Robustness of Plant Recognition Based on Active Shape Models

Maria Johansson
Halmstad University
Box 823
S-301 18 Halmstad
SWEDEN
Maria.Johansson@ide.hh.se

Henning T. Søgaard
Danish Institute of Agricultural Sciences
Research Centre Bygholm
Box 536
DK-8700 Horsens
DENMARK
HenningT.Sogaard@agrsci.dk

Albert-Jan Baerveldt
Halmstad University
Box 823
S-301 18 Halmstad
SWEDEN
Albert-Jan.Baerveldt@ide.hh.se

***Abstract* – Computer vision is often used in research of plant recognition, as a camera is a very powerful sensor at a relatively low cost. Recently, the use of active shape models for plant recognition has shown promising results when used for off-line weed mapping. The advantage is that active shape models cover the whole plant and do not require any segmentation of the image. Our objective is to use active shape models in an on-line weeding robot, which requires real-time implementation. Therefore, the number of searches required to obtain the best match needs to be minimized. This makes demands on where to start the searches relatively to the object to be recognized. To start a search, an initial active shape model must be placed at the starting position of the search. The initial model is described by a set of pose parameters, namely the position, the rotation angle and the scale. Preliminary values of these parameters must be estimated prior to the search. The sensitivity of the search result, depending on these parameters, is analyzed in order to investigate how accurate the initial model needs to be placed. This paper presents an investigation of the robustness of the search result depending on these parameters when using active shape models for classifying sugar beet plants. An active shape model has been built, using 14 images of sugar beet plants. An evaluation of the sensitivity in position, rotation angle and scale is performed on a test set of 50 images. The test results show a high correspondence between the plant and the reshaped model if the size parameter sets an initial model that is less than twice the plant size, the position is centered within a radius of 23% of the plant size and the rotation angle of the initial model is +/- 18.5 degrees from the plant angle.**

I. INTRODUCTION

Active shape models are reshapeable templates that only can reshape according to some predefined criteria and have been used in many applications for interpreting medical images [1]. The models are suitable for modelling changeable objects such as hands, insects, leafs, etc. In the area of plant recognition the active shape models have been used for classifying weed species [2]. Søgaard has created models for approximately 20 weed species and has during classification tests got a performance rate of 80% when classifying what weed specie a plant belongs to. So far, this classification has only been used off-line on pre-sampled images. To be able to use the active shape models in a real-time system together with a weeding robot for sugar beet fields [3], the numbers of searches required to get a good match need to be minimized. Minimizing the number of required searches places demands on the placement of the initial model prior to the searches. The placement of the initial model is described by a set of pose parameters that need to be estimated prior to the search. The pose parameters are the position (x, y coordinates within the image), the rotation angle and the scale.

II. OBJECTIVES

The objective of this study is to use an active shape model for recognition of sugar beet plants in order to analyse how the search result depends on the initial pose parameters of the model.

III. METHOD

A. Building the model

The active shape model consists of a number of points where each point belongs to a specific part of the object to be modelled. Statistics describing how the positions of the points can vary relatively to each other is achieved through a set of labelled training images, where a number of points have been marked along the edges of the training object. In this case, 80 points are used to describe sugar beet plants having the first two true leaves developed. Since the plants have four leaves, two cotyledons and two true leaves, 20 points are used for each leaf. Points 0-19 describe the upper leaf (true leaf), points 20-39 the right leaf (cotyledon), points 40-59 the bottom leaf (true leaf) and points 60-79 the left leaf (cotyledon). A labelled training image can be seen in Fig 1.

To get a larger training set and a more symmetric model, the images are mirrored horizontally and vertically. Fig 2 shows the four resulting images from one of the images in the training set. The original image is in the upper left corner.

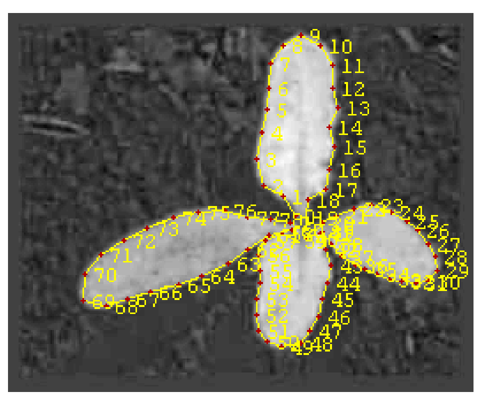

Fig 1. Labelled training image.
Points 9, 29, 49 and 69 mark the tip of each leaf.

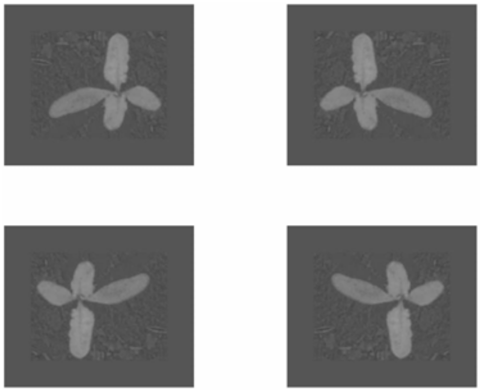

Fig 2. Example of how a larger training set can be obtained by using mirrored copies of the original images.
Upper left: original image. Upper right: vertically mirrored. Lower left: vertically and horizontally mirrored. Lower right: horizontally mirrored.

The active shape model is built when all training images are labelled. Building the model results in a mean model and a number of modes of variation. The mean model is the shape created by the mean position for each point. The modes of variation describe how the model can reshape itself and are the outcome of a principal component analysis of the deviations between the marked points and the mean model. Based on the active shape model, new model shapes can be constructed from the mean model and a number of shape parameters controlling one mode of variation each.

Fig 3 shows how the model reshapes when changing values for one shape parameter at a time. Each row in the matrix represents one mode of variation. The modes are sorted in significance with the most significant mode on the lowest row. The middle column shows the mean model and the other columns show the reshaped model when decreasing/increasing the shape parameters.

More information of how to build the model and how the modes work can be found in [1], [4] and [5].

B. Training set and test set

The training set and the test set consist of images of sugar beet plants with the first two true leaves developed. The original images are in RGB but since the active shape models work on grey level, the images are converted to grey by I = 2G – R - B.

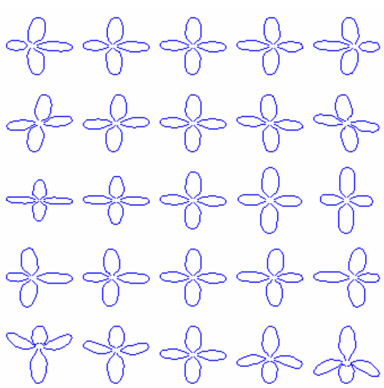

Fig 3. The first five modes of the model

Each image contains one sugar beet plant and no weed.

The training set consists of 56 images, 14 original images and the 42 mirrored copies. All the training plants have the same rotation, the cotyledons horizontal and the first true leaves vertical. The test set consists of 50 images and these plants all have different rotations.

For comparison there is also a set of 50 images where each image contains one weed plant and no sugar beet plants.

C. Searching an image for a plant

Prior to a search the model is placed over the plant at a given position (x, y coordinates) and with a given rotation angle and scale. When a search is performed, the model is adapted to the object by changing these pose parameters and by reshaping the model according to the limitations in the modes of variation.

To quantify how well the model manages to adapt to the plant, the ratio between two areas is calculated. The area ratio is the ratio of the number of pixels covered by both plant and model to the number of pixels covered by plant or model (intersection/union) [2]. This area ratio will tend to one for good matches and to zero for bad matches. It has in [2] been used as a good feature for classifying weed species.

Fig 4 and Fig 5 show the reshaped model after a search. The line marks the contour of the reshaped model. In Fig 4 the reshaped model does not follow the contour of the sugar beet plant. The area ratio in this search is 48%. Fig 5 on the other hand, shows a search result where the area ratio is 93% and the reshaped model follows the contour of the sugar beet plant closely. For comparison, initial tests with the sugar beet model show that the mean area ratio for the weed test set is 60% and for the sugar beet test set 77%.

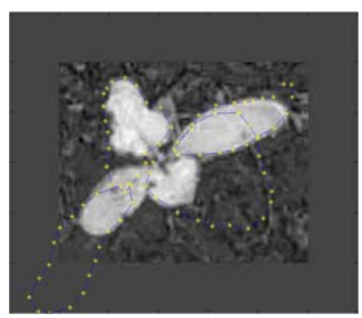

Fig 4. Deformed model after search. Area ratio 48%

Fig 5. Deformed model after search. Area ratio 93%

D. Test parameters for placing the model

Testing the parameters for placing the model becomes a multidimensional problem with four dimensions: position (two dimensions), rotation and scale. To find the sensitivity of the parameters, they need to be tested one by one. When investigating each parameter, the remaining parameters are set to some predefined values.

- The predefined value for the scale value is set to the ratio of the longest side of the bounding box[1] for the plant in the test to the longest side of the bounding box for the mean model.
- For the rotation angle, eight predefined values are used; 45, 90, 135, 180, 225, 270, 315 and 360 degrees.
- For the position, three different predefined values are used; bounding box centre, centre of gravity for the segmented sugar beet and root position. Bounding box centre and centre of gravity are calculated from the threshold image, where the background is separated from the plant. The root position, where the stem is expected to enter the soil, is marked manually for all images in the test set.

E. Testing the area ratio with respect to variations in scale

When testing the scale parameter the predefined values for position and rotation are used. Scales from 0.1 to 5 are used, where the step size for scale values less than 1 is 0.1, and for scales above 1 is 0.5. For each position and scale, all eight rotation angles are tested and the result from the rotation angle with the highest area ratio is stored for evaluation.

F. Testing the area ratio with respect to variations in rotation angle

The parameter rotation angle is tested for the predefined values of position and scale. To test the rotation angle the model is rotated one degree at a time, from 1 to 360 degrees. For each predefined position the area ratio for all rotation angles is stored for evaluation.

G. Testing the area ratio with respect to variations in position

The sensitivity to the initial position is tested by placing the model at different positions within the bounding box of the plant. For the first search the centre of the bounding box is used as the initial model position and, subsequently, positions closer to the edges of the bounding box is used. The offset to the centre is measured in percent of half the bounding box size. The used percentages are 1, 2, 3, 4, 5, 10, 25, 50 and 75. Fig 6 shows where in the bounding box the model is positioned.

For the parameter scale, the predefined value is used, and all eight predefined rotation angles are tested. The rotation angle that gives the highest area ratio is stored for evaluation.

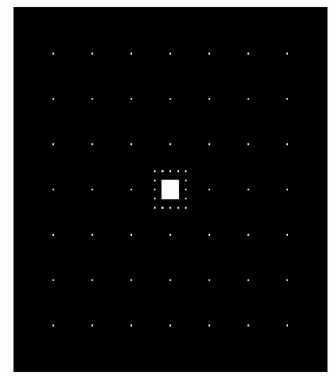

Fig 6. Positioning of the initial model within the bounding box. The black area is the area covered by the bounding box and the white spots marks where the centre of the initial model will be positioned. The white rectangle in the centre represents 11 x 11 marks and marks the positions for -5% to 5% in vertical and horizontal directions. The surrounding rectangle is the 10% markers and the rest are the 25%-75% markers.

IV. RESULT

A. Grading the area ratios

From manually examination of the search results the area ratio is empirically divided into three grades; high, medium and low. The limits for the grades are

- High: 80-100%
- Medium: 60-80%
- Low: 0-60%

B. Result for parameter scale

The result for all three predefined positions is plotted in Fig 7. The y-axis represents the mean value of the area ratio of the sugar beet images in the test set and the x-axis represents the scale. For a scale less than 2, these images give a high area ratio. The scale value only scales the model. To compare the size of the mean model to the size of the plants in the test set, the ratio of the longest side of the bounding box for the plant to the longest side of the bounding box for

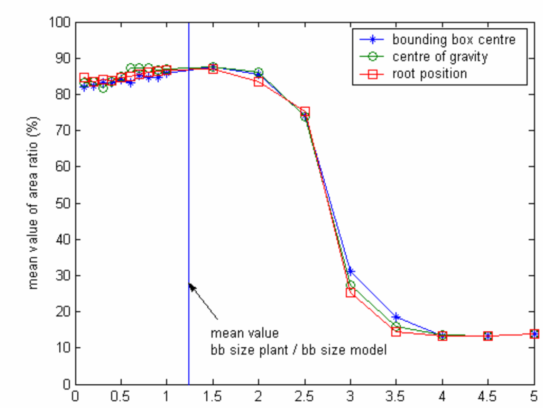

Fig 7. Plot of mean value of area ratio when the scale parameter is set to values from 0.1 to 5

[1] The bounding box is the smallest rectangle that encloses the object.

the mean model is marked in the plot. The mean value of this ratio over the images in the test set is 1.24 and the standard deviation is 0.11. Comparing the mean value of this ratio to the test results show that a scale corresponding to a model that is up to almost twice the plant size gives a high area ratio.

C. *Result for the rotation parameter*

Plotting the area ratio versus the rotation angle for one image gives a result as Fig 8. This graph would be expected to be periodic with a period of 180 degrees since the model is invariant to rotations of 180 degrees. This is, however, not the case since the number of points per leaf is even and the points do not become symmetrically distributed around the leaf (1 point at tip of leaf, 9 points on one side and 10 points on the other side).

The maximum area ratio for each plant in the test set ranges between 72% and 94%. To find out how much the model can be rotated, the regions that have an area ratio larger than a certain threshold are calculated for each image. For each image, the largest region indicates the tolerance for that image. The threshold is set to 70% to get a tolerance level greater than zero for all images.

In Fig 9 the tolerance is plotted. On the x-axis is the size of the largest region with area ratios over 70% and on the y-axis is how many images that have regions of at least that size. The results for all three predefined positions are plotted in the same graph. Using the bounding box centre as initial model position, the angle tolerance is +/- 18.5 degrees for 90% of the images. With the centre of gravity and the root position as initial model position the corresponding angle tolerances are +/- 16.5 degrees and +/- 12 degrees, respectively.

D. *Result for parameter position*

The result from testing the relationship between the initial model position and the search result can be seen in Fig 10. The area ratio is plotted as a function of the percentage distance between the initial model position and the centre of the plant bounding box (mean result over all the test images). The area ratio is generally decreasing with increasing

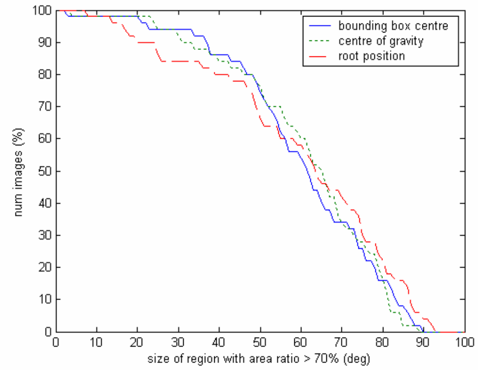

Fig 9. Number of images where largest region with area ratio larger than 70% is larger than the value at the x-axis.

distance, but up to a distance of 46% of half the bounding box length the area ratio is at least 80%.

For the test set, the distance from bounding box centre to the root position has a mean value of 8% with a standard deviation of 6% and the mean value for the distance from bounding box centre to centre of gravity is 15% with a standard deviation of 12%. Comparing these numbers with Fig 10 shows that using the root position or the centre of gravity will most often result in high area ratio.

V. SUGGESTED METHODS

Previous applications say very little about how the initial model was placed in automatic applications. Most reports do not mention how the initial model is placed while other asks the user to point out specific points in the images [6]. The strategies are dependent on the applications, the objects to be recognized and the extent of human interaction.

In this work, the search result was least sensitive to the scale parameter, while it is more sensitive to the rotation angle and the position. According to the test results, the area ratio will be high if the scale parameter results in a model that is smaller than twice the plant size, the model centre is placed within 46% from the bounding box centre and the rotation angle is within +/- 18.5 degrees from the best rotation angle.

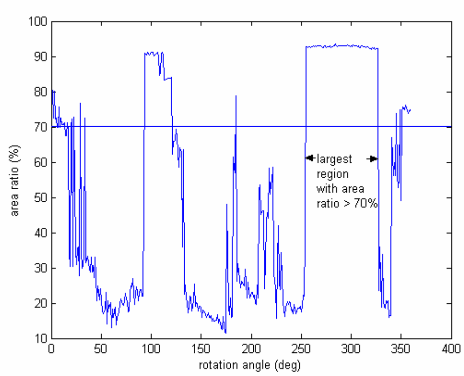

Fig 8. Plot of area ratio for each degree when rotating model 1-360 degrees

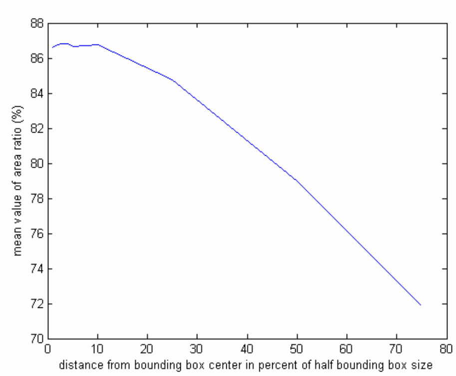

Fig 10. Area ratio as a function of the distance from the bounding box centre.

A. *Estimating the scale parameter*

Bounding box relation – The scale estimate could be based on the size of the bounding box of the plant by using the ratio between the longest side of the bounding box of the plant and the longest side of the bounding box of the mean model. In this way we should always have a model that is smaller than twice the size of the plant of our choice.

Manual estimate - Another approach is to manually estimate the mean value of the size of the plants to be tested. Then the scale should be set to a value that gives a model the same size as most of the plants, e.g. the estimated mean plant size divided by the mean model size. When using this approach, there is no need to calculate the bounding box of the plant, which can be hard to do when not knowing exactly what objects in the image that belongs to the same plant.

B. *Estimating the rotation parameter*

From the test results concerning the rotation parameter, the conclusion is that it is more important to have a good estimate of this parameter than of the scale parameter. Moreover, if the rotation angle is correct, with respect to where the model is positioned, the search will give a high area ratio.

Largest leaf - A value of the rotation angle may be calculated from the angle of the main axis of the largest leaf. Fig 11 shows a comparison of this angle to the result of the rotation test. The estimated angle of the plant is derived from the rotation test by identifying the centre of the largest angle region resulting in area ratios greater than 70%, see Fig 8. For this comparison, the largest leaf has been found by manually marking the leaves and picking the leaf with the largest area in the segmented image. When comparing the angle of the largest leaf to the estimated angle of the plant, only 46% of the plants in the test set will be within the limit +/- 18.5 degrees. If the model is rotated 180 degrees, 66% of the plants will be within the given limits. This result is too low to be the only way of determining the rotation angle.

Extra rotations - The result can be improved by repeating the search with different rotation angles, but then the number of searches will increase.

C. *Estimating the position parameter*

RTK-GPS - When the plants are sown, a RTK-GPS can register where the seed is placed. The RTK-GPS has an error of a few centimetres and the seed might move a little bit from where the precision seed drill puts it. Further more the sugar beet plant may not emerge exactly where the seed was placed due to the soil structure. These “noise” effects result in a displacement of the sugar beet plant of approximately 32-59 mm from where the RTK-GPS indicates that the seed was placed [7]. A sugar beet with the first two true leaf developed has a size of approximately 5 cm. This means that the model might be misplaced with up to over 200 % which is more than the limit of 46%, thus this method is not recommended.

Bounding box centre - If the bounding box of the plant is found a search from the centre of the bounding box will give a high area ratio. The problem is to find the correct bounding box when there is more than one plant in the image and the image also contains weeds. The bounding box can also be distorted if the plant is occluded by other leaves. An occluded plant can give a displacement of the bounding box centre by more than 100%.

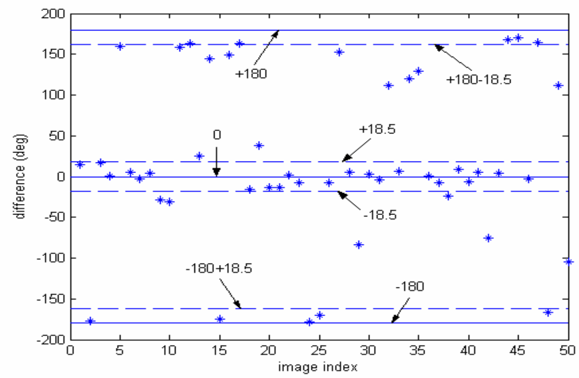

Fig 11. Difference between the angle of largest leaf and the estimated angle of the plant. The tolerance of +/-18.5% is marked.

Centre of gravity - Using the centre of gravity for the segmented plant as position gives in the test similar result as the bounding box centre. It is more time consuming to calculate the centre of gravity than the bounding box centre. Consequently, in the choice between these two methods, the bounding box centre is the preferred.

D. *Test with suggested methods*

For positioning of the initial model the *Bounding box centre* is used. For estimation of the scale parameter the method *Bounding box relation* with the scale factor set to 0.8 is used. To get the best match the rotation angle is a combination of the suggested methods *Largest leaf* and *Extra rotations*. In the test different number of rotations and offset angles are used. Up to 16 rotations are tested and as offset angle either the largest leaf angle or a zero angle is used. The rotation angle that gives the highest area ratio is chosen as the best match.

Fig 12 shows the result from eight tests. The different tests have different rotation angles as follows:

- Test 1. 2 rotations, rotation angle = offset angle + n * 180 degrees, n = 1…2, offset angle = 0
- Test 2. 4 rotations, rotation angle = offset angle + n * 90 degrees, n = 1…4, offset angle = 0
- Test 3. 8 rotations, rotation angle = offset angle + n * 45 degrees, n = 1…8, offset angle = 0
- Test 4. 16 rotations, rotation angle = offset angle + n * 22.5 degrees, n = 1…16, offset angle = 0
- Test 5. 1 rotation, rotation angle = angle largest leaf
- Test 6. 2 rotations, rotation angle = offset angle + n * 180 degrees, n = 1…2, offset angle = angle largest leaf
- Test 7. 4 rotations, rotation angle = offset angle + n * 90 degrees, n = 1…4, offset angle = angle largest leaf
- Test 8. 8 rotations, rotation angle = offset angle + n * 45 degrees, n = 1…8, offset angle = angle largest leaf

It appears that test 3, 4, 7 and 8 have a high area ratio for at least 80% of the images. Comparing these tests show that:

- For test 3 and 8 the number of images with high area ratio differs with 2%. These tests have the same number of rotations, but test 8 also uses the angle of the largest leaf, therefore, in a real-time perspective, test 3 is better than test 8.
- Test 4 has the highest number of images, 98%, with high area ratio but it takes the most number of rotations, 16.
- The result for test 3 is 4% better than for test 7. Test 7 uses fewer rotations, 4 compared to 8, but instead the angle of the largest leaf has to be calculated.

VI. CONCLUSIONS AND OUTLOOK

The method of using active shape models for plant recognition has shown promising results for off-line weed mapping. To get the method to work in a real-time system, working with sugar beets, the number of searches needs to be minimized. This makes demands on the initial placement of the model in relation to the object to be recognised.

This paper shows an investigation of the initial pose parameters: position, scale and rotation angle. According to the test results, if the scale parameter makes the model smaller than twice the size of the plant to be tested, the position is centred within a radius of 46% of half the bounding box size and the rotation angle is within +/- 18.5 degrees from the plant angle the search will result in a high area ratio.

If the scale limit is exceeded, the search will most probably fail, no matter what the other two parameters are set to. Exceeding the limit for the rotation angle will lead to an unsteady result. Some angles will give a high area ratio while other angles will result in a low area ratio. The position parameter will give a decrease in the search result the farther away from the plant centre the model is positioned.

The results of the study lead to a proposal for estimation of the initial pose parameters. As a position estimate the centre of the bounding box of the plant is proposed. As an estimate of the scale, the ratio between the size of the plant bounding box and the size of the model bounding box is proposed. Finally it is proposed to use the angle of the largest leaf and some derived angles as estimates of the model angle.

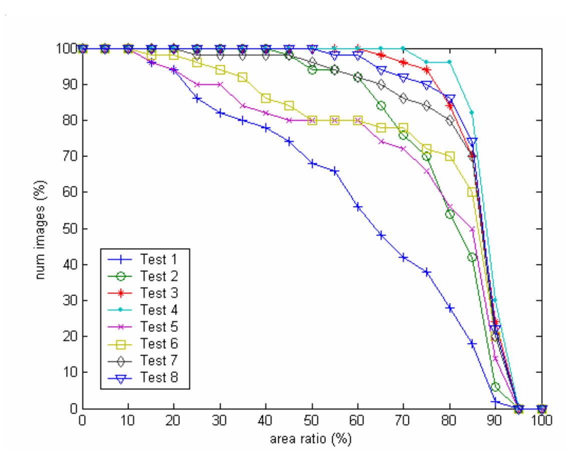

Fig 12 Result from test with suggested methods. Scale and position parameters are the same for all curves. Different initial offset angles and number of rotations is used for rotation angle parameter.

The suggested methods depend on if a bounding box can be found for the plant. Finding the bounding box can be difficult since we do not know what objects belongs to what plant.

Using the largest leaf for finding the rotation angle can also be tricky. It is likely that two or more leaves are connected and they may also be occluded by weeds.

It is proposed that the search is repeated with different initial rotation angles. The number of extra rotations can be decreased if the model is symmetric. The model will be more symmetric if the number of points per leaf is odd and the points are evenly distributed along the edges of the leaf. A symmetric model gives the same search result for the rotation angle 1-180 degrees as for 181-360 degrees.

The test images did not contain weeds so further research is needed to determine how well the suggested methods are working on images that contain both sugar beet plants and weeds.

VII. ACKNOWLEDGEMENTS

This work was carried out in the Mech-Weed project. The project is currently sponsored by SJV, Statens Jordbruksverk – Swedish Board of Agriculture, and SLF Stiftelsen Lantbruksforskning – Swedish Farmers' Foundation for Agricultural research.

VIII. REFERENCES

[1] T.F. Cootes, A. Hill, C.J. Taylor and J. Haslam, "The use of active shape models for locating structures in medical images", *Image and Vision Computing*, Vol. 12, No 6, July 1994, pp. 355-366.

[2] H.T. Søgaard and T. Heisel, "Weed Classification by Active Shape Models", in *AgEng 2002, Budapest, Hungary*, 02-AE-004.

[3] B. Åstrand and A-J. Baerveldt, "An Agricultural Mobile Robot with Vision-based Perception for Mechanical Weed Control", *Autonomous Robots*, Vol 13, 2002, pp. 21-35.

[4] T.F. Cootes and C.J. Taylor, "Active Shape Models – 'smart snakes'", in *Proc. British Machine Vision Conference*, Springer-Verlag, 1992, pp.266-275.

[5] T.F. Cootes, C.J. Taylor, D.H. Cooper and J. Graham, "Active Shape Models - Their Training and Application", *Computer Vision and Image Understanding*, Vol 61, No 1, January 1995, pp. 38-59.

[6] P.P. Smyth, C.J. Taylor and J.E. Adams, "Automatic Measurement of Vertebral Shape Using Active Shape Models", in *Proceedings British Machine Vision Conference*, Edinburgh, Great Britain, September 1996.

[7] H-W. Griepentrog, M. Nørremark, H. Nielsen, S. Blackmore, "Individual Plant Care in Cropping Systems", in *4th European Conference on Precision Agriculture*, Berlin, Germany, July 2003.