

i. Summary of the original classifiers and sensor system

With all classifiers, leaf occlusion or overlap increases the likely hood of misclassifications. As the H sensor distinguishes between plant type based on leaf shape and size, if the shape of individual leaves or plants cannot be determined then the accuracy will decrease.

There are classifiers designed for use in different light conditions, the x.0 classifiers are suited to high light conditions and the x.1 classifiers are more suited to cloudy conditions.

WW3.0 and WW3.1 classify into two categories, TRZAW and DICOT (cereal and broadleaf weeds), these are best suited to small cereal plants, growth stage less than 4 leaf (Zadocks GS14) and dicot weeds. They do not classify well when the crop row is thick, as the cereal leaves that overlap create a large “blob” in the image and this is more often than not classified as DICOT. These classifiers can be used to identify broadleaf weeds in cereal crops and are also effective at identifying grass weeds in pulse crops. In the pulse crop scenario the thick crop row creating a large blob is less of an issue because it is DICOT in the crop row and therefore will likely be classified correctly. Issues can arise when grass weeds are so thick that they create a complete green mat on the soil surface as then they are classified DICOT or when they only occur in the crop row because of leaf occlusion. This can occur when pre emergent grass herbicides are used as the inter-row in this situation can be weed free.

WW2.0 and WW2.1 classifies into three categories, TRZAW, MOCOT and DICOT (cereals, small grass weeds and broadleaf weeds), as with WW3.0 and WW3.1, it is suited to small cereal plants less than 4 leaf (Zadocks GS14). It uses grass weed leaf size and shape to separate small grass weeds from cereal crop leaves. The classifier frequently misclassifies the end of the crop leaf as grass weeds. This is because in the image, the tip of the crop leaf is separated from the rest of the leaf where the twist in the crop leaf is edge on to the sensor. Classification of DICOT leaves is good. This classifier creates more individual features or segments than the WW3.0 and WW3.1 classifiers and so the image can be broken up into more pieces for more accurate classification. These classifiers can be used to identify broadleaf weeds in cereal crops and is also effective at identifying grass weeds in pulse crops. For the purposes of our early assessments TRZAW and MOCOT have been used as the same class. The assessments have concentrated on the separation of dicot and monocot plant types.

MAIS2.2 is designed for use in maize crops, it classifies into three categories, MOCOT, ZEAMX and DICOT. This classifier can be used in larger cereal crops, it has been used in an oat hay crop to identify wild radish. It was more effective than the WW2.x or WW3.x classifiers because of the broader leaf of the oat crop. For the purposes of our early assessments ZEAMX and MOCOT have been treated as the same class and combined into 1 class. The assessments have concentrated on the separation of dicot and monocot plant types.

RAPS1.0 and RAPS1.1 are designed for use in canola crops, it classifies into three categories, BRSNN, MOCOT and DICOT. It has been effective in our canola crops at identifying grass weeds however leaf occlusion within the crop row makes it difficult to identify grass weeds growing under or immediately adjacent to the canola plants. For the purposes of our early assessments BRSNN and

DICOT have been treated as the same class. Assessments have concentrated on the separation of dicot and monocot plant types.

All classifiers have the ability to classify features or segments into a separate category, UNDEF. This is the category that is assigned if a feature or segment does not fit into any of the normal classes.

The H sensor is capable of taking an image 24cm * 37cm, processing it and making a decision up to 10 times per second depending on the amount of green material in the image. The output from the sensor is then given in the form of percent classification per image. For example an image may contain 2.2% DICOT (broadleaf weed) and 30.5% TRZAW (cereal crop). From these figures a decision can be made to spray or not to spray or the raw data can be used to generate a weed density map.

The per cent area of the field that is captured in the images depends on three factors. The image capture rate, the speed of travel and the swath width of the camera passes. With an average of 8 images per second traveling at 12km/h and a camera swath width of 6m approximately 4% of the field is captured in the images. This infers that the 4% capture represents the remaining 96% of the paddock. To capture 100% of the field in the imagery at an image capture rate of 10 images per second a travel speed of 8.5km/h would be required with a camera mounted every 37cm.

Questions still to be answered;

What proportion of the field needs to be captured to make a reliable representation of the weed distribution in the field?

Does it make a difference what weed species you are targeting?

What level of misclassification is acceptable?

‘Classification errors of weed species or class can fall into two categories; omission and commission. Omission occurs when a plant belonging to a certain class is classified as something else, commission occurs when a plant from another class is classified as the class of interest (Lamb & Brown 2001).’

ii. Method of assessments

Images were collected from a series of crop paddocks at a timing when a significant level of weed infestation could be identified. This corresponded to a range of growth stages in the following crops from 2014 - 2016; wheat, barley, field pea, faba bean, lentil, lupin, canola, chickpea and oat. Several of each crop species were targeted. All images were collected by mounting the H sensor to a vehicle. Some entire paddocks were scanned by following seeder or sprayer tracks in transects across the paddock, in some paddocks specific areas of the paddock were targeted where weed densities were high or varied. All images captured were recorded on portable hard drives for future analysis.

Analysis 1 - The Plant Feature Method

The first analysis for a series of images involves taking a subset of 50 – 100 of the pre-recorded images and running them through the H sensor software with the classifier to be tested. This process generates a file which contains a list of all of the individual features or plant segments for each image. This file is then used in a second software package to manually label all of those individual

features with the correct classification. The output from the sensor can then be compared to the output from the labelling software to give an objective assessment of the accuracy of the classifier. This method is a good way to measure incremental improvements to a classifier. For example the accuracy of classification may be 85% of features correctly classified, an incremental improvement of 5% may be difficult to judge by eye on the sensor but it would be detected in this way.

One flaw in this method occurs when a weed and a crop plant occur in the same feature. In this case the individual plant parts cannot be separated and therefore an accurate assessment cannot be made. In these situations the labelling software has the ability to label these features as UNDEF, or undefined.

Analysis 2 - The Threshold Method

The second analysis involves making assessments based on the entire image rather than the individual features within the image, and what decision would be made from the sensor output. The decision may be to spray an area when the H sensor returns a DICOT value of greater than 0.1%. This value can be manipulated depending on the size of the weed and or the tolerance to the weed in question. This is a crude method of analysis but gives another measure of the accuracy and the usability of the data generated by the sensor. Because this method returns a value based on an entire image the levels of accuracy that are expressed are different than those in Analysis 1. The analysis assumes that when an image returns a percentage of the classification of interest greater than the threshold value that the decision will be made to spray the area. This gives an output of the number of images where the wrong decision will be made in both the case of under spray and over spray. It is important to note that the percentage of under spray is based on the number of images in the series that contain weeds and the percentage overspray is based on the number of images with no weeds. The total number of images is not used because it is not possible to under spray an area with no weeds and it is not possible to overspray an area with weeds.

The results from this analysis vary depending on the threshold value that is set. Therefore the data that has been presented is based on a threshold value that achieves as close to 5% under spray value possible. 5% under spray was considered a reasonable and realistic tolerance. 0% under spray is not targeted because in many situations when 0% is under sprayed close to 100% is over sprayed. This means that it would be difficult to compare classifiers.

An assessment of the effect of the proportion of weed to weed-free images has also been conducted. This was achieved by taking a series of 1000 images that contained 72 images with broadleaf weeds (BLW). Images without weeds were removed from the series so that the total number of images was reduced to 500 and 250 but the number of images with BLW was maintained at 72. Comparisons of the results were then made.

Analysis 3 - Spatial Analysis

The spatial analysis involves comparing the sensor output to physical plant count data taken from either the field using GPS locations or from plant counts taken from the images collected during the scanning process. Five images and five 0.1m² quadrat counts are taken from a given point and the paddock, the images are processed through the H sensor and a regression analysis is calculated from

the data. Alternatively a map is generated from the sensor data and the sensor value for a given point is recorded and tested with regression against quadrat or image plant counts.

iii. Classifier training method

The process of training the sensor or improving a particular Classifier involves taking a representative series of 100 – 200 images with a mixture crop and weed plants and running them through the H sensor software with the classifier to be improved. This process generates a file which contains a list of all of the individual features or plant segments for each image. This file is then used in a second software package to manually label all representative individual features with the correct classification. The output from the second software package along with the raw images is then sent to Agricon for further processing. Analysis 1 is then used to make objective assessments of the difference between the original and updated classifiers.

iv. Sensor scanning area calculations

The H sensor operates by taking still photographs and processing the imagery in real time. The sensor is able to capture and process up to 10 images per second and each image is approximately 0.075 m². However as the number of leaf shapes in the image increase, so too does the processing requirements and processing time, this reduces the image frequency in areas where weed density is high.

The still photograph method of scanning means that not all of the paddock can be scanned. Therefore, the results from the scanned area must be interpolated over the remaining paddock area. Table iv demonstrates how much of a given paddock will be scanned for a given speed and camera spacing. Agricon suggest operating four sensors on a 24m boom, giving 6m camera spacing. If the camera spacing is increased to one sensor per boom, in an Australian context where many booms are 36m, the area scanned decreases by a factor of 6.

Table iv: The percentage of a paddock scanned with the H sensor for a range of image capture rates and ground speeds at camera spacing 6 and 36m.

% area scanned at camera spacing 6m								
Images/second		4	5	6	7	8	9	10
km/h	10	2.1%	2.7%	3.2%	3.7%	4.3%	4.8%	5.3%
	11	1.9%	2.4%	2.9%	3.4%	3.9%	4.4%	4.8%
	12	1.8%	2.2%	2.7%	3.1%	3.6%	4.0%	4.4%
	13	1.6%	2.0%	2.5%	2.9%	3.3%	3.7%	4.1%
	14	1.5%	1.9%	2.3%	2.7%	3.0%	3.4%	3.8%
	15	1.4%	1.8%	2.1%	2.5%	2.8%	3.2%	3.6%
	16	1.3%	1.7%	2.0%	2.3%	2.7%	3.0%	3.3%

% area scanned at camera spacing 36m								
Images/second		4	5	6	7	8	9	10
km/h	10	0.3%	0.4%	0.5%	0.6%	0.7%	0.8%	0.9%
	11	0.3%	0.4%	0.5%	0.5%	0.6%	0.7%	0.8%
	12	0.3%	0.4%	0.4%	0.5%	0.6%	0.6%	0.7%
	13	0.3%	0.3%	0.4%	0.5%	0.5%	0.6%	0.7%
	14	0.2%	0.3%	0.4%	0.4%	0.5%	0.6%	0.6%
	15	0.2%	0.3%	0.3%	0.4%	0.5%	0.5%	0.6%
	16	0.2%	0.3%	0.3%	0.4%	0.4%	0.5%	0.5%

v. Summary of results: Analysis 1 - The Plant Feature Method

The following tables show a summary of results for the best classifiers for a given image series for dicot crops with ryegrass as weeds and monocot crops with broadleaf weeds.

For more detailed information refer to the corresponding report section in the document body.

Table v a: The percent of correctly labelled plant features for ryegrass in dicot crops.

Report section		1	2	3	4	5	6	13
Crop		Canola	Lupin	Faba bean	Field pea	Lentil	Chickpea	Canola
Weeds		Ryegrass	Ryegrass	Ryegrass	Ryegrass	Ryegrass	Ryegrass	Ryegrass
Best classifier		WW3.1	WW3.0	WW3.0	WW3.0	W3.0	WW3.0	WW3.1
Number of segments	Type of segment							
	All	96	78	88.1	90	91	95	86
	Weed	91	85	86	93	90	90	76
	Crop	99	72	98.5	91	93	97	96
	Undefined	*	64	*	*	53	97	95
Area of image (Pixel number)	All	93	80	93	99	73	96	97
	Weed	59	76	63	96	82	93	85
	Crop	97	86	99.9	99	98	97	98
	Undefined	*	32	*	*	3	55	100

Table v b: The percent of correctly labelled plant features for dicot weeds in monocot crops.

Report section	7	7	8	8	8	9	15	10	10
Crop	Barley (GS12)	Barley (GS22)	Wheat (GS22)	Wheat	Wheat	Oat	Wheat	Wheat in canola stubble	Wheat with canola stubble removed
Weeds	Wild radish	Wild radish	Tares, bifora, medic	Wild radish	Wild radish	Wild radish	Wild radish	Tares, bifora, medic	Tares, bifora, medic
Best classifier	WW2.0	WW2.0	WW3.1	WW3.0	WW1.0	WW3.0	WW2.1	WW2.0	WW2.0
Type of segment									
All	78	75	78	79	84	79	81	57	47
Weed	70	75	87	96	88	99	77	46	45
Crop	96	78	74	66	61	56	82	69	49
Undefined	*	5	36	23	0	66	6	5	43
All	88	80	35	68	88	18	88	37	41
Weed	87	88	76	89	95	96	69	25	26
Crop	91	66	86	51	52	9	95	87	84
Undefined	*	1	6	39	0	75	0	3	4

vi. Summary of Results: Analysis 2 - The Threshold Method

Table vi a: Summary of results for the most accurate classifier for grass weeds in broadleaf crops using analysis 2 - the threshold method.

Report section	1	2	3	4	5	6	13
Crop	Canola	Lupin	Faba bean	Field pea	Lentil	Chickpea	Canola
Weeds	Ryegrass	Ryegrass	Ryegrass	Ryegass	Ryegrass	Ryegrass	Ryegrass, brome, wheat
Best Classifier	WW1.0	WW3.1	RAPS1.0Ud	WW3.0	WW3.1	WW3.0	RAPS1.0
Total images in series	256	161	437	80	299	60	216
Images under sprayed (%)	6	5	46	3	5	3	8
Images over sprayed (%)	53	55	14	53	70	54	75

Table vi b: Summary of results for the most accurate classifier for broadleaf weeds in grass crops using analysis 2 - the threshold method.

Report section	7	7	8	8	9	16
Crop	Barley	Barley	Wheat	Wheat	Oat	Wheat
Weeds	Wild radish	Wild radish	Tares, bifora, medic	Wild radish	Wild radish	Field pea, tares, wild radish
Best Classifier	WW2.0	RAPS1.0	MAIS2.3	WW3.0	MAIS2.2	WW3.1
Total images in series	144	1000	500	998	300	300
Images under sprayed (%)	6	6	5	5	6	6
Images over sprayed (%)	59	90	68	88	81	57

vii. Summary of Results: Analysis 3 - Spatial Analysis

Table vii: The R squared values of regressions between H sensor output and weed plant density for the tested classifiers.

Test	Crop	Weed	Classifier	R2 value	Comments
1	Canola	Ryegrass	WW3.0	0.84	
2	Lupin	Ryegrass	WW3.0	0.03	Lupin leaflets look similar to grass
3	Bean	Ryegrass	WW3.1	0.04	Small ryegrass could not be detected
4	Chickpea	Ryegrass	RAPS1.0	0.46	
5	Barley	Radish	WW2.0	0.66	
6	Barley	Radish	WW2.0	0.85	
7a	Wheat	Bifora, Medic, Tares	WW2.0	0.34	
7b	Wheat	Bifora, Medic, Tares	WW2.0	0.76	Outliers removed
8	Wheat	Radish	WW2.0	0.02	Too much leaf occlusion to identify radish
9a	Bean	Ryegrass	RAPS1.0	0.002	Small ryegrass could not be detected
9b	Bean	Ryegrass	RAPS1.0	0.43	Outliers removed
10a	Bean	Ryegrass	WW3.1	0.0003	Small ryegrass could not be detected and thick ryegrass areas identified as dicot
10b	Bean	Ryegrass	WW3.1	0.05	Outliers removed, thick ryegrass areas identified as dicot

8 Wheat, tares (*Ciccia sativa ssp nigra*), bifora (*Bifora testiculata*), snail medic (*Medicago scutellate*)

Grower: Dare, Paddock: Long

Wheat growth stage: 4 leaf 2 tillers

Weed growth stage: tares, up to 7 node, bifora, 2 true leaves, medic, up to 8 node

Initially the WW3.1 classifier was used for analysis one and found to be not working well due to high levels of leaf occlusion and miss classification of the thick wheat row. After contacting Agricon, WW1.0 was used and updated with training to improve the classification. The overall number of segments that were classified correctly increased with the updated WW1.0 compared to the original WW1.0, however this still did not meet the overall level of accuracy achieved with WW3.1. Although the overall accuracy was not as great with the updated WW1.0 classifier, this classifier had the highest accuracy when looking at only the segments labelled as broadleaf weeds (91%). Because the overall accuracy for WW1.0 update is only 33%, it means that many of the non broad leaf weed segments are classified incorrectly as broad leaf which will lead to an over spray decision in the absence of weeds.

Table 8a results from analysis one for classifiers WW1.0, WW1.0Update, WW2.0 and WW3.1, results expressed as percent of segments and segment area correct (Dare, Long Pdk).

		Classifier							
		Original WW1.0		Updated WW1.0		WW2.0		WW3.1	
		Number of segments	Area (Pixels)	Number of segments	Area (Pixels)	Number of segments	Area (Pixels)	Number of segments	Area (Pixels)
All segments	T total segments/area labelled manually	8108	27083955	8108	27083955	8108	27083955	4346	10848079
	Segments/area classified correctly	3703	12844133	4260	8977466	4327	8763419	3385	3773714
	% correct	46%	47%	53%	33%	53%	32%	78%	35%
Segments labelled Crop	T total segments/area labelled manually	4059	8181180	4059	8181180	4073	8183204	1834	1233805
	Segments/area classified correctly	788	774140	971	5379357	2037	7430513	1597	935306
	% correct	19%	9%	24%	66%	50%	91%	87%	76%
Segments labelled Broadleaf weed	T total segments/area labelled manually	3576	2862974	3576	2862974	3582	2864531	2297	2736264
	Segments/area classified correctly	2764	2153768	3268	2540253	2283	1114174	1695	2345494
	% correct	77%	75%	91%	89%	64%	39%	74%	86%
Segments labelled UNDEF	T total segments/area labelled manually	453	16036220	453	16036220	453	16036220	192	6792846
	Segments/area classified correctly	151	9916225	21	1057856	7	218732	70	407750
	% correct	33%	62%	5%	7%	2%	1%	36%	6%

Table 8b, results from analysis two expressed as total correct, under spray, and over sprayed images. Under spray % is calculated from images with weeds that are not sprayed and over spray % is calculated from images without weeds that are sprayed (Dare, Long Pdk).

Total number of images	Images with weeds	Images without weeds	Classifier	Threshold value %	Number of images sprayed	% Overall correct	% Under sprayed	% Over sprayed
500	449 90%	51 10%	WW1.0	0.01	361	75.6	24.0	29.3
			WW3.0	7.00	772	88.2	5.0	87.8
			WW3.1	4.50	474	87.4	5.2	95.1
			MAIS2.2	0.20	469	88.4	5.2	95.1
			MAIS2.3	0.33	463	89.6	5.2	68.3
			RAPS1.0	0.07	477	87.2	5.0	100.0
			RAPS1.1	0.13	475	87.6	5.0	95.1




Figure 8a, map of H sensor output (WW2.0) for tillering wheat (Dare, Long, 2015) and broadleaf weeds including tares, bifora and volunteer canola. Point data shows locations and broadleaf weed plant density from quadrat counts (plants/m²). Top legend = broadleaf weed plant density at quadrat points, bottom legend = H sensor DICOT output WW2.0.

Analysis 3 shows that the raw data from the sensor was not well correlated with actual weed density. For each point on the graph in figure 8b the H sensor DICOT pixel area is produced from five sensor images. At six points one of the five images produced a very high DICOT sensor output (> 10%). When the images with DICOT pixel area > 10% were removed from the data the relationship improved to produce an R^2 value of 0.76 (Figure 8c). However this relationship relies on the single point of 450 plants/m², without this point the relationship is poor.




Figure 8b, broadleaf weed plant density compared to H Sensor output of DICOT pixel (% of image) raw data $y = 0.0286x + 0.6931$, $R^2 = 0.34$.




Figure 8c, broadleaf weed plant density compared to H Sensor output of DICOT pixel (% of image) with the data edited to remove any DICOT pixel area above 10%, $y = 0.0035x + 0.0525$, $R^2 = 0.76$.

Reasons the sensor performed poorly in 4 leaf cereal

The sensor has two problems in cereals at around the 4 leaf growth stage.

1. The first problem is that as the leaves are overlapping they cannot be identified individually and therefore the shapes of the wheat leaves and plants are not consistent. The sensor relies on features of plants of the same type having similar shapes to make the classification. The result of this is often some images of wheat being classified as large broadleaf weeds (figure 8d).
2. With large amounts of overlap weeds, in this case cotyledon wild radish, are occluded by the cereal leaf and the leaf shape cannot be distinguished for classification.




Figure 8d) example image of 4 leaf wheat and medic (left) from H sensor and the result image from WW3.0 classifier showing leaf occlusion and misclassification (right)