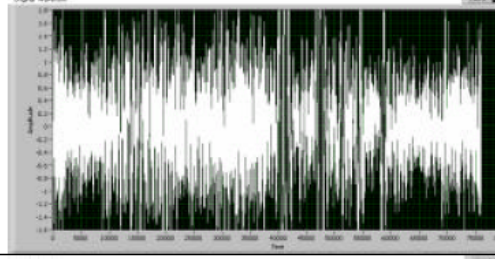
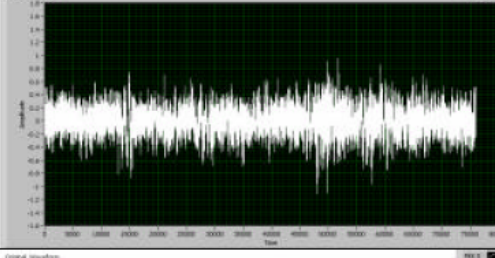
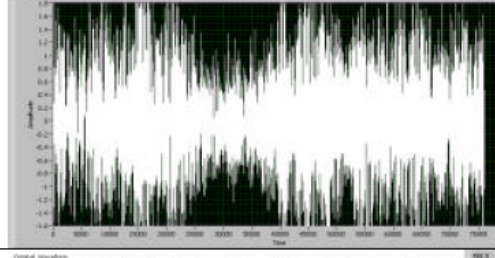
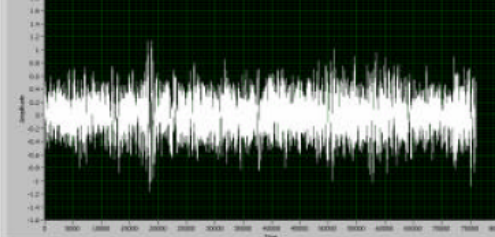
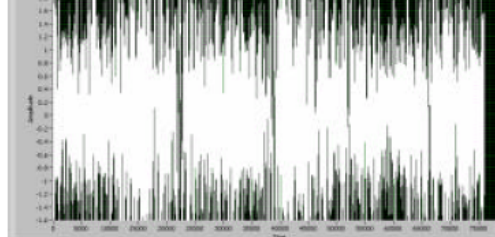
Speed (mph)	Pavement	Gravel
10	 <p>Original vibration plot for Pavement at 10 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>	 <p>Original vibration plot for Gravel at 10 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>
15	 <p>Original vibration plot for Pavement at 15 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>	 <p>Original vibration plot for Gravel at 15 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>
20	 <p>Original vibration plot for Pavement at 20 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>	 <p>Original vibration plot for Gravel at 20 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>
25	 <p>Original vibration plot for Pavement at 25 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>	 <p>Original vibration plot for Gravel at 25 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>
30	 <p>Original vibration plot for Pavement at 30 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>	 <p>Original vibration plot for Gravel at 30 mph. The y-axis is labeled 'Amplitude' and ranges from -1.6 to 1.6. The x-axis is labeled 'Time' and ranges from 0 to 8000. The plot shows a dense, noisy signal centered around zero, with a peak amplitude of approximately 0.4.</p>

Table 13. Spectrum of the gravel and pavement examples on Table 2.

Speed (mph)	Pavement	Gravel
10		
15		
20		
25		
30		

Table 14. Wave form of the g's of gravel and pavement. Each wave form represents a period of time of 0.01 sec.


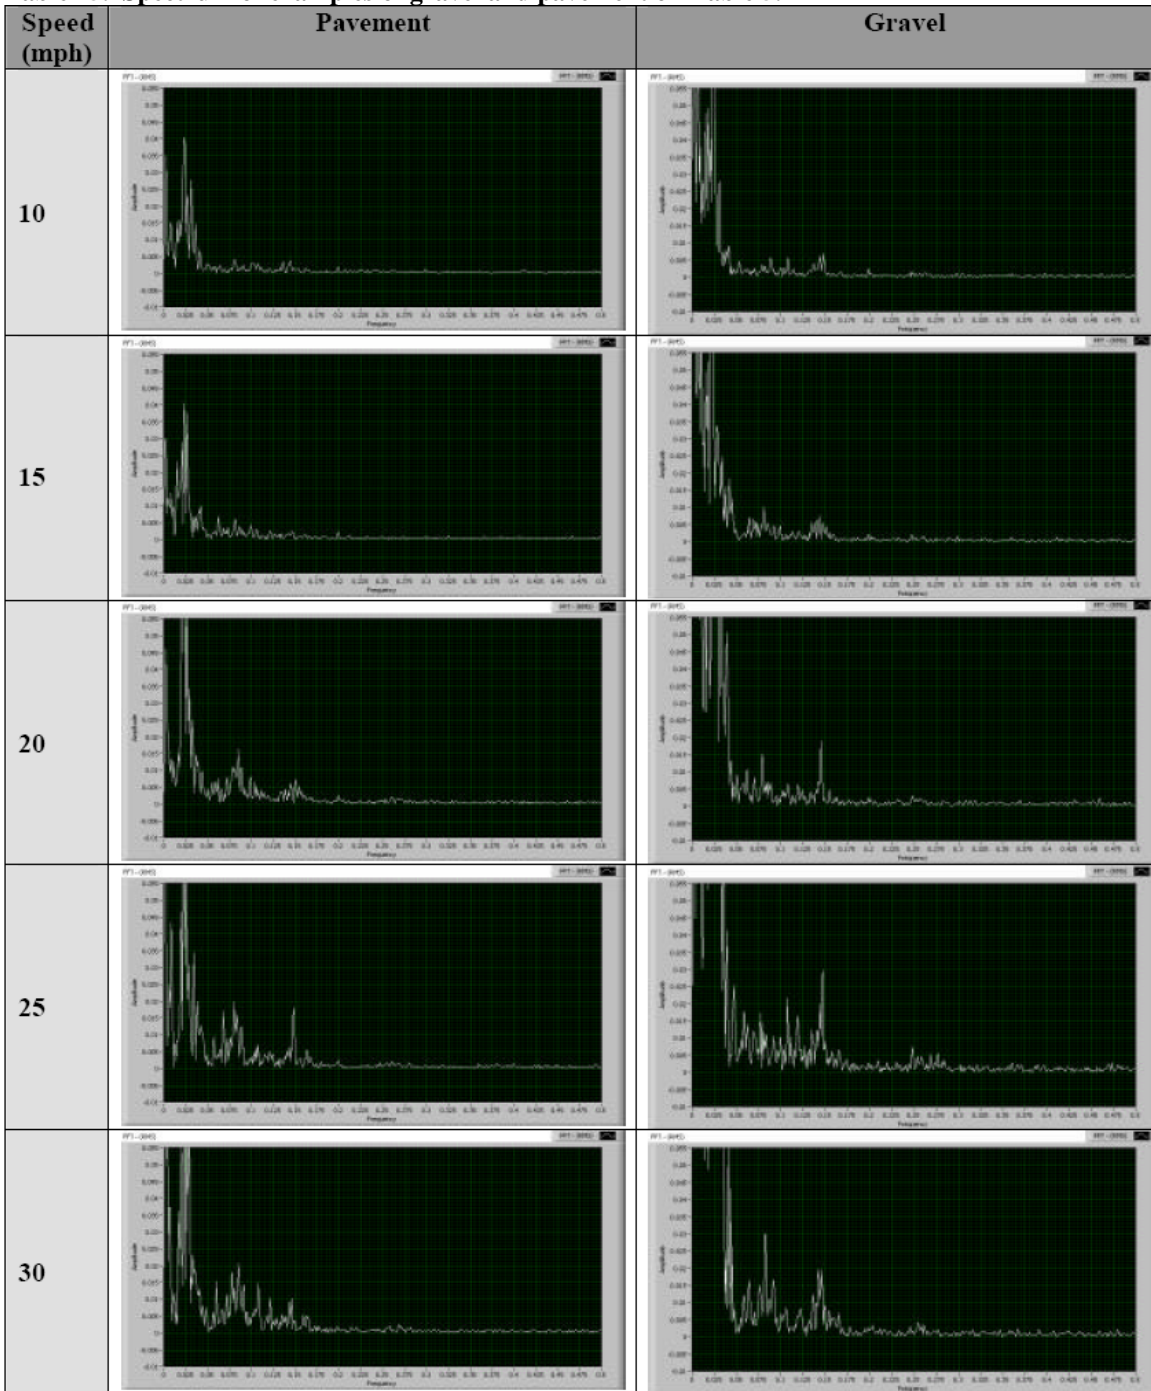


Table 15. Spectrum of examples of gravel and pavement on Table 5.



5.3.2 The 10th of a second analysis

Since we are required to identify the road surface every 10th of a second, we made an analysis of segments with that period of time for all the speeds previously mentioned. We analyzed 76 segments for each combination of terrain and speed. Table 16 summarizes the distribution of the samples used in the analysis.

Table 16. Distribution of the samples or segments used on the analysis.

Terrain\Speed	10 mph	15 mph	20 mph	25 mph	30 mph	Total
Pavement	76	76	76	76	76	380
Gravel	76	76	76	76	76	380
Total Samples						730

Each segment has only 1000 elements (g's). For each speed, the amplitudes of the waveforms of both terrains are different. Although, as in the larger segments on Table 12, the amplitude of pavement at 30 mph is very similar to the amplitude of gravel at 10 mph. However, their spectrums are different (as shown in Table 15).

5.3.3 Time domain

Since we want to know how different the amplitudes of the samples are, we use the attribute **Sum_Abs_Amplitude** (the sum of the absolute value of all the g's on the sample or segment) to compare those segments. Figure 30 shows the **Sum_Abs_Amplitude** of all the samples grouped by speed. We can see, in Figure 30, that for each speed tested, basically the samples of pavement have a **Sum_Abs_Amplitude** lower than the samples of gravel. We can consider that all the samples having a **Sum_Abs_Amplitude** bigger than 1174 are gravel samples (gravel safe zone), and that all the samples having a **Sum_Abs_Amplitude** smaller or equal than 497 are pavement samples (pavement safe zone). However, several samples of both gravel and pavement have a **Sum_Abs_Amplitude** bigger than 497 and smaller than 1174 (confusion zone). Basically the samples on the confusion zone are gravel samples at speeds 15 mph and less and pavement samples at speeds equal or higher than 20 mph. In total 316 samples out of 760 are on the confusion zone.

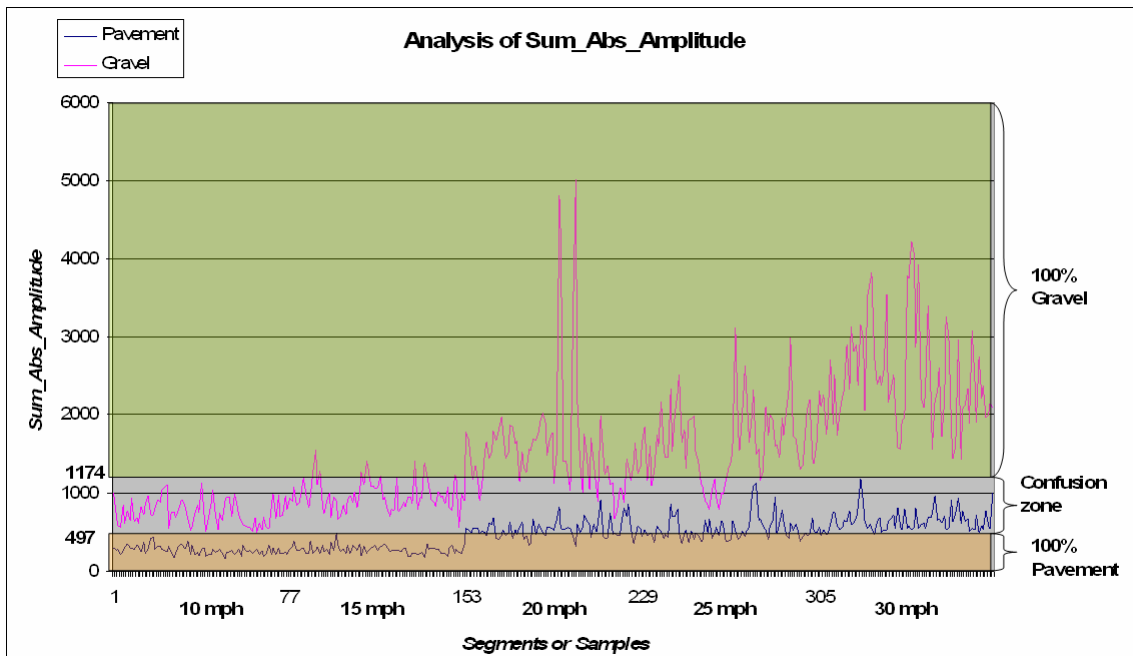


Figure 30. Analysis of Sum_Abs_Amplitude, it contains 76 samples for each speed and 316 samples out of 760 are on the confusion zone.

5.3.4 Spectrum analysis

The spectrum of each sample contains 500 points. We can observe, on Table 13 and Table 15, two things; first the amplitude of lower frequencies is bigger than the amplitude of the higher frequencies. Second, the ratio of low to high frequencies is higher for gravel samples than for pavement samples. We defined the attribute **FFTRatio** as the ratio of low and high frequencies. $FFTRatio = ValueA / ValueB$, where **ValueA** is the sum of the amplitude of the first 55 points and **ValueB** is the sum of the point 56 to 170 of the spectrum for a given sample. Figure 31 shows the ranges of **ValueA** and **ValueB**. Figure 32, Figure 33 and Figure 34 show those 3 attributes for each sample. In general, for pavement the range of **FFTRatio** decreases as the speed increases. On the other hand, for gravel the **FFTRatio** is on the same range for all the speeds. The shape of Figure 33 for **ValueA** is very similar to the shape of Figure 32. To obtain the result on Table 11. Results obtained using accelerometer system to distinguish gravel from pavement. We classified 380 samples of each terrain at different speeds., we use a procedure described in the next section that makes use of both attributes, **FFTRatio** and **ValueA**.

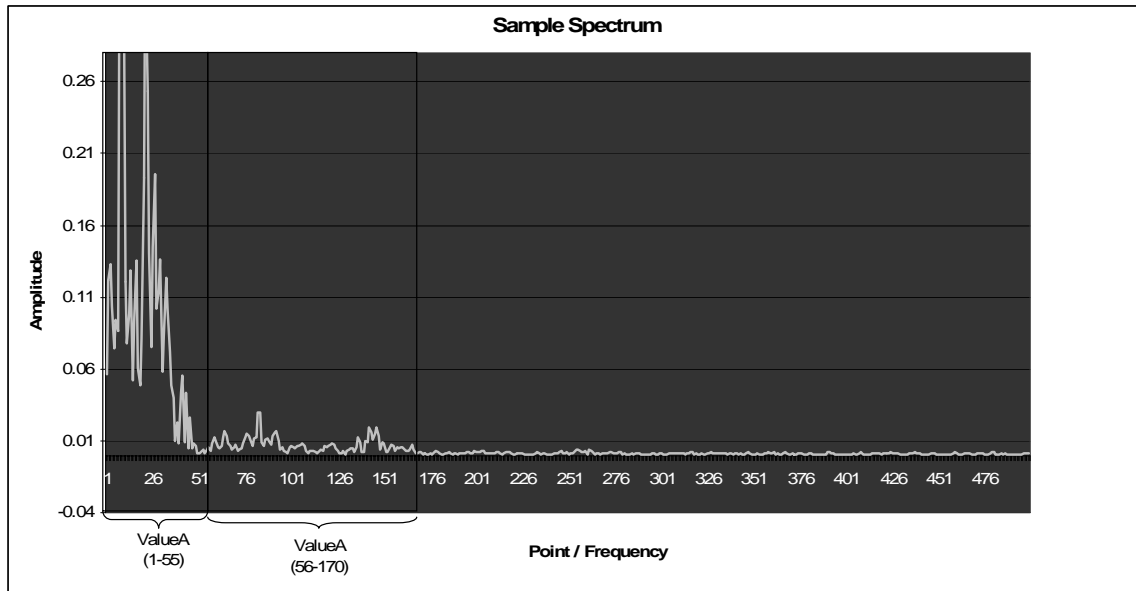


Figure 31. Ranges for ValueA and ValueB, both used to calculate FFTRatio.




Figure 32. FFTRatio of all the samples. FFTRatio decreases as the speed increases.




Figure 33. ValueA of all the samples.




Figure 34. ValueB of all the samples.

5.3.5 The classifier procedure

We created a classifier based on **FFTRatio** and **ValueA** that can generate the results in Table 11. This procedure classifies as gravel everything that has either **ValueA** > 2 or **ValueA** > 1.049 and **FFTRatio** > 4 and anything else is classified as pavement. Figure 35 shows this procedure.

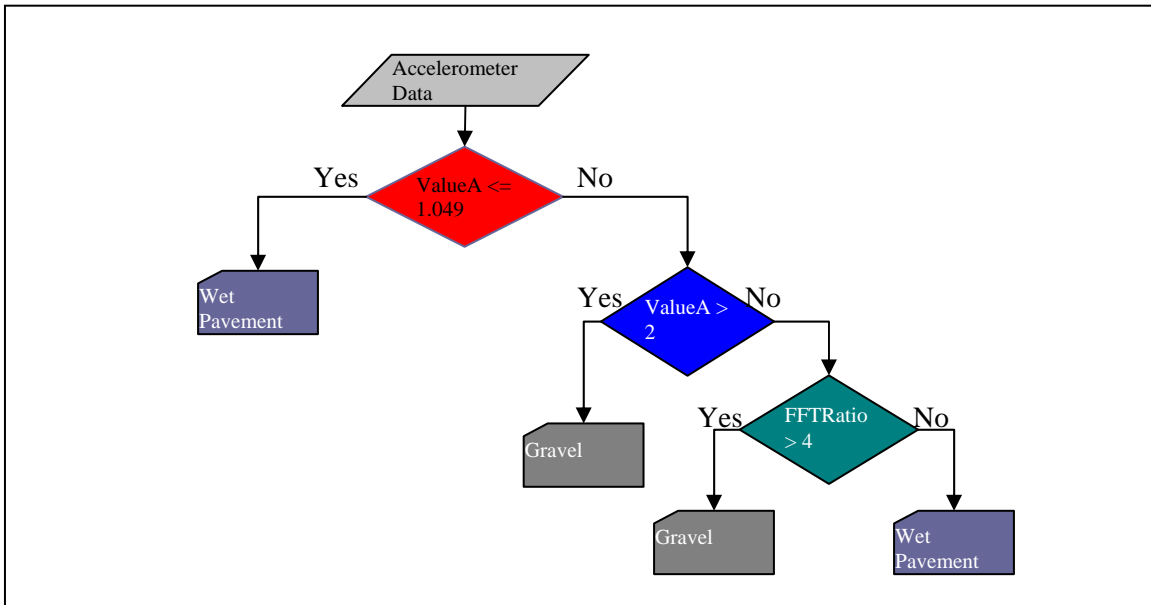


Figure 35. Procedure to distinguish gravel from pavement using the accelerometer.

ValueA is very similar to **Sum_Abs_Amplitude**. It has both, a gravel safe zone and a pavement safe zone. The gravel safe zone has 2 as the lower boundary and the pavement safe zone of **ValueA** is on the range 0 to 1.049 (Figure 36 shows both zones). The confusion zone is on the range $1.049 < \text{ValueA} < 2$ and contains gravel samples at speeds 10 and 15, and pavement samples at speeds 20 and above. The pavement samples on the confusion zone have an **FFTRatio** < 4 because of the speed (the higher the speed the lower the **FFTRatio**), and the gravel samples must have **FFTRatio** > 4 (Figure 37 shows the thresholds on **FFTRatio**). Notice that most of the pavement samples with **FFTRatio** > 4 are on the lower speeds and they have a **ValueA** ≤ 1.049 , so they are classified as pavement; and most of the gravel samples with **FFTRatio** < 4 have **ValueA** > 2 , so they are classified as gravel.




Figure 36. ValueA zones.




Figure 37. Threshold of FFTRatio used on the classifier procedure.

6 Classification

The classification process is decomposed into a preprocessing phase, a learning phase, and an execution phase. The preprocessing phase occurs in order to tame the vast amounts of data coming into the system by preserving only the values found useful in discriminating between terrain types. The learning phase then uses these filter values to generate rules that will effectively classify the terrain type and the precipitation on the terrain. The execution phase or the real-time terrain classification is the phase in which effective and real-time classification of terrain may ensue.

6.1 Preprocessing: Extracting data of interest from the sensors

There is a plethora of data being supplied by the three sensing components of the JIDOKA architecture approximately every 10 milliseconds. Of this vast amount of perpetual information, only a subset is necessary for distinguishing among the various types of terrains and the different precipitation possibilities that may exist. The preprocessing phase extracts this subset of data from the sea of information being supplied by the sensors and forwards this filtered data set to the classification phase.

The pulse radar supplies an analog signal that is then transformed into a digital waveform. Of this digitized data, the actual information of interest is the echoed transmission, so the original transmission waveform is removed and discarded. The echoed transmission then spans the next 640 points. The filtered pulse radar data set is comprised of the first of these 640 points, and then every twentieth point until the 640th point is reached. These 33 points are then passed to the classifier for terrain identification.

The original camera image has dimensions 240 by 360. The camera will be mounted such that the center of the image should only contain the approaching terrain. Thus the central 80 by 110 pixels will be extracted from the original image. These pixels will then be used to calculate the mean and variance of the red green blue color scale contained in the terrain. These six data points will then be passed to the classifier in order to aid in terrain distinction.

The accelerometer returns a sample set consisting of 250 samples per 10 milliseconds. For each sample set a variety of expressive values are calculated. These values are the accumulated sum, the accumulated high frequency FFT values, and the accumulated low frequency FFT value.

For the accumulated sum, the absolute values of all 250 amplitudes are averaged together to generate an average vertical displacement. This value is useful in distinguishing terrains if the speed of the vehicle is equal for each terrain type being discriminated. If the speeds change, then the FFT must be executed in order to observe the various frequencies represented in the data.

Once the FFT has been calculated, the waveform is divided into two sets, the first 15 points of the FFT waveform and the final 39 points of the waveform. The value called AccValueA is calculated by summing the amplitudes of the first 15 points of the FFT waveform, the higher frequencies represented in the data set. Similarly, the value

AccValueB is calculated by summing the amplitudes of the last 39 points of the FFT waveform, or the lower frequencies represented in the data set. Once the two values are calculated, they are compared via the AccFFTRatio value which is defined as the ratio of AccValueA / AccValueB.

Once the data is filtered, the resulting values are more expressive of the discriminating aspects of terrain and precipitation necessary for classification. Yet in order to generate effective classification, these values must be used for the training of the classification rules. This training must persist until testing reveals satisfactory results to within the required accuracy. This process of learning classification rules is detailed in the following section.

6.2 Learning: Classification Rules

The learning phase of the JIDOKA classification process relies on the C4.5 application (the C4.5 application is an actual implementation of the C4.5 algorithm by Ross Quinlan and builds Classification/Decision trees) for generating rules for classifying terrains from a sample data set. The rule induction is based on a decision tree analysis of quantified characterization of the terrain provided by the three sensors, camera, radar and accelerometer. The C4.5 application execution may be decomposed into two sequential processes, those of training and testing. The training process actually generates the rules and the testing process implements those rules on an independent testing data set in order to quantify each rule's discriminating capabilities so that appropriate rules may be elected.

6.2.1 Training

Once the data has been filtered, it is then used to train the classification system. This training is defined by the generation of rules designed to exploit the distinctions between terrain types exposed in the data sets used for training. These rules are created using a tree-based rule generation algorithm (C4.5) in which classes are characterizing by tree branches. For example, if a distinguishing quality is that snow has very not mean values for the amount of red in the image, then a tree node will compare the mean red value as less than a certain low threshold. The left hand branch of this node represents true outcomes to this comparison and the right hand branch represents false. Thus, the left hand branch characterizes the class of snowy terrains. The training process is described in more detail in Chapter 7.

6.2.2 Testing

Once a set of rules is generated, these rules must be tested against independent data sets in order to assess whether they ensure adequate class discrimination. This assessment is made via confusion matrices designed to describe the discriminatory potential of the rules generated from the training set.

Each rule is evaluated in terms of confidence, support and capture. The confidence of a rule represents the accuracy with which the rule will correctly identify a particular terrain. For example, if a rule correctly identifies 8 terrain tests as dry dirt and incorrectly identifies 2 terrain tests as dry dirt, then the rule has a confidence of 80%.

Support measures the percentage of training data for which the rule applies. This metric is used to express how applicable the rule is to the problem. For example, if a rule

is only applicable in a single, rare situation, then its support will be a very low percentage, and indicates that the rule may not provide enough discrimination.

The capture metric conveys how good the rule is at identifying a specific class. This measurement will be highly dependent on the distribution of the values within the problem space. For example, if all wet pavement data are close together in the domain space, then a rule may have a high capture rate.

In order to provide a concise understanding of the caliber of discrimination that a rule generates, all three of these metrics must be considered. As an example, let us consider a rule that was created and tested with the following results.

	Confidence	Support	Capture
snow	100%	25%	100%

The rule corresponding to the percentages is as follows:

```
If CamRedVar < 43 then Species_name = snow.
```

This rule says that if the red variance data provided by the camera (CamRedVar) is less than 43, then the terrain (Species_name) is snow. The 100% capture measurement means that this rule captures all of the snow terrain in the test data. The support value conveys the fact that this rule applies to 25% of the test data. The confidence value states that of all the test data for which the red variance is less than 43, all of them are snow. Together these results show that $\frac{1}{4}$ of the test data is of snow, and all data comprising this $\frac{1}{4}$ partition is identified by the above rule. Also, no test data that is not in this $\frac{1}{4}$ partition is classified as snow. In the JIDOKA system, the confidence measurement is prioritized over support and capture values. The process of testing is described in greater detail in Chapter 7.

6.3 Real-Time Terrain Classification

Also referred to as the execution phase, the real-time terrain classification phase begins once testing has revealed satisfactory classification capabilities of the generated rule set. The active JIDOKA terrain detection system may then be invoked with the specified degree of accuracy.

PART TWO: System Learning

7 Training and Testing

The process of training consists on two steps. First, collect samples of the desired terrains to identify, we recommend at least two hundred (as showed in the section “Saving Data”). And second, create the set of rules to discriminate one terrain from another, using the C4.5 application. Following those two steps, the system is able to learn how to classify a new type of terrain.

7.1 Saving Data

Once the program is running and all the sensors are working, you can start to save samples of the current terrain by pressing the “Save Data” button located at the top left corner on the user interface (see next image). When the “Save Data” button is on, the program starts saving preprocessed data of the terrain from all the sensors.

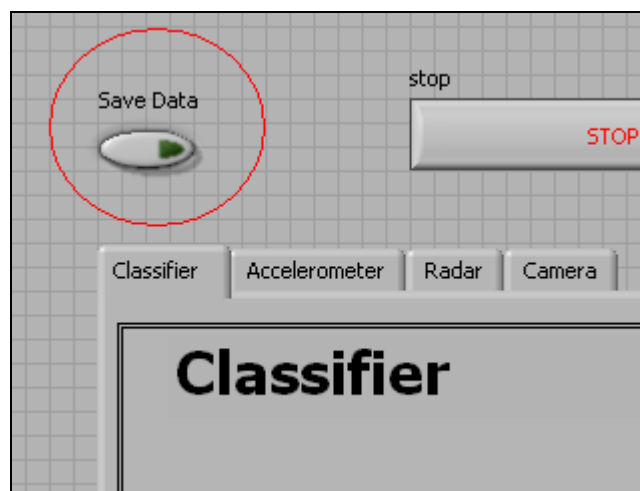


Figure 11: The “Save Data” button is located at the upper-left corner of the tab control.

The collected data is saved in a directory called “DATA” inside the JIDOKA directory (usually “C:\Jidoka\DATA”). On the DATA folder there are three subfolders, the “all” folder (where the preprocessed data from all the sensors is saved), “Camera” folder (where pictures of the terrain are saved), and the “Classifier” folder (where the classification results are saved). Inside the “all” folder you will find files with the pattern Samples-*-*-*.txt as name, where the first * will replace the number of month, the second * the day and the third * is the minute where the data has been taken. The “Camera” folder you will find files with the pattern camera-*.bmp (where * is replaced by a number) as name, one file or picture for each sample of the terrain. Meanwhile on the “Classifier” subfolder you will find files with the pattern class-*_*_.csv as name (the * are replaced as in the “all” subfolder), containing the preprocessed data of each sample of the terrain (attributes from each sensor).

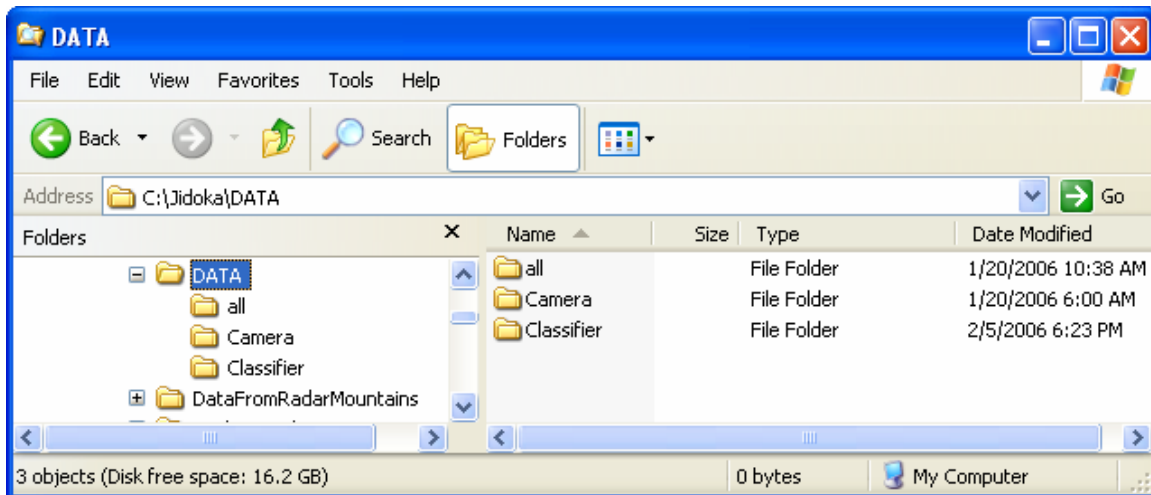


Figure 12: The DATA directory structure.

When you have the desired number of samples of the terrain (we recommend at least 250: 200 for training and 50 for testing), you can compile the files on the “all” directory using the program and add the data (in order) and pasting the data on the C4.5 application.

Both folders, “all” and “Camera”, must be empty before starting to save data from another terrain. Once you have compiled the files for a specific terrain, you should move the files on the subfolders of the DATA directory.



Figure 13: The “all” directory content.

Once you have samples of all the terrains you want to be able to distinguish, and after the compilation of the files on the “all” directory, you have to use the *compiled* files to create a set of rules using the C4.5 application.

7.2 Invoking the C4.5 Application

The C4.5 is an MS Excel file that contains an implementation of the C4.5 algorithm. This implementation was done by Angshuman Saha, and it processes the attributes from the set of sensors to create a set of rules that will be used by the classifier program. The Excel file (the one containing the application) has seven spread sheets or tabs. The first tab titled “ReadMe” contains a tutorial of the application and explains some concepts and the functionality. The tab titled “UserInput” is where you should set the parameters and options explained on “ReadMe” and then after load your training samples (explained later) you should click on the button “Build Tree” to generate the set of rules. The samples, previously connected (section 6.1), have to be loaded on the tab “Data”. The best way to do this is put first a set of data for each terrain, as training data and then put another set for each one, as testing data; in that way you can validate the rules. The tabs “Tree” and “Node” show the resultant tree and the selected node respectively, after the process of building the tree is done. The “Result” tab present a summary of the data used and the most important a confusion matrix showing the accuracy of the rules in both the training and the testing data. Finally in the tab “Rules”, a rule summary table is showed including the support, confidence, capture (concepts explained on the “ReadMe” tab) and the rules themselves.

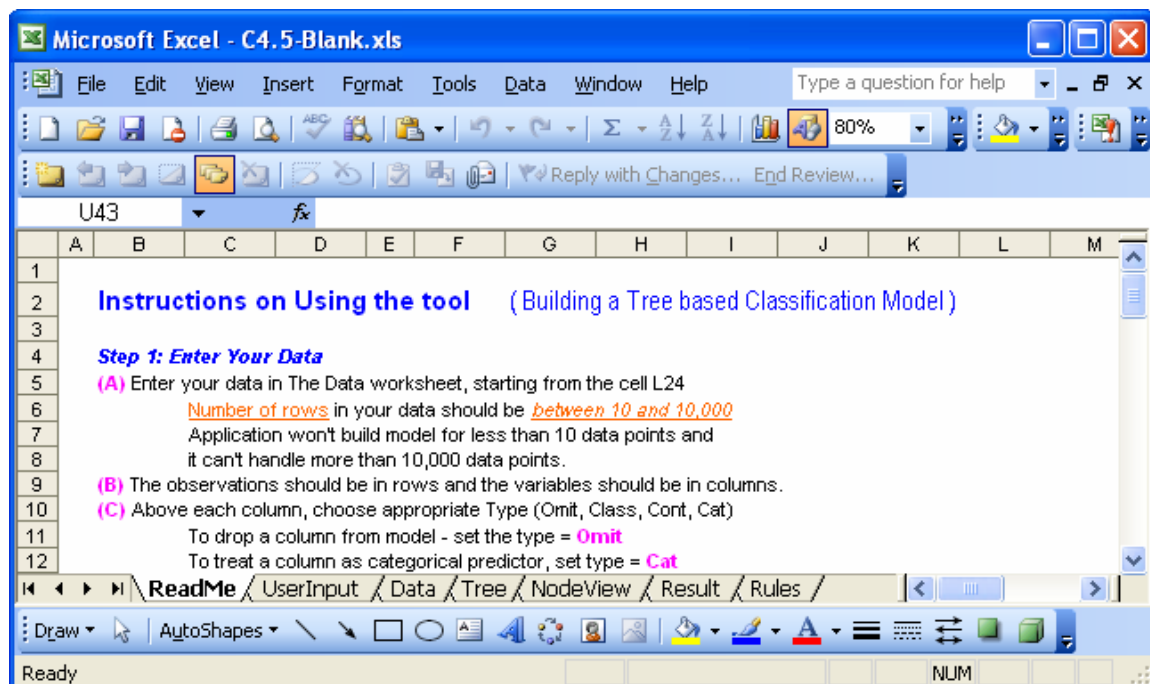


Figure14: The tabs included on the C4.5 application.

The set of rules will enable the classifier to distinguish among the terrains that have been used to create the rules, meaning that you have to put samples for all the terrains that you want to identify.

7.3 Creating Rules

On the “Data” tab, you have to load the content of the compiled files for each terrain. That means that for each terrain you should have a file containing the compiled attributes of each sample. For example, if you want to be able to distinguish among gravel, pavement, dirt, and snow; you have to have a compiled file for each terrain. Continuing with the example, if you took 250 samples for each terrain, each compiled file should have 250 rows, with each row representing a sample. We recommend using 200 samples (rows) of each terrain and 50 for training. It could be that the compiled files are tab separated text files, in which case you can open them using Excel and then select the first 200 rows and paste them in the “Data” tab of the C4.5 application. The first column of the compiled files corresponds to the “Point1” attribute, so you should paste the data on the “N” column on the “Data” tab. After the data is pasted you should add the name of the class (gravel, pavement, etc.) on the column labeled “Species_name” (column “M”) and a corresponding number of class on the column labeled “Species_No” (column L). After copying the 200 samples of the compiled files, you can add the remaining 50 samples of each terrain following the training data they are going to work as testing data.

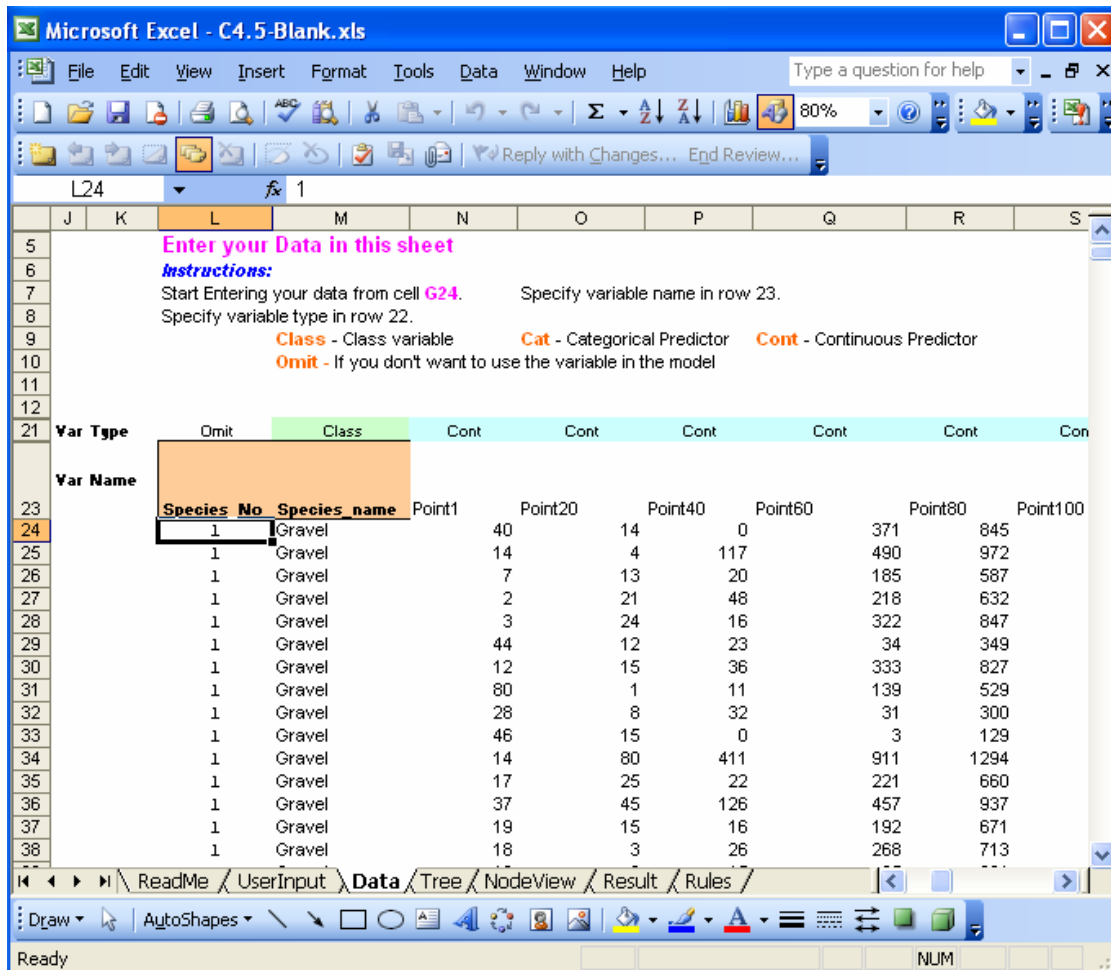


Figure 15: The “Data” spreadsheet.

When the “Data” spreadsheet is ready, you should go to the “UserInput” tab and do the following. First modify the parameter on the “Option 2” that says “Use last #### rows of the data as validation set” where #### has to be 200 for our example (50 for each terrain and 4 terrains). Second, click the “Build Tree” button to start the tree generation process. *This process would take approximately 30 minutes for our example in an Intel Pentium M laptop at 2.0 MHz. and 1GB of RAM.*

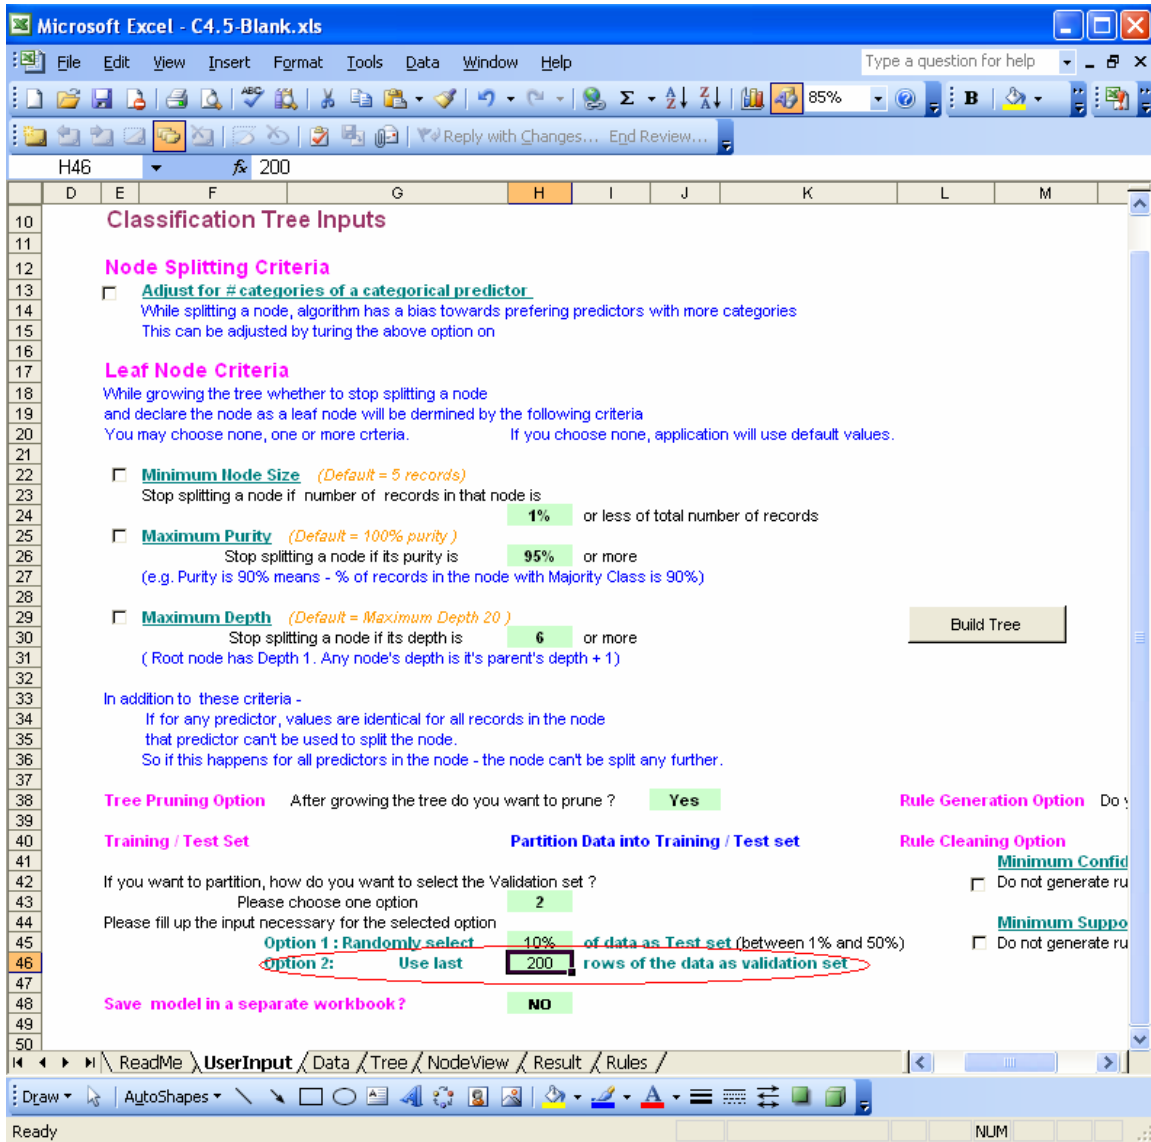


Figure 16: The “UserInput” spreadsheet

The application generates the decision rules. However, the rules are basically for viewing and gaining insight into the classification problem (the application has no capability of using the rules to classify new data points).

7.4 Implementing Dynamically Generated Rules

After the set of rules is created you have to create the file “rulesFile3.txt”, the one that contains the rules used by the classifier. This file has to be placed on the JIDOKA folder (usually “C:\Jidoka\”). Each rule generated by the C4.5 application has three attributes Support, Confidence and Capture (as showed in figure 17). The rules have to be sorted by confidence (this attribute measures the accuracy of the rule.). After the sorting, the rule number zero has to be moved to the bottom.

Rule Summary Table						# Rules
Rule ID	Class	Length	Support	Confidence	Capture	24
0	dirt	0	100.0%	25.0%	100.0%	
1	snow	1	25.0%	100.0%	100.0%	
2	pavement	3	9.9%	95.0%	37.6%	
3	dirt	1	22.8%	85.8%	78.4%	
4	dirt	2	19.4%	98.0%	76.1%	
5	dirt	1	25.0%	82.4%	82.4%	
6	gravel	2	3.6%	73.0%	10.6%	
7	dirt	2	1.5%	80.0%	4.7%	
8	dirt	2	16.9%	94.2%	63.5%	
9	pavement	1	32.2%	67.7%	87.1%	
10	dirt	1	2.0%	65.0%	5.1%	
11	dirt	2	14.6%	71.8%	42.0%	
12	pavement	3	21.9%	96.4%	84.3%	
13	pavement	2	16.3%	75.3%	49.0%	
14	dirt	2	21.5%	79.0%	67.8%	
15	dirt	2	0.7%	85.7%	2.4%	
16	gravel	3	24.2%	83.8%	81.2%	
17	gravel	3	19.0%	93.8%	71.4%	
18	dirt	1	33.5%	56.7%	76.1%	
19	pavement	3	17.4%	96.0%	66.7%	
20	dirt	2	2.2%	95.5%	8.2%	
21	gravel	2	15.9%	93.2%	59.2%	
22	gravel	2	2.5%	84.6%	8.6%	
23	dirt	3	13.6%	77.0%	42.0%	
24	gravel	3	15.1%	68.8%	41.6%	

Figure 17: Rule summary table, generated by the C4.5 application.

```

rulesFile3.txt - Notepad
File Edit Format View Help
Rule1 IF CamRedVar < 43 THEN Species_name = PS
Rule4 IF CamRedAvg < 27 AND CamRedvar >= 110 THEN Species_name = DD
Rule12 IF CamRedAvg < 39 AND CamRedvar < 126 AND Point140 >= 1564 THEN Species_name = PD
Rule19 IF CamRedAvg < 39 AND CamRedAvg >= 27 AND Point160 >= 1602 THEN Species_name = PD
Rule20 IF CamRedVar >= 132 AND Point160 >= 1634 THEN Species_name = DD
Rule2 IF CamRedVar < 110 AND CamRedvar >= 43 AND Point140 >= 1561 THEN Species_name = PD
Rule8 IF CamRedAvg < 27 AND Point160 < 1592 THEN Species_name = DD
Rule17 IF CamRedAvg >= 34 AND CamRedvar >= 94 AND Point140 < 1564 THEN Species_name = GD
Rule21 IF CamBlueAvg >= 41 AND CamRedAvg < 64 THEN Species_name = GD
Rule3 IF CamGreenAvg < 25 THEN Species_name = DD
Rule15 IF CamRedAvg >= 64 AND CamRedvar >= 102 THEN Species_name = DD
Rule22 IF CamRedAvg < 41 AND CamRedAvg >= 39 THEN Species_name = GD
Rule16 IF CamBluevar >= 106 AND CamRedAvg >= 27 AND Point160 < 1589 THEN Species_name = GD
Rule5 IF CamRedAvg < 27 THEN Species_name = DD
Rule7 IF CamRedAvg >= 61 AND CamRedvar >= 43 THEN Species_name = DD
Rule14 IF CamBlueAvg < 36 AND Point140 < 1564 THEN Species_name = DD
Rule23 IF CamBlueAvg < 36 AND Point240 < 1069 AND Point40 < 88 THEN Species_name = DD
Rule13 IF Point140 >= 1564 AND Point20 < 15 THEN Species_name = PD
Rule6 IF CamRedvar >= 43 AND Point280 < 33 THEN Species_name = GD
Rule11 IF CamRedAvg < 39 AND Point160 < 1433 THEN Species_name = DD
Rule24 IF CamRedAvg >= 27 AND Point160 < 1589 AND Point620 < 11 THEN Species_name = GD
Rule9 IF Point160 >= 1592 THEN Species_name = PD
Rule10 IF CamRedvar >= 143 THEN Species_name = DD
Rule18 IF CamRedvar >= 126 THEN Species_name = DD
Rule0 IF CamRedAvg >= 0 THEN Species_name = DD

```

Figure 18: Example of a rules file “rulesFile3.txt”.

The final rule file has to contain the Rule_ID followed by the rule. Each Boolean element of the rule has to be separated by tabs as showed in the figure 18. You have to add “IF CamRedAvg >= 0 THEN Species_name =” to the Rule 0, since basically this rule does not have a head. Finally it is important that the names of the classes are as shown in Table 2.

Table 3: Class names for the rules file “rulesFile3.txt”.

Terrain	Name for rule file
Pavement Dry	PD
Pavement Wet	PW
Pavement Snow	PS
Pavement Ice	PI
Gravel Dry	GD
Gravel Wet	GW
Gravel Snow	GS
Gravel Ice	GI
Dirt Dry	DD
Dirt Wet	DW
Dirt Snow	DS
Dirt Ice	DI

PART THREE: System Assessment

8 Results

The C4.5 application generates a set of rules with corresponding assessments of confidence, support and capture. Within these rules, a single set of data may be classified as more than one type of terrain. This ambiguity must be clarified by relying on the characteristics of the rule set. In this section we discuss strategies for administering the rules to the data being gathered. The results generated by each strategy are explained in the following subsections. Confusion matrices for all results are included.

Given the set of n rules generated by the C4.5 application, we needed to find a mechanism for applying the rules in an efficient manner that would optimize accurate classification. We tried four approaches. These are discussed below.

8.1 *The rule with highest confidence first*

Our initial attempt was to sort the rules based on their confidence only, such that the first rule considered has the highest confidence of all rules in the set. The classifier evaluates the rules in this order, and the first rule whose antecedent evaluates to true has the highest confidence of any rule in the rule set that will evaluate to true for this data. Hence, this rule will determine the terrain type.

In exploring the accuracy of this approach, the experiment consisted of 1029 training data samples. Of each data sample, 39 different predictors or attributes are used for classification. Refer to section 6.1: *Preprocessing: Extracting data of interest from the sensors* for an explanation of the filtering of these predictor attributes. These 39 attributes are extracted as 33 points from the radar data and 6 values from the camera data. This approach of sorting by confidence resulted in the successful classification of 94.22% of the trained data. A confusion matrix of the results is shown in Figure 38.

Classification Tree Model

Number of Training observations	1020
Number of Predictors(Attributes)	39
Number of Classes	4
% Correctly Classified On Training Data	94.22%
% Missclassified On Training Data	5.78%

Confusion Matrix

Training Data

True Class	Predicted Class				
	dirt	gravel	pavement	snow	
dirt	217	30	8		255
gravel	4	244	7		255
pavement	6	4	245		255
snow				255	255
	262	251	252	255	961

Figure 38. Results for rule analysis approach of sorting by rule confidence.

8.2 Voting

In the voting approach, all rules in the rule set are evaluated. If any rule evaluated to true, then that rule would essentially be voting for the terrain type that it classifies. These votes are then tallied, and the terrain type with the most votes wins. So, the terrain type classified by the most rules that evaluate to true wins the election, and is thus the classified terrain for the data sample.

This approach did not perform as well as the previous approach of sorting by confidence. The accuracy of voting was observed at 86.67%. This deterioration was a result of the imbalance of rule distribution per terrain in the rule set. For example, the snow terrain is classified based on only one rule, while there are 13 rules that classify

dirt. Yet, the one rule that classifies snow has a confidence of 100%, while the rules that classify dirt range from 25% to 98%. So, for example, there may be a situation in which the one snow rule evaluates to true, and three low confidence dirt rules evaluate to true. In this case, the voting mechanism will classify the terrain as dirt even though the single snow rule had a much higher confidence simply because there were more dirt rules that evaluated to true than there were snow rules in the rule set. See Figure 39 for a confusion matrix of the results for the voting approach.

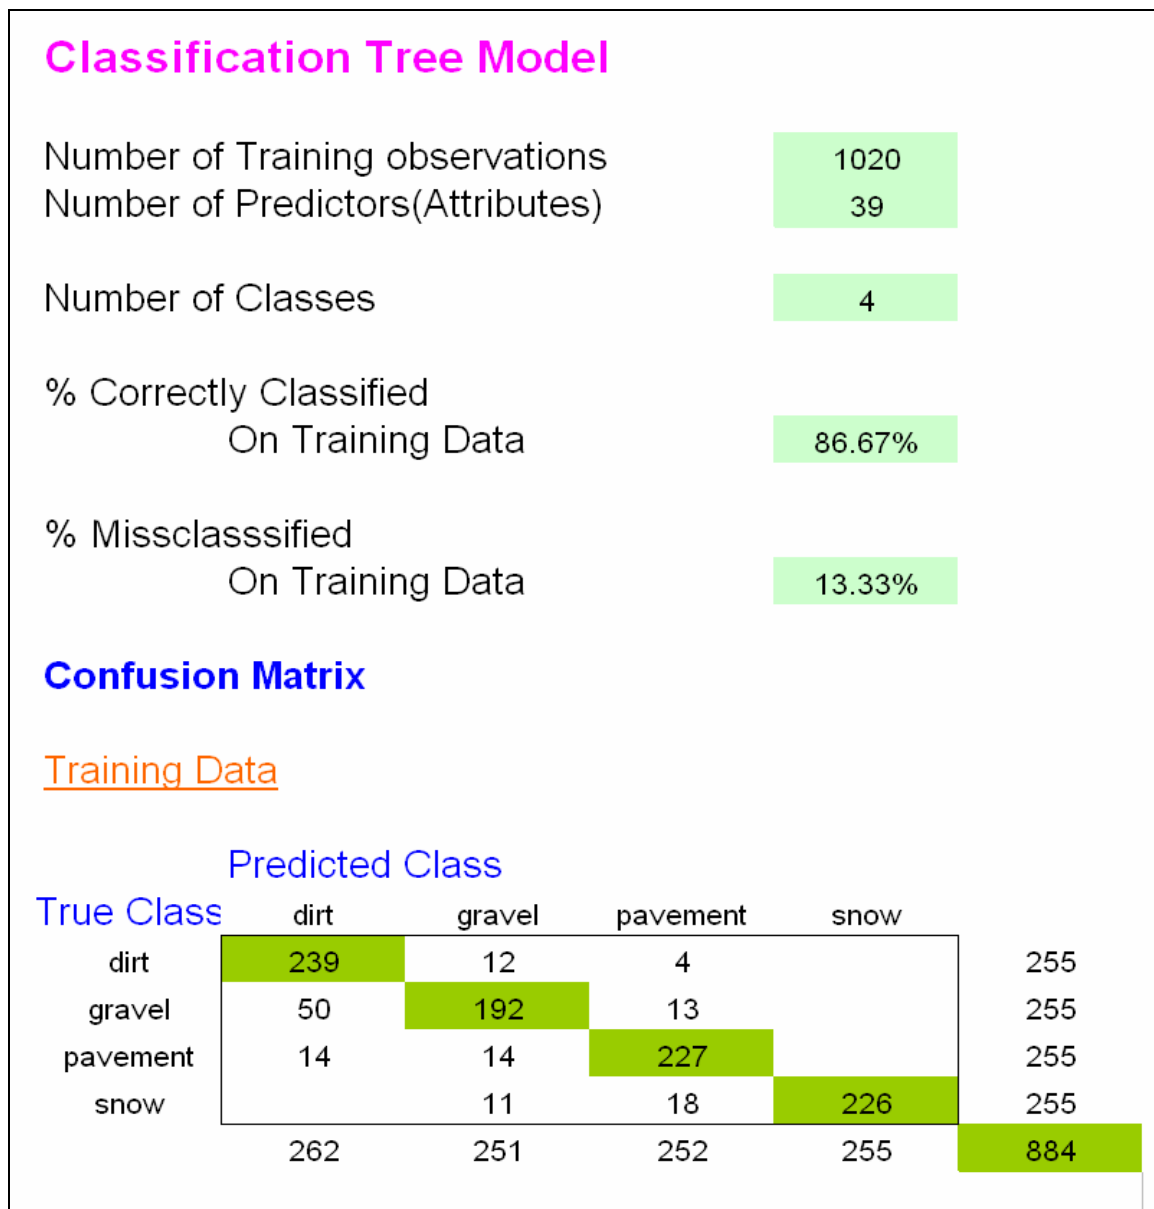


Figure 39. Results for rule analysis approach of voting.

8.3 Weighted Rules

In order to correct for the deficiency of the voting approach as described in the previous section, a weight assessment was incorporated into the voting mechanism. Each rule was attributed a weight that was calculated based on the rule's support, confidence and capture. This weight is calculated by the following equation:

$$Weight = Support \times Confidence \times Capture$$

As in the voting approach, we evaluate all the rules such that each rule that evaluates to true is able to vote for the terrain type that it classifies. In the weighted rule approach, though, each vote has a value equivalent to the weight as calculated by the above equation rather than having a value of one as the vote had in the voting approach. The experiments revealed an accuracy of 90.49% for the weighted rules approach.

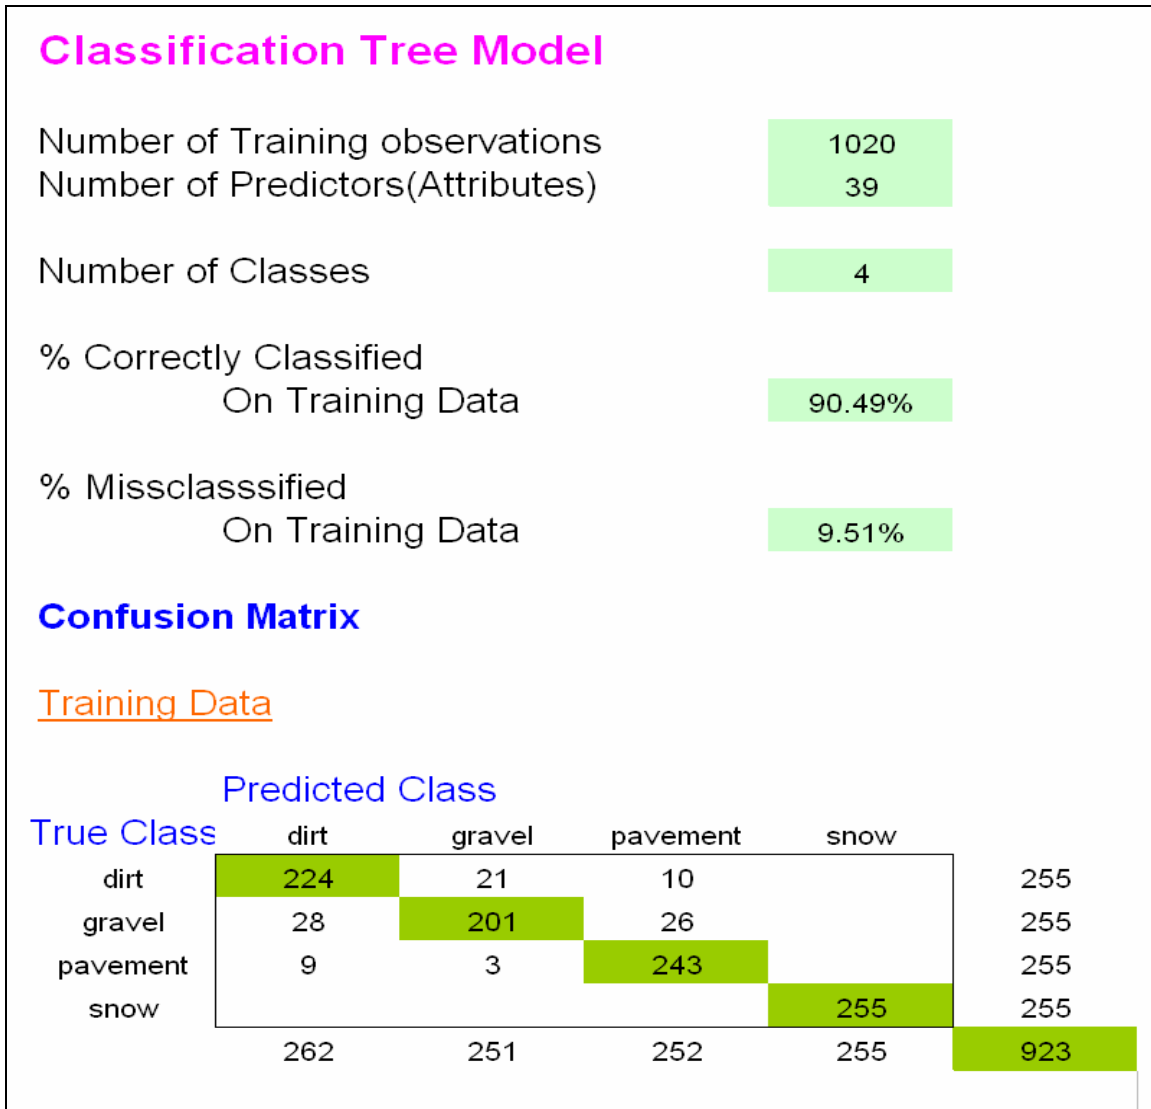


Figure 40. Results for rule analysis of weighted rules approach.

8.4 Normalized Weight

The weighted rules approach corrects for the voting approach's deficiency of not incorporating support, confidence and capture into its assessment of the rules, but it still has the problem that each terrain is not equally represented in the number of rules that classify it in the rule set. In order to correct for this, we propose normalizing the weights of each rule such that the sum of the weights of all rules that classify a certain terrain add up to 1. Implementing this strategy will result in the equal representation based on rule count of each terrain in the rule set. For example, since snow only has one rule in the rule set, this rule's weight will be 1. Similarly, since dirt has 13 rules that classify it, the sum of the weight for each of these rules will be 1.

More specifically, we first assigned a weight to each rule, using the product of the support, confidence and capture percentages as was used in the weighted rules approach. This weight is called *SCC_WEIGHT*. Then for each class of terrain we summed the weights of all the rules that classify that terrain. This sum is referred to as the *SUM_CLASS_WEIGHT*. Now that we have a *SUM_CLASS_WEIGHT* for each class of terrain, we assign a new normalized weight to each rule. This new rule weight is called the *FINAL_WEIGHT*. The *FINAL_WEIGHT* is calculated as follow.

$$Final_Weight = \frac{SCC_Weight \times 100}{SUM_CLASS_WEIGHT}$$

This approach produced a classification accuracy of 90.78%. See Figure 41 for complete details.

Classification Tree Model

Number of Training observations	1020
Number of Predictors(Attributes)	39
Number of Classes	4
% Correctly Classified On Training Data	90.78%
% Missclassified On Training Data	9.22%

Confusion Matrix

Training Data

True Class	Predicted Class				
	dirt	gravel	pavement	snow	
dirt	210	31	14		255
gravel	4	225	26		255
pavement	5	14	236		255
snow				255	255
	262	251	252	255	926

Figure 41. Results for rule analysis of normalized weighted rules approach.

9 Conclusions

In this research, we designed a road surface detection system that utilizes a camera, a pulse radar, and an accelerometer as sensor units for detecting the characteristics of the terrain on which a vehicle is driving. This terrain detection capability is essential to attaining the greater JIDOKA vision of providing vehicles with computationally enhanced structural awareness. The need for accurate terrain detection is due to the significant effects of terrain on the durability and reliability of a vehicle's overall functional integrity.

Enabling satisfactory terrain detection presented quite a challenge. Our research culminated in the application of the C4.5 algorithm to generate a rule set based on a training set of data samples. These rules are then strategically evaluated in order to provide adequate classification.

The best results were generated by the rule analysis approach of ordering the rules by confidence, and then classifying based on the first rule that evaluates to true. In other words, the rule with the highest confidence that evaluates to true is used to classify the terrain. This approach produced an accuracy of 94.22%, which exceeded the next highest accuracy of an approach by 3.44%. These results based purely on sensor data processing are exceptional. The development of the JIDOKA vision is well on its way.

9.1 Future Improvements to Terrain Detection

The results produced in the work are based entirely on the processing of sensor data. It is easily conceivable that incorporating even basic contextual information into the classification algorithm will improve the accuracy of the system. For example, adding an awareness of recent terrain classifications may assist in eliminating disparate classifications.

Another improvement may be the evolution of the rule set. This may be implemented by reserving actual data and then periodically using this data as a training set in order to reorient the thresholds that delineate terrain types. These renovations of the rules will never be drastic, because the accuracy of the new rule set is based on the classifications of the previous rule set, but periodically making small adjustments will maintain an accurate representation of the terrain by the rules.