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 think 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 oaten 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 with in 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 0.1m^2 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 of 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 m^2. 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**
	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%
km/h	**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**
	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%
km/h	**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
	Type of segment							
Number of segments	**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									
Number of segments	**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
Area of image (Pixel number)	**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 senor 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 occulsion 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 idenified as dicot
10b	Bean	Ryegrass	WW3.1	0.05	Outliers removed, thick ryegrass areas identified as dicot

13 Canola and grasses

Grower: J Venning

Paddock: Copleys Middle

Growth stages: Canola 3-4 true leaves, ryegrass 1-3 leaf, brome grass 1-2 leaf, wheat 2 leaf-mid tillering

This series of images was collected from a 3 leaf canola crop growing on a heavy wheat stubble from 2015 and was not inter-row sown. There was approximately 5t/ha of stubble present and it was laying on the ground over most of the paddock, although in some small areas it was still standing.

Classification of some of the straw in the images as grass has caused some over spray of images in the series. 96 of the 216 images had at least 1 piece of straw identified in the image as green plant material and almost all of these were classified as grass weeds. When the image contained grass weeds the misclassification of straw did not cause any issues as the decision to spray remained the same. However if there was no actual grass in the image then the result to spray the image was incorrect. In almost all cases of this problem the features in question appear in the upper left quarter of the image. Agricon are aware of the problem and in their situation they do not believe it to be an issue. However I believe that in this situation it would lead to significant increases in overspray decisions. As has been demonstrated in previous image series the WWx.1 classifiers produce more of this misclassified straw.

Table 13a shows the results from Analysis 1 for the WW3.1 and RAPS1.0Update classifiers. It demonstrates that as with the first image series (001), generally both of these classifiers perform well at separating canola and ryegrass when looking at individual features. Accuracy of identifying small broadleaf weeds however is poor with only 35% of the broadleaf weed pixel area being classified correctly using RAPS1.0Update. For this reason a second table has been created where both canola and broadleaf weeds have been treated as the same plant type (as was done for series 001). Although analysis 1 shows that the sensor can be accurate at identifying most of the individual features, analysis 2 shows that the misclassification is still significant enough to result in the wrong spray decision with the amount of images that are incorrectly classified as having grass in them high at 64.8% for WW3.1 and 60.2% for RAPS1.0Update.

Table 13a, Results of Analysis 1 for series 030 for the classifiers WW3.1 and RAPS1.0Update.

		Classifier			
		WW3.1		RAPS1.0Update	
		Number of segments	Area (Pixels)	Number of segments	Area (Pixels)
All Segments	Total segments/area labelled manually	2634	8138304	5036	14879920
	Segments/area classified correctly	2260	7886933	3279	12652934
	% correct	86%	97%	65%	85%
Segments labelled Grass	Total segments/area labelled manually	375	257037	428	292465
	Segments/area classified correctly	285	217630	272	216914
	% correct	76%	85%	64%	74%
Segments labelled DICOT	Total segments/area labelled manually	2043	7825165	364	532230
	Segments/area classified correctly	1955	7660204	210	184205
	% correct	96%	98%	58%	35%
Segments labelled BRSNN (canola)	Total segments/area labelled manually	-		4153	13970374
	Segments/area classified correctly			2794	12179022
	% correct			67%	87%
Segments labelled DICOT and BRSNN	Total segments/area labelled manually	-		4517	14502604
	Segments/area classified correctly			3004	12363227
	% correct			67%	85%
Segments labelled Straw	Total segments/area labelled manually	195	29934	91	12058
	Segments/area classified correctly	0	0	0	0
	% correct	0%	0%	0%	0%
Segments Labelled UNDEF	Total segments/area labelled manually	21	26138	3	72793
	Segments/area classified correctly	20	9099	3	72793
	% correct	95%	35%	100%	100%

Table 13b, Results of Analysis 1 for series 030 for the classifiers WW3.1 and RAPS1.0Update. The results for RAPS1.0Update have been converted so that canola (BRSNN) and broadleaf weeds (DICOT) have been combined. This is done so that both classifiers are only using two categories.

		Classifier			
		WW3.1		RAPS1.0Update	
		Number of segments	Area (Pixels)	Number of segments	Area (Pixels)
All Segments	Total segments/area labelled manually	2634	8138304	5036	14879920
	Segments/area classified correctly	2260	7886933	3958	14551444
	% correct	86%	97%	79%	98%
Segments labelled Grass	Total segments/area labelled manually	375	257037	428	292465
	Segments/area classified correctly	285	217630	272	216914
	% correct	76%	85%	64%	74%
Segments labelled DICOT	Total segments/area labelled manually	2043	7825165	4153	13970374
	Segments/area classified correctly	1955	7660204	3958	13741059
	% correct	96%	98%	95%	98%
Segments labelled Straw	Total segments/area labelled manually	195	29934	91	12058
	Segments/area classified correctly	0	0	0	0
	% correct	0%	0%	0%	0%
Segments Labelled UNDEF	Total segments/area labelled manually	21	26138	3	72793
	Segments/area classified correctly	20	9099	3	72793
	% correct	95%	35%	100%	100%

Table 13c, Results from Analysis two for image series 030 for classifiers RAPS1.0 original and update and WW3.0 and 3.1

Total number of images	Images with weeds	Images without weeds	Classifier	Threshold value	Number of images sprayed	% Overall Correct	% Under sprayed	% Over sprayed
216	108 (50%)	108 (50%)	RAPS1.0 Original	0.01	180	58.3%	8.3%	75.0%
			RAPS1.0Update	0.01	157	62.5%	14.8%	60.2%
			WW3.0	0.01	138	63.9%	22.2%	50.0%
			WW3.1	0.01	160	59.3%	16.7%	64.8%