Pattern Analysis and Applications
Industrial and Commercial Application

Bayesian Classification and Unsupervised Learning for Isolating Weeds in Row Crops

François-Michel De Rainville · Audrey Durand · Félix-Antoine Fortin · Kevin Tanguy · Xavier Maldague · Bernard Panneton · Marie-Josée Simard

Received: 2 November 2011 / Accepted: 9 October 2012

Abstract This paper presents a weed/crop classification method using computer vision and morphological analysis. Subsequent supervised and unsupervised learning methods are applied to extract dominant morphological characteristics of weeds present in corn and soybean fields. The novelty of the presented technique resides in the feature extraction process that is based on spatial localization of vegetation in fields. Features from the weed leaf area distribution are extracted from the cultivation inter-rows, then features from the crop are inferred from the mixture model equation. Those extracted features are then passed to a naive bayesian classifier and a gaussian mixture clustering algorithm to discriminate weed from crop plant. The presented technique correctly classifies an average of 94% of corn and soybean plants and 85% of the weed (multiple species) without any prior knowledge on the species present in the field.

F.-M. De Rainville · A. Durand · F.-A. Fortin · K. Tanguy · X. Maldague
Laboratoire de vision et systèmes numériques
Département de génie électrique et de génie informatique
Université Laval, Québec (Qc), Canada G1V 0A6
E-mail: francois-michel.de-rainville.1@ulaval.ca, audrey.durand.2@ulaval.ca, felix-antoine.fortin.1@ulaval.ca, kevin.tanguy.1@ulaval.ca, xavier.maldague@gel.ulaval.ca

B. Panneton
Centre de recherche et de développement en horticulture
Agriculture et Agroalimentaire Canada
430 boul. Gouin, St-Jean-sur-Richelieu (Qc), Canada J3B 3E6
E-mail: bernard.panneton@agr.gc.ca

M.-J. Simard
Centre de recherche et de développement sur les sols et les grandes cultures
Agriculture et Agroalimentaire Canada
2560 boul. Hochelaga, Québec (Qc), Canada G1V 2J3
E-mail: marie-josee.simard@agr.gc.ca

Keywords Bayesian classification, Machine learning, Weed detection, Computer vision, Plant morphology, Corn and soybean fields

1 Introduction

In agriculture, weed control is a critical operation for maintaining crop yields. Under conventional farming practices, this is done by spraying herbicide on the entire surface of the field. According to CropLife Canada[1], herbicides account for 76% of the 1.91 billion CAD invested in pesticide in 2009 in Canada alone. In corn and soybean fields, herbicide is applied over the entire field even if weeds cover non-uniformly less than 30% of the field area. Furthermore, herbicides are known to have potential harmful consequences on the environment [5, 8]. To improve their efficiency, herbicides should be applied with the help of precision agriculture technologies specifically on weed-infested sites rather than uniformly on the field [23]. The main drawback for the widespread acceptance of this approach is the lack of economical and reliable application-specific sensors for automated data collection [21]. For post-emergence herbicide application, sensors capable of detecting weeds in presence of crops would be necessary to perform site-specific application or spot-spraying [12]. Computer vision has been intensively used for this kind of automatic visual detection task in multiple domains, from industrial inspection to process control. It is therefore the appropriate choice when it comes to identify, at *high speed*, portions of a field that contain weeds. Moreover, with cameras placed in front of a tractor, a proper detection system

[1] http://www.croplife.ca/web/english/plant_science_industry/

would be able to modulate the herbicide dose spread in function of the weed infestation level.

There exist two major trends to assess where herbicide should be applied using computer vision. The first one aims at identifying each particular plant as either weed or crop. Aitkenhead et al [1] present two methods to differentiate carrot (*Daucus carota* L.) plant from ryegrass (*Lolium perenne*) and Fat Hen (*Chenopodium album*) weeds. The first approach directly uses the leaf shape as the only discriminant and gives correct classification rates between 52% and 74%. The second method consists in a self-organizing neural network that discriminates plant types with an accuracy above 75% without prior knowledge. Hemming and Rath [9] use fuzzy logic to combine different measures of the structure of each plant in a digital image to classify them as either weed or crop. Their method, tested under field conditions, classifies between 51% and 95% of the plants correctly.

The second trend tries to identify patterns on crop rows and inter-rows to locate weed spots. Tellaeche et al [22] divide digital images of the field into cells delimited by crop row centres and fixed horizontal space. Each cell content is then provided to a trained support vector machine classifier to identify whether it should be sprayed with herbicide or not. They correctly identify 85% of the cells to be sprayed. Yang et al [24] propose a method which uses fuzzy rules on the coverage and weed patchiness to calculate the amount of herbicide to apply.

In this paper, we propose an intermediate approach that uses information from rows and inter-rows to learn the difference between crop plants and weeds. A naive bayesian classifier is used to discriminate crops from weeds based on statistics computed from row and inter-row leaf area. An unsupervised learning procedure is applied afterward to refine the preceding classification. The proposed technique novelty resides in exploiting the structure of the data to extract otherwise unavailable information on the vegetation. Two major advantages arise from this technique: no prior knowledge on crop or weed species is necessary nor does it require any prior training. Furthermore, in our real-time herbicide application context, methods requiring prior training are unacceptable in terms of work and financial effort.

The rest of the paper is organized as follow. First, Section 2 presents in details the problem addressed by this paper and the methodology adopted to resolve it. Then, Section 3 describes the whole feature extraction process, while Section 4 presents the naive bayesian classifier and the gaussian mixture model used to refine the classification. Finally, results of the application of our method on images of corn and soybean weed infested fields are presented in Section 5.

2 Rationale

As stated above, determining whether or not herbicides should be applied on a specific part of a field can significantly reduce the financial and environmental impact of weed treatment. This paper addresses the problem of discriminating weeds from crops using computer vision and machine learning in order to evaluate the degree of infestation in a section of a field.

The proposed technique applies to fields where crops are seeded in rows. All photographs are acquired early in the growing season corresponding to the appropriate phenological stage for post-emergence herbicide application. At this time, there is little foliage overlap since all plants are usually relatively small. Hence, the green area observed on digital images taken horizontally can be representative of the leaf area index (leaf area per unit ground area) [13].

The approach is described by the classification process given on Figure 1. Processes are given in closed boxes while data appears in open ones. The system takes as input a colour image of a section of a field and it outputs the number of pixels respectively associated to weeds or crop plants. The whole operation is divided into three steps: segmentation, feature extraction, and classification. Each part will be described in the following sections. The image acquisition process is currently not fully integrated to our system, it is therefore not represented in Figure 1. However, since the images are essential to our system, the next section will described how they are acquired.

2.1 Acquisition

All images are taken using a camera mounted on a capturing/carrying custom mobile platform as presented in Panneton and Brouillard [16]. The self-propelled platform consists of a pyramid-shaped aluminum chassis covered with opaque shading material and equipped with two high power flash lights combined to a custom reflector to ensure uniform illumination and minimum shadow for all pictures. An operator steers the unit and monitors its speed. The camera is linked to a distance sensor located on one wheel so that it takes a picture every three meters. The device is built to hold the camera at its apex (about 2.5 m above the ground) with its optical axis perpendicular to the ground, covering a rectangular area of 3 m by 2 m with each picture (4 rows of corn at 75 cm spacing). A Nikon D100 RGB digital

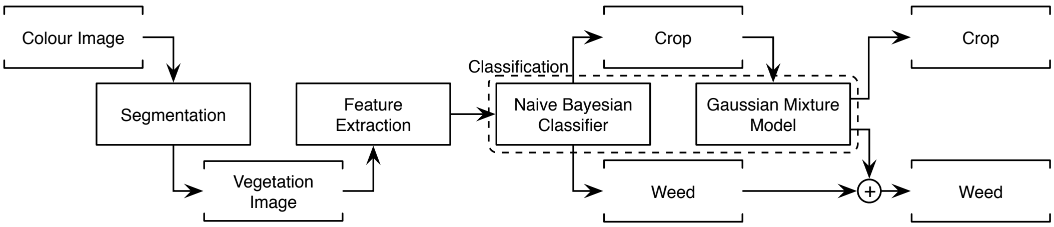

Fig. 1 Complete classification process from field images to weed/crop

camera is used. Around 6 million pixels (3034 × 2024) covers the 3 m by 2 m area giving a resolution of approximately 1 px/mm^2.

2.2 Segmentation

The colour image is segmented to split the vegetation from the ground. The algorithm, based on the principal component analysis, extracts the component associated with the vegetation in the image. The Otsu method [14] is then used to find a threshold that will discriminate between vegetation and soil based on that component score. The image is segmented by masking all the soil pixels, leaving the vegetation pixels to their original colour. Because the current work focuses on the feature extraction and the classification processes, the segmentation procedure is not fully presented and can be found in Panneton [15] and Panneton et al [17]. In fact, the segmentation method could be replaced by an alternate one as long as it can efficiently mask the ground leaving only vegetation pixels in the output image.

2.3 Feature Extraction

Once the image are segmented, the important features can be extracted. In the proposed approach, the feature extraction process is influenced by the spatial location of the plants. Thus the first step consists in identifying the crop rows in the image. This is done using the Hough transform that finds the straight lines made by the densest vegetation. Interestingly, vegetation lines are found by such a procedure even if alignment is not perfect with respect to the field geometry or image acquisition process. These straight lines are more likely to be crop plants because they are seeded in rows using a mechanical planter. Row borders are found by analyzing the plants that cross the row centre lines. Then, based on our primary assumption, every plant outside the rows is considered to be weed. The data provided by these plants outside the rows is used to compute the probability density functions (PDFs) of morphological characteristics of weeds. Considering that the plants inside the rows are a mixture of weeds and crop plants, the inside row data is used to infer the PDFs of morphological characteristics of crop based on the weed models. This novel feature inference method is described in details in Section 3 along with the row identification process.

2.4 Classification

Now that the morphological characteristics PDFs have been found for weeds and crop plants, a naive bayesian classifier based on the extracted features is used to differentiate the two classes of plants that are in the rows. It outputs two sets of labelled blobs. The set that has been classified as crop plants is sent into a second classifier in order to identify leftover weeds using a gaussian mixture model. Weeds found by the gaussian mixture model are added to the ones previously found. Finally, the weed cover field-section percentage can be estimated using the preceding classification. The complete classification process is described in details in Section 4.

3 Feature Extraction

As aforementioned, the proposed approach for feature extraction is based on the spatial position of plants, i.e. the presence of crop rows. Therefore, the first step is detecting the rows. Then, features from plants inside and outside the rows are extracted. This allows modelling the PDFs of weed features, and inferring the PDFs associated to features of crop. The whole feature extraction process is detailed in the following sections.

3.1 Row Detection

Row information is crucial for gathering weed properties. This information will be used later to make assumptions on the nature of plants according to their relative position to the rows. In the proposed method, the rows are detected using the Hough transform on segmented and binarized images of the field. This technique has already been used in Asif et al [3], Jones et al

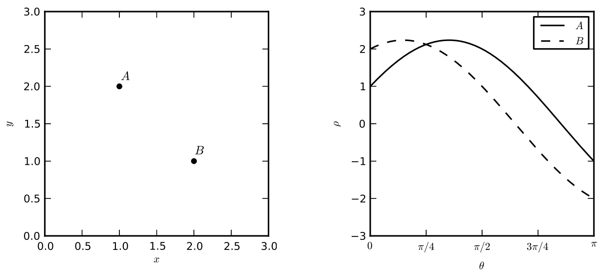

(a) Points in the x, y space.

(b) Parameters of every line going through A or B in the ρ, θ space.

Fig. 2 Hough transform of an image containing two points. The intersection of the two lines in (b) represents the parameters of the *single line* in (a) going through both points A and B

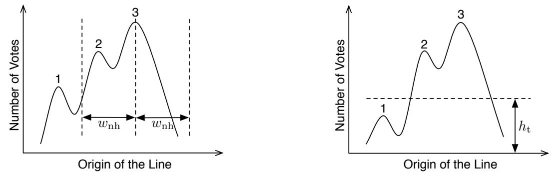

(a) Neighbouring size, peak number 2 is discarded because of its proximity to peak number 3.

(b) Threshold, peak number 1 is too low to be considered.

Fig. 3 Rules to find the Hough peaks

[11], Slaughter et al [20], Tellaeche et al [22], and Åstrand and Baerveldt [25]. The current section explains how the Hough transform is used for detecting the row position, followed by the computation of row boundaries and details about the assumptions made on plants located outside the rows.

3.1.1 Row Position

Introduced by Hough [10], the Hough transform is generally used in computer vision to find patterns in an image via a voting procedure. In this case, it is used to find straight lines defined by aligned crop plants in a segmented image of a field section. When detecting lines in an image that contains two points $A = (x_A, y_A)$ and $B = (x_B, y_B)$ (Figure 2(a)), the transform first finds every possible line that passes by A, i.e. every pair (ρ, θ) that satisfies the line equation $\rho = x_A \cos\theta + y_A \sin\theta$, where $\rho \in \mathbb{R}$ is the distance between the line and the origin and $\theta \in [0, \ldots, \pi]$ is the angle of the vector from the origin to the closest point of the line. The pairs giving those lines are located on the continuous line of Figure 2(b). Then, it finds every line that goes through B, satisfying the equation $\rho = x_B \cos\theta + y_B \sin\theta$ (dashed line of Figure 2(b)). Finally, every pair (ρ, θ) present on the two lines of Figure 2(b) is given a vote. A single pair has two votes, it represents the most dominant straight line present in the image, i.e. the one that crosses both points, A and B, in the image space. Generalizing this idea to more natural binary images allows to locate the points (pixels) that are relatively align on straight lines even if they are discontinuous. The most dominant lines are represented by the combination of parameters ρ and θ with the most votes.

In our application, the parameters with the most votes represent the lines with denser vegetation in the image. We only consider the pairs (ρ, θ) with $\theta = 0$ and $\rho = [0, 0.5, \ldots, \sqrt{\text{image width}^2 + \text{image height}^2}]$. From this group of parameters, a subset representing row centres are selected following three simple rules:

1. each maximum must be separated by a neighbouring size of w_{nh};
2. a maximum must be higher than a threshold h_{t};
3. a maximum of n_{p} greater peaks are kept.

Rules one and two are illustrated in Figures 3(a) and 3(b) on a one parameter Hough transform. Using the first rule, maximum number 2 is eliminated because it is too close to the higher maximum number 3. Using the second rule, maximum number 1 is not considered because it is under the threshold. The selection of w_{nh} and h_{t} is based on the input image, while n_{p} has to be set according to the maximum number of crop rows that can be found in an image, as shown by Equations 1 and 2.

$$w_{\mathrm{nh}} = \frac{\text{image width}}{2 n_{\mathrm{p}}} \tag{1}$$

$$h_{\mathrm{t}} = \frac{\max(\text{votes})}{4} \tag{2}$$

3.1.2 Row Boudaries

The boundaries of a row are computed using the bounding box of each blob that crosses the row centre line c. The centre lines are given by the Hough transform. The left width w_{l} and the right width w_{r} are computed using the average distances from the row centre line to the left and right ends of each bounding box.

A blob is considered to be part of a row if at least one pixel from that blob is contained in $[c - w_{\mathrm{l}}, c + w_{\mathrm{r}}]$. Figure 4 presents an example of blobs that are inside and outside a row.

3.1.3 Row Relative Assumptions

Proper detection of rows is important because it allows the formulation of hypotheses on the blobs according

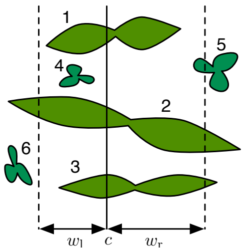

Fig. 4 Blobs in and out a row. Blobs 1, 2 and 3 are used to determine w_l and w_r. Every blob is inside the row except for blob 6

to their position with respect to the rows. Indeed, the proposed method freely identifies a plant outside a row as a weed. Rows have to be properly located; if a row is not detected, its entire content will be classified as weeds and will bias the weed PDF towards the crop PDF. Further classification presented in Section 4 will be applied only to blobs located inside the rows.

3.2 Modelling Probability Density Functions

PDFs are modelled for some morphological characteristics of the two groups of plants based on their position on the image: weeds only between the rows; and crop plants and weeds combined inside the rows. The PDFs must be estimated since the process that generated the characteristics of each plant is unknown. The estimation of the PDFs is made using kernel density estimation, also known as Parzen-Rosenblatt window [18, 19].

From a set of independent and identically distributed random variables, the form of the function $p(x)$ that generated n samples x_i can be approximated by

$$\hat{p}(x) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right), \tag{3}$$

where $K(\cdot)$ is a kernel function and h is a smoothing parameter. A gaussian kernel is used for modelling the morphological features of the plants. Since measured characteristics are always positive, we apply a constraint on the probabilities by applying a transformation. The density is estimated with the data scaled logarithmically. When computing the probability for a certain value, the result is transformed back to the original scale. h is set in function of the gaussian optimal smoothing rule of thumb described in detail in Bowman and Azzalini [4] along with the log scale transformation for bounded data. The estimator $\hat{p}(x)$ is used to evaluate PDFs both inside and outside the rows given respectively by $\hat{p}(x|\mathcal{I})$ and $\hat{p}(x|\mathcal{O})$.

3.3 Inferring Probabilities

The goal of the classification is to distinguish crops from weeds and vice versa. Unfortunately, the information about crop plants cannot be inferred directly from the images. However, weed cover characteristics can be deduced as they are the only plants outside the row boundaries. Since weeds located outside the row boundaries are probably similar to weeds located inside the rows, it is possible to subtract the PDF estimated from outside the rows to the one estimated inside the rows to assess the crop PDF.

Two things are known in terms of conditional probability: the probability of occurrence of a characteristic x knowing the plant is a weed $p(x|\mathcal{W})$ and the probability of occurrence of a characteristic x knowing the plant is a weed or a crop $p(x|\mathcal{C} \cup \mathcal{W})$. This allows us to calculate the probability $p(x|\mathcal{C})$ of a characteristic x knowing the plant is a cultivated one, with the mixture model equation

$$p(x) = \sum_{i=1}^{k} p(\mathcal{G}_i) p(x|\mathcal{G}_i), \tag{4}$$

where $p(x)$ is the probability of occurrence of x, $p(\mathcal{G}_i)$ is the probability of occurrence of class $\mathcal{G}_i$ and $p(x|\mathcal{G}_i)$ the probability of occurrence of x knowing it is part of class $\mathcal{G}_i$. Applying Equation 4 to the two classes problem mentioned above gives

$$p(x|\mathcal{C} \cup \mathcal{W}) = p(\mathcal{C})p(x|\mathcal{C}) + p(\mathcal{W})p(x|\mathcal{W}). \tag{5}$$

Isolating $p(x|\mathcal{C})$ in Equation 5, we find

$$p(x|\mathcal{C}) = \frac{p(x|\mathcal{C} \cup \mathcal{W}) - p(\mathcal{W})p(x|\mathcal{W})}{p(\mathcal{C})}. \tag{6}$$

In this last formulation the priors $p(\mathcal{C})$ and $p(\mathcal{W})$ are still missing. They can be estimated using the pixel density of weed per area unit outside the rows ρ_o and the pixel density of the whole vegetation inside the rows ρ_i. Assuming that the weed pixel densities outside and inside the rows are similar, ρ_i should be equal to ρ_o plus the crop density. Weed and crop pixel densities inside a row are then given by

$$\rho_\mathcal{C} = \rho_\mathrm{i} - \rho_\mathrm{o}, \tag{7}$$

$$\rho_\mathcal{W} = \rho_\mathrm{o}. \tag{8}$$

The priors are then normalized to

$$\hat{p}(\mathcal{C}) = \frac{\rho_\mathcal{C}}{\rho_\mathrm{i}}, \tag{9}$$

$$\hat{p}(\mathcal{W}) = \frac{\rho_\mathcal{W}}{\rho_\mathrm{i}}. \tag{10}$$

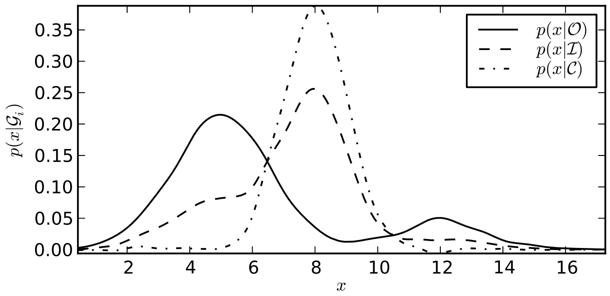

Fig. 5 Estimation of the PDFs for crops $p(x|\mathcal{C})$ based on the data outside $p(x|\mathcal{O})$ and inside $p(x|\mathcal{I})$ the rows. It can be seen as subtracting (according to the priors) the density $p(x|\mathcal{O})$ from $p(x|\mathcal{I})$ resulting in the density $p(x|\mathcal{C})$ using Equation 6.

If this estimation fails, i.e. pixel density outside the rows is superior to pixel density inside the rows, an equiprobable estimator is assumed: $\hat{p}(\mathcal{C}) = \hat{p}(\mathcal{W}) = 0.5$.

Figure 5 presents how the inferring process is used to estimate the PDF for crop plants $p(x|\mathcal{C})$ and weeds $p(x|\mathcal{W})$. First, the PDFs for the vegetation inside $p(x|\mathcal{I})$ and outside $p(x|\mathcal{O})$ the rows are modelled by a Parzen-Rosenblatt window using the respective data for a characteristic x. The rows contain both crops and weeds, while the inter-rows contain only weeds. Thus, $p(x|\mathcal{I})$ and $p(x|\mathcal{O})$ are respectively equal to $p(x|\mathcal{C} \cup \mathcal{W})$ and $p(x|\mathcal{W})$. Then, using the priors given by Equations 9 and 10 in the inference Equation 6, we find the crops hidden PDF.

4 Classification

The classification of vegetation present in the rows is based on the PDFs inferred from the morphological features of each blob. Two probabilities (weed and crop) by feature are computed using the Bayes' theorem and the final classification is done by choosing the higher of those two probabilities. In order to refine the classification, a technique based on gaussian mixture models is employed to filter out from the group of plants classified as crop, the weeds that were wrongly classified using the naive bayesian approach. The newly identified weeds are added to the already classified ones. The present section describes those two classification methods.

4.1 Naive Bayesian Classification

Bayesian classification is done by using the preceding inferred probabilities taken from a single image and by applying them to the unlabelled data in order to establish to which class it belongs. The Bayes' theorem gives the probability $p(\mathcal{G}_i|x)$ of belonging to class $\mathcal{G}_i$ knowing the observed variable x, the prior probability $p(\mathcal{G}_i)$, and the marginal probability $p(x)$ of observing x. It is defined as

$$p(\mathcal{G}_i|x) = \frac{p(\mathcal{G}_i)p(x|\mathcal{G}_i)}{p(x)}. \tag{11}$$

From the Bayes' theorem, a naive bayesian classifier will attribute a label based on the higher probability given by the application of Equation 11 for $i = 1 \dots k$. Note that this naive bayesian classifier can easily be extended to the multiple feature case.

Applied to the current problem and using the estimators described in Section 3.3, the classifier becomes

$$p(\mathcal{C}|x) = \frac{\hat{p}(\mathcal{C})\hat{p}(x|\mathcal{C})}{p(x)} \qquad \text{and} \tag{12}$$

$$p(\mathcal{W}|x) = \frac{\hat{p}(\mathcal{W})\hat{p}(x|\mathcal{W})}{p(x)}, \tag{13}$$

with the class given by $\arg\max_{\mathcal{G}_i \in \{\mathcal{C},\mathcal{W}\}} p(\mathcal{G}_i|\mathbf{x})$, where $\hat{p}(x|\mathcal{C})$ is given by Equation 6 with $p(\mathcal{W}) = \hat{p}(\mathcal{W})$, $p(\mathcal{C}) = \hat{p}(\mathcal{C})$, $p(x|\mathcal{W}) = \hat{p}(x|\mathcal{O})$, and $p(x|\mathcal{C} \cup \mathcal{W}) = \hat{p}(x|\mathcal{I})$. For more details on bayesian classification the reader is referred to Alpaydin [2].

4.2 Gaussian Mixture Model

Unsupervised learning is generally used to classify or identify data generated by different, a priori unknown, phenomena. After applying the naive bayesian classification, it is highly probable that some weeds were classified as crops. The goal of the gaussian mixture model (GMM) is to refine this classification by excluding, from the classified crop data set, the data that does not fit the crop general distribution.

The GMM identifies, in the provided data set $\mathcal{D}$, k underlying gaussian distributions that are most likely to reproduce $\mathcal{D}$. Equation 4 gives the probability density function given k gaussian distributions $\mathcal{G}_i$. The fitting of the distributions – finding the appropriate means and standard deviations – is usually done with the expectation-maximization (EM) algorithm [6]. This algorithm requires the number of inherent distributions to be predefined. Figueiredo and Jain [7] propose an algorithm that identifies the number of distributions that best fit the data set while also precluding the possibility for the EM algorithm to converge to a boundary configuration.

The algorithm from Figueiredo and Jain [7] is used in the current method to identify the underlying distributions of the crop blobs area. Figure 6 presents the distribution of the area of each blob classified as crop, after using the bayesian classification of Section 4, for

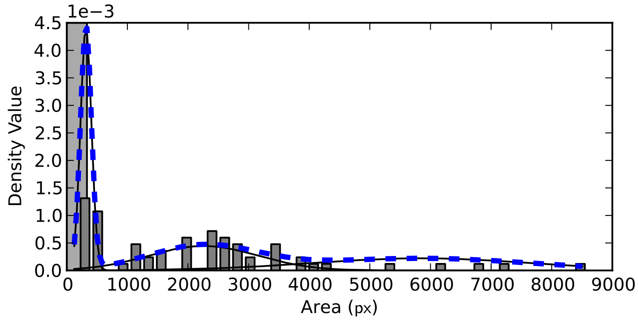

Fig. 6 Probability density function of the areas classed as crop using three gaussians. The thin lines are each gaussian, the dashed thick line is the mixture model and the bars are the underlying histogram of occurrence. The shaded rectangle represents the part of the crops that are relabelled as weed.

a single image. The first gaussian, with mean about 334 px, is most likely to be produced by a different phenomenon than the two next that are more similar. Based on the assumption that weed blobs are generally smaller than crop blobs, any group of pixels having an area smaller than the smallest estimated mean is relabelled as a weed. In Figure 6, all blobs with an area falling in the shaded rectangle (between 0 and 334 px) are relabelled into the weed class.

We found experimentally that classifying every blob belonging to the first gaussian distribution significantly decreased the percentage of well classified crop blobs while not increasing the percentage of well classified weed blobs. For example, in Figure 6, this latter classification would have reclassified every blob with an area up to 585 px into the weed class. However, the set of blobs with an area between 334 and 585 px contains a majority of crop plants. Being less aggressive in the post-bayesian classification, using the smaller mean as criterion, allows us to maintain a very high percentage of properly classified crop blobs while increasing substantially the percentage of accurately classified weed blobs.

5 Experiments

This section presents our method for classifying crops and weeds applied over a set of 149 images taken from two different crops: corn (C) and soybean (S), at the second to fourth leaf stage of corn and at the second to third trifoliate leaf stage of soybean. Each crop grows on two different sites located in the province of Québec (Canada): Acadie (A) and Beaumont (B). In the images, three different levels of infestation were identified: low (L), medium (M), and high (H). Examples of each infestation level for corn and soybean fields are presented in Figures 7 and 8. In the current section, images are identified using a three-letter acronym representing respectively the site, the crop, and the level of infestation. For example, the acronym ASL is used for site Acadie (A) soybean fields (S) with a low weed infestation (L). Following are presented the experimental procedure, the measures used for comparing the methods and the results of the proposed classifier for both classifying the data and estimating the weed cover in parts of the fields.

5.1 Experimental Procedure

The input images are segmented providing vegetation only images. The resulting segmented images are then filtered using an opening morphological operation with a disc of 1 px radius[2] in order to remove most of the segmentation noise. In fact, 55% of the 872 000 plant blobs found in the 149 images have an area inferior to 4 px, but they represent only 1.1% of the total blob surface. Clearing these blobs prevents their small size from disturbing the statistics used by the bayesian classifier without risking the loss of important data.

Next, row detection, using the Hough transform, is applied to find the position and bounds of each row. The value of n_{p} is set as the number of visible rows in each image, inducing at the same time w_{nh}. A 4-connected scheme is used to identify blobs in the images, in order to limit overlapping between groups of plants that would be considered a single blob under a 8-connected scheme (Figures 9(e) and 9(f)). PDFs of these features for weed are then computed using the plants outside of rows. The PDFs associated to these features for the crop are then inferred using the data inside the rows.

The following morphological properties are the features that were tested:

Area – the number of pixels in a blob (Figure 9(a));
Compactness – the shape descriptor being independent of any linear transformation and defined as $\frac{P^2}{A}$, using the perimeter P and the area A (Figures 9(b) and 9(c));
Major Axis – the value of the greater axis of the ellipse with the same normalized second moment than the blob (Figure 9(d)).

The probabilities computed for weed and crop are both used as the input of the naive bayesian classifier which outputs the class for each plant. All blobs classified as crop are then passed to the GMM to refine their classification. At this step, only the area is used as a feature. Every blob with an area smaller than the smallest mean found by the gaussian mixture algorithm

[2] A disc with a radius of 1 px is a cross with a diameter of 3 px.

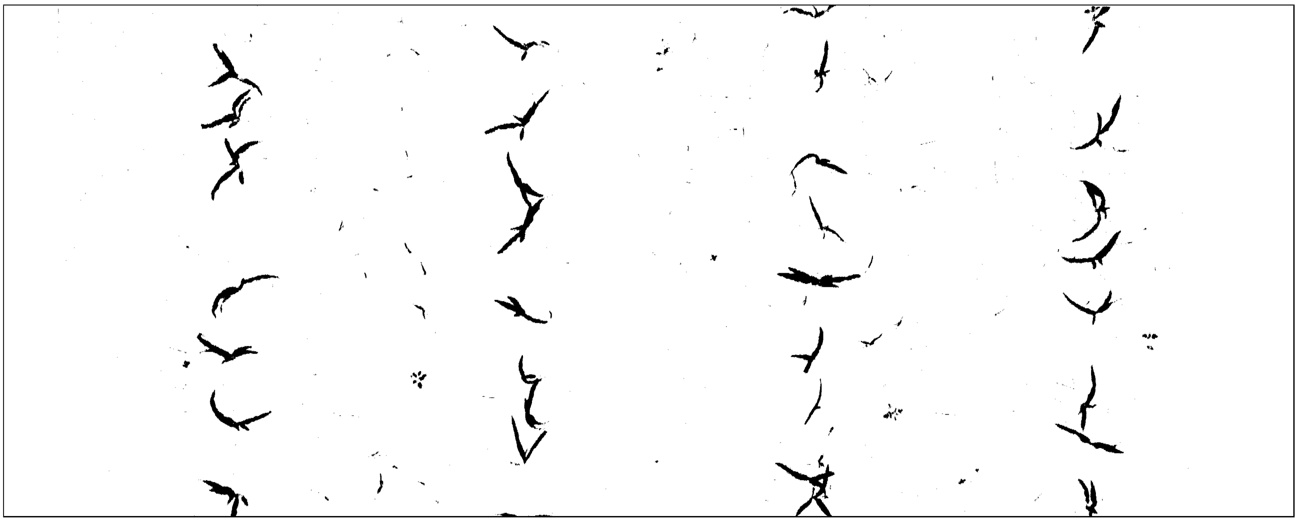

(a) Low infestation

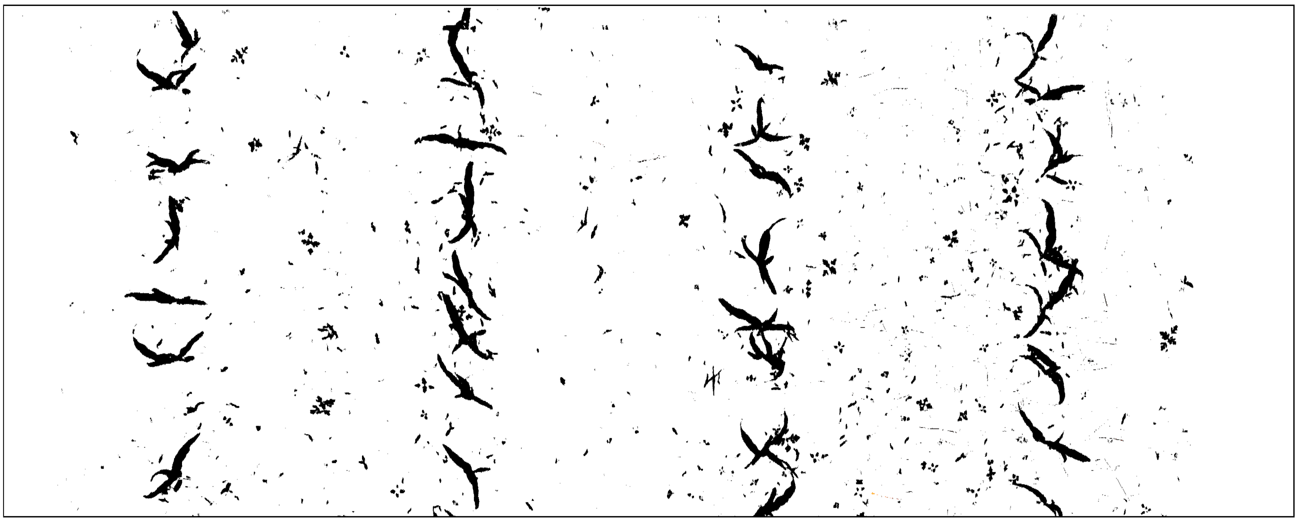

(b) Medium infestation

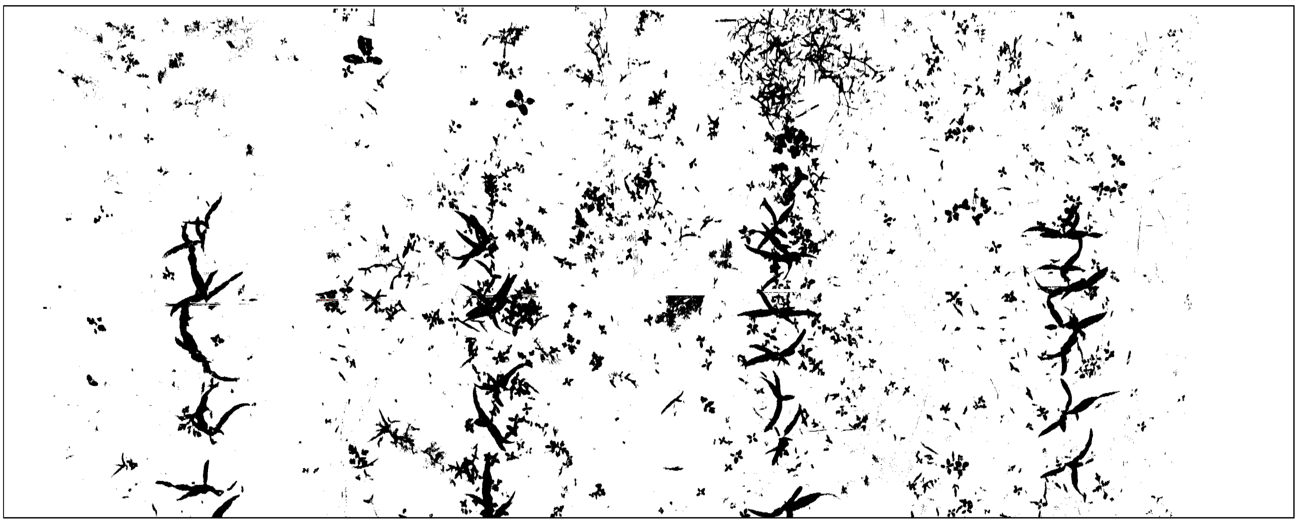

(c) High infestation

Fig. 7 Infestation levels in corn fields

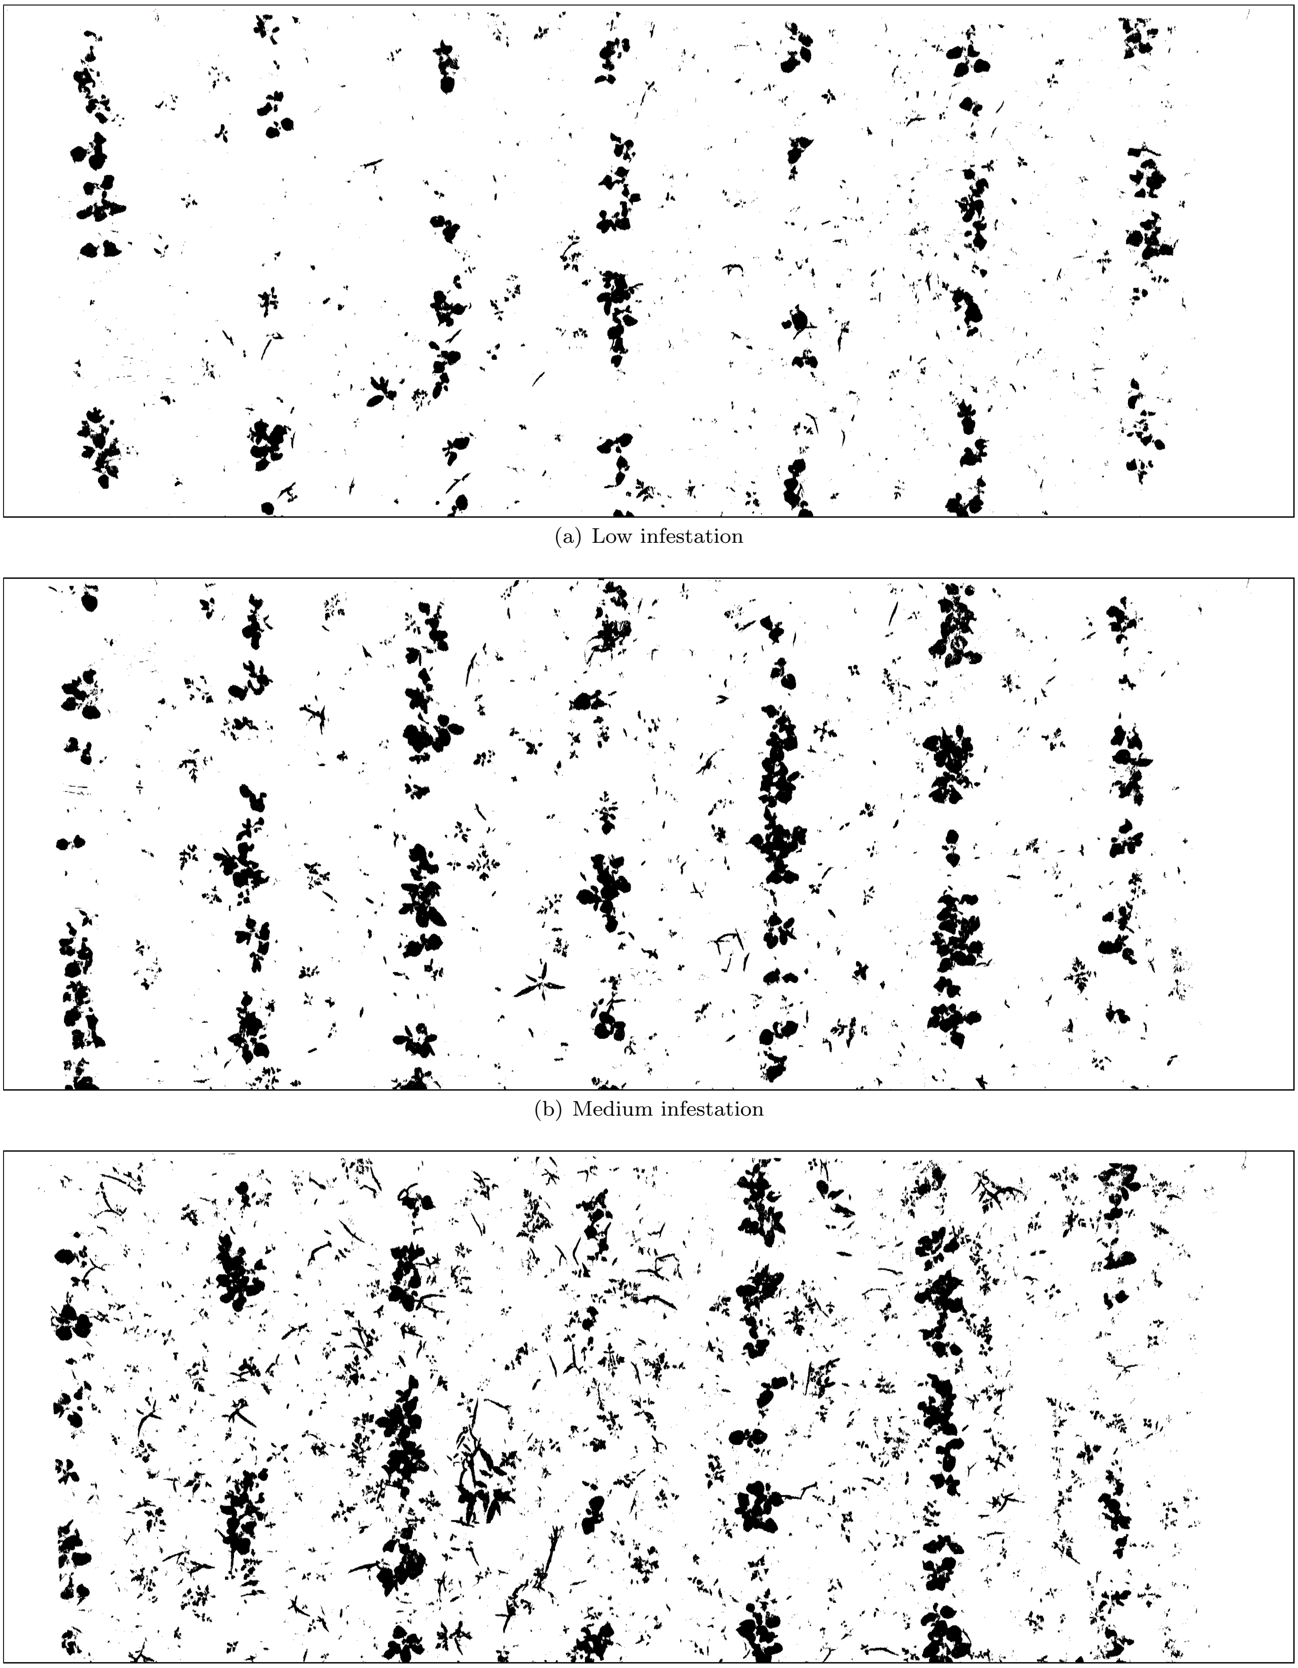

(a) Low infestation

(b) Medium infestation

(c) High infestation

Fig. 8 Infestation levels in soy fields

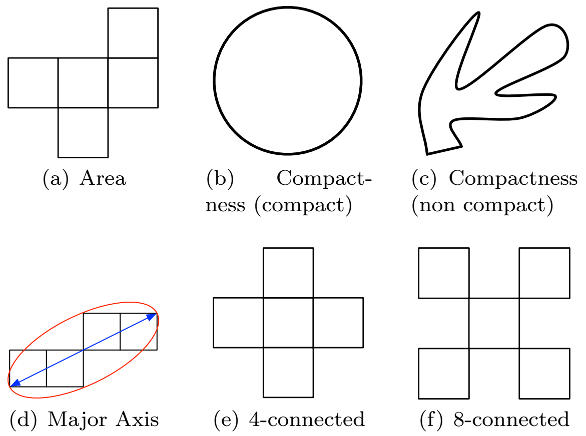

(a) Area (b) Compactness (compact) (c) Compactness (non compact)

(d) Major Axis (e) 4-connected (f) 8-connected

Fig. 9 Representation of the different morphological properties used to distinguish between crops and weeds, (a) represents a blob having an area of 5 pixels, (b) and (c) are examples of compact and non compact forms, (d) is a representation of the major axis, and (e) and (f) are blobs under a 4-connected and 8-connected scheme respectively.

is relabelled as weed. The final output consists in two sets of labelled blobs, one for crops and one for weeds.

5.2 Measures

In the remaining of this section the following measures are used, with the cardinality $|\mathcal{A}|$ giving the number of pixels contained in a set $\mathcal{A}$:

- $|\mathcal{I}|$ is the total number of pixels in the image;
- $|\mathcal{C}^*|$ is the actual number of crop pixels, with $\mathcal{C}^*$ the set of perfectly classified crop pixels;
- $|\mathcal{W}^*|$ is the actual number of weed pixels, with $\mathcal{W}^*$ the set of perfectly classified weed pixels;
- $|\mathcal{C} \cap \mathcal{C}^*|$ is the number of pixels correctly classified as crops, with $\mathcal{C}$ the set of crop pixels classified by the proposed method;
- $|(\mathcal{V} \backslash \mathcal{C}) \cap \mathcal{W}^*|$ is the number of pixels correctly classified as weeds, with $\mathcal{V} \backslash \mathcal{C}$ the set of all vegetation pixels excluding those contained in $\mathcal{C}$.

In each image, weed pixels have been identified by a weed scientist. These were used directly to determine $\mathcal{C}^*$ and $\mathcal{W}^*$. The use of $\mathcal{V} \backslash \mathcal{C}$ instead of $\mathcal{W}$ allows us to reintroduce, as weed, the small blobs that have been removed in the filtering phase of the method. The area of these blobs, although small for the learning phase, remain important when estimating the weed cover.

The classification performance for crop κ_{c} is given by

$$\kappa_{\mathrm{c}} = \frac{|\mathcal{C} \cap \mathcal{C}^*|}{|\mathcal{C}^*|}. \tag{14}$$

Corollary, the classification performance for weed κ_{w} corresponds to

$$\kappa_{\mathrm{w}} = \frac{|(\mathcal{V} \backslash \mathcal{C}) \cap \mathcal{W}^*|}{|\mathcal{W}^*|}. \tag{15}$$

Also, the global classification performance κ is given by

$$\kappa = \frac{|\mathcal{C} \cap \mathcal{C}^*| + |(\mathcal{V} \backslash \mathcal{C}) \cap \mathcal{W}^*|}{|\mathcal{C}^*| + |\mathcal{W}^*|}. \tag{16}$$

Finally, the relative weed cover estimation χ_{w} is given by

$$\chi_{\mathrm{w}} = \frac{|(\mathcal{V} \backslash \mathcal{C}) \cap \mathcal{W}^*|}{|\mathcal{I}|}. \tag{17}$$

Results for a particular method are denoted using the exponent b for the naive bayesian classification and gm for the GMM. For example, $\kappa^{\mathrm{b,gm}}$ denotes the global classification performance for the bayesian classification followed by the GMM technique. In order to strictly demonstrate the classification efficiency, we worked on the subset of images for which row positionning was successful. This subset contains 129 of the 149 images. The images that have been discarded had either too many weeds or too few crop plants, making it impossible for the Hough method to identify the correct peaks. This is not a concern since we present a classification method that depends on prior knowledge provided by the row identification method.

5.3 Classification Results

Table 1 gives the global classification average performance measured with and without the GMM, for each morphological properties used in the naive bayesian classifier. The combination of properties are not considered since this classifier requires independent properties and these are not. Knowing that the performance κ^{gm} when using only the GMM is 88.9%, it is clear that combining the bayesian classification and the gaussian mixture improves the classification rate. Moreover, it can be seen that the best performance, 90.8%, is obtained using either the area or the major axis. However, the smallest standard deviation, 4.4%, is obtained when using the area. It should be noted that these results are specific to the current experimental problem and different properties could perform better in another context. A Student's t-test reveals that the hypotheses $\kappa^{\mathrm{b,gm}} > \kappa^{\mathrm{gm}}$ and $\kappa^{\mathrm{b,gm}} > \kappa^{\mathrm{b}}$ have a significance level higher than 99.95% each when using the area. Using both method successively is thus significantly better than using them separately. In consequence, the rest of this paper will discuss this combined method.

Table 1 Global classification average and standard deviation performance using different properties for the bayesian classifier. Classification percentage κ^{gm} is 88.9% with a std dev. of 5.0.

Property	κ^{b} (std dev.)	$\kappa^{\mathrm{b,gm}}$ (std dev.)
Area	90.5% (4.5)	90.8% (4.4)
Compactness	89.7% (5.2)	89.9% (5.1)
Major Axis	90.5% (4.6)	90.8% (4.5)

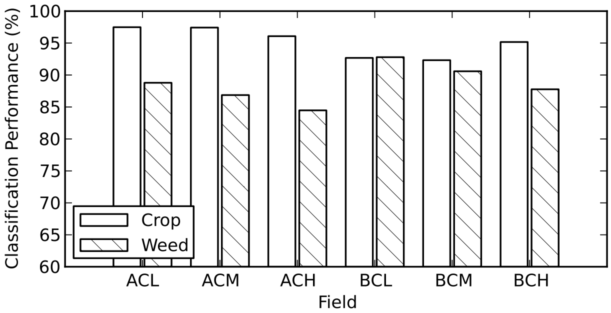

(a) Corn

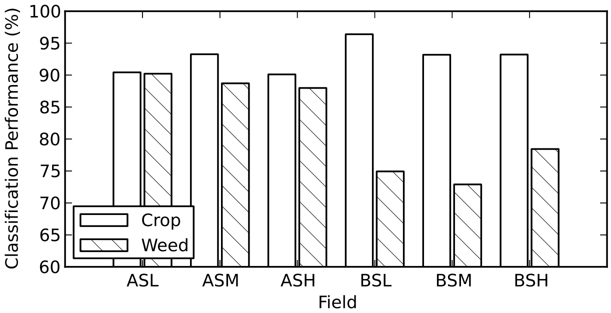

(b) Soy

Fig. 10 Average classification performances $\kappa_{\mathrm{c}}^{\mathrm{b,gm}}$ and $\kappa_{\mathrm{w}}^{\mathrm{b,gm}}$ for different crops and levels of infestation using the area as only characteristic

Figure 10 shows the average percentage of good classification for crop and weed at different infestation levels. Each subfigure represents for each field, one of the two crops, corn or soy. The results were obtained using the area as the only feature. One could notice that the crops are generally better identified than the weeds. A possible cause for this is that cultivated plants are more uniform in size since there is only one type of crop in a field and the crop was all planted at once on the same day (i.e. same age); contrary to the multiple weed species that each have their own emergence time and pattern. In other words, the fuzziness induced by the different kinds of weeds does not allow the classifier to fit a narrow distribution around their characteristics. The widespread distributions create uncertainty that lower the probability of a blob to be classified as a weed, favouring the crop class.

Figure 12 shows the resulting classification for a single row of corn with medium weed infestation. We observe that the classification is generally good and most

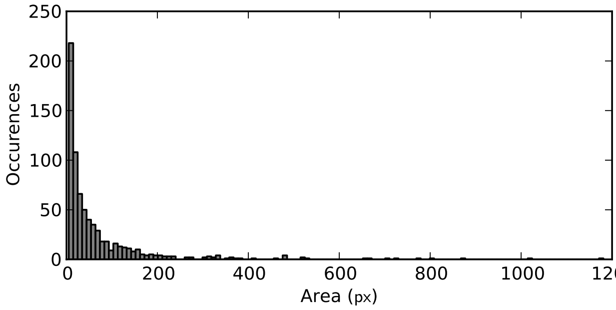

Fig. 11 Histogram of the inter-row weeds area

of the weeds have been identified correctly. However, it is possible to identify three types of error identified as A, B and C. The first type, A, corresponds to a very small part of the crops that have either been separated of the main plant during the segmentation process or misclassified by the specialist. In either case, the area of those blobs is very similar to other weeds, therefore letting us strongly believe that any other classification method would take a similar decision. The second type of error, B, occurs when a weed have been mistaken for a crop. Knowing that blob B has an area of 872 px, Figure 11 shows that there are very few weeds between the rows that have a similar area. Since few of this weed type reside outside of the rows, the inference mechanism can explain this type of error. The third type of error, C, is caused by a weed that is merged with a crop. Since decisions in our method are based on blobs, this sort of error cannot be avoided.

Together, the three types of error create the following confusion. In the corn field images, 12% of the weeds were classified as crop plants and 6% of crop plants were classified as weeds. In the soybean field images, these are 18% and 6% respectively. Our results cannot be compared directly to any other published work. However, Hemming and Rath [9] presented a technique close to ours where images are taken much closer to plants in cabbage and carrot cultures. They achieved a misclassification error of about 11% and 12% for cabbage culture, while about 34% and 31% for carrot culture. This shows that our technique compares readily to previous work in term of classification rate for differentiating weeds from crops.

5.4 Weed Cover Estimation

The plant classification method allows weed cover estimation. A common unpublished technique is to esti-

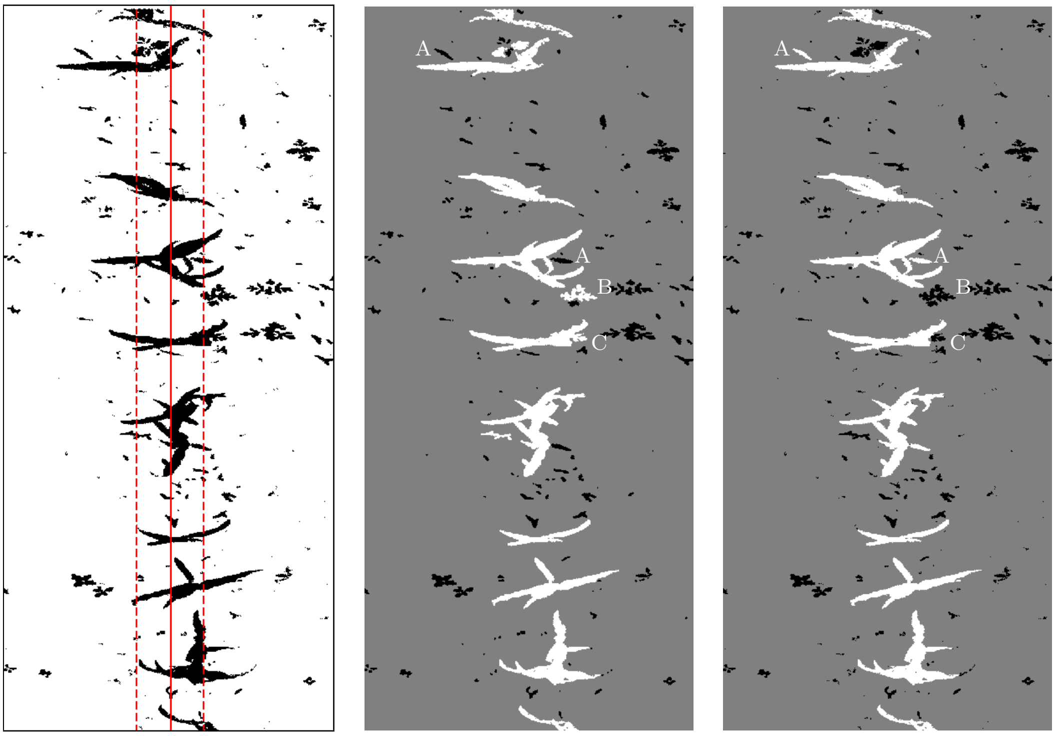

(a) Segmented Image (b) Classification (c) Truth

Fig. 12 Output of the classification for a single row in a corn field, (a) the segmented image with centre of row marked by a full line and boundaries marked by dashed lines, (b) the result of the classification and (c) the truth image, where white blobs are the corn plants and black blobs are weeds.

mate the position and bounds of the rows[3] and compute the number of pixels between each row, assuming that every non-background pixel belongs to a weed. The coverage is estimated by the ratio between the number of plant pixels in inter-rows and the inter-row total area. This method will be referred to as the proportional method (prop.).

Figure 13 presents the absolute error in weed cover estimation using the bayesian classification in comparison to the proportional method and the true coverage. As we can see, on each field the estimation using the bayesian classification is closer to the true value than the proportional method. In fact, the coverage estimation error is on average 3.7 times less with our method than with the proportional method. We note that the quantity of weed in an image is always underestimated using the proportional method while it is underestimated in 69% of the images using the bayesian classification. However, since our estimation is quite closer to the real value, we can conclude that our method offers a good compromise with minimal error when compared to the common trivial approach.

[3] With this method, the boudaries of each row are computed using the absolute minimum and maximum of the bounding box of the blobs that cross the centre line.

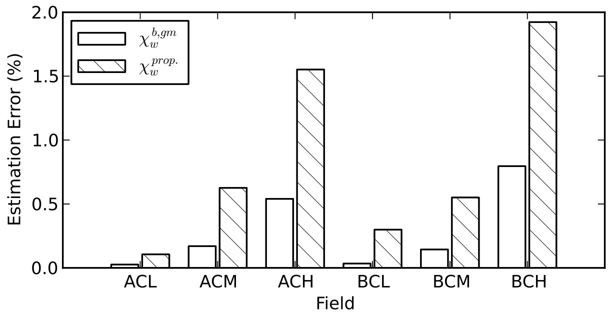

(a) Corn

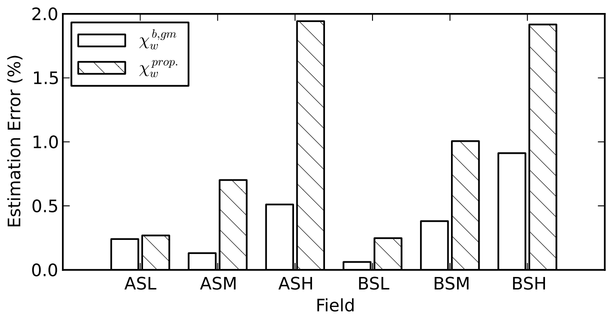

(b) Soy

Fig. 13 Average estimation errors in percentage $\chi_{\mathrm{w}}^{\mathrm{b,gm}}$ and $\chi_{\mathrm{w}}^{\mathrm{prop.}}$ for different crops and levels of infestation using the area as only characteristic for the bayesian classifier

The weed cover estimation measure is subject to bias and should be taken with care. The coverage could be computed perfectly although the classification error could not be negligible. This is because the area of weed classified as crop could be compensated by the area of crop classified as weed. We think that this bias may occur in our experiments but the probability that it is corrupting our results is very small given our classification results.

6 Conclusion

In this paper we addressed the problem of evaluating the degree of infestation in a section of a field represented by a single image containing one or more crop rows. We presented an original technique for extracting data in order to discriminate weeds from crops based solely on the information available in the image. We have shown that this technique provides a high classification rate on two distinct vegetation types and on various levels of infestation. We have also illustrated that the classification leads to a very good estimation of the weed coverage on field sections. Both of these results are steps toward smart herbicide application on crop fields, thus helping to reduce the overall amount of sprayed herbicide.

The presented technique has two major advantages: no prior training nor prior knowledge on crop or weed species are necessary. These constitute a requirement in a real-time herbicide application context. Such a system could be realized with an acquisition system, namely a set of cheap cameras, placed in front of the tractor and the herbicide applicator placed behind. The information would serve to modulate the herbicide dose passing through the nozzles when they go over the scanned area. The number of sensors could be increased to cover the entire width of a boom (6 m or more) and the number of nozzles could be increased (each nozzle covering a smaller area) for greater precision in the application of the herbicide. Tractors speed is reasonably low when applying herbicide (around $10\,\mathrm{km/h} \approx 3\,\mathrm{m/s}$), which gives our method approximately 1 s to acquire and process an image, and activate the required spraying nozzles, if the system is 3 m long. Parallelism can also be exploited as each row is treated independently, reducing the time required to process an image.

The proposed method also has many other advantages. Its intuitive probabilist reasoning makes it easy to understand and adaptable to any sort of cultivation. It is easily extensible, providing new characteristics to the process is painless and can bring novel information to the classifier. A possible venue for new type of information could imply using stereoscopic vision or time of flight sensors to retrieve the height of each pixel. It would also be possible to use ultra-violet/infrared sensors to capture the non-visible signature of the plants. Finally, the bayesian classifier also accounts for an acceptability parameter that could be tuned to favour the classification of weed or crop. One that prefers missing wort over missing weed can adjust this parameter to fit his needs.

The present work can be extended in several ways. First, it would be interesting to evaluate the performance on images of entire rows instead of sections. The herbicide application setup could be rearranged so that it would not interfere with the real-time usage. Second, cumulative knowledge in a single field could be implemented so that the system could reuse knowledge aquired earlier when processing images on the same field. This could improve the quality of the inference. Third, in order to fully benefit from the information brought by multiple characteristics, the bayesian classifier could be modified to account for the interdependence between each selected property. Fourth, classifying weeds in a single class stretch out the weed probability distribution. This could be avoided by identifying, in an unsupervised manner, multiple weed types, allowing to tighten the distribution around those types. It could also be considered that weed grows only between crops, thus reducing the available area and their relative density in crop rows. Finally, a multiple classifiers voting technique could help boosting the classification success rate.

Acknowledgements The authors wish to thank B. Panneton and M.-J. Simard from Agriculture and Agri-Food Canada for their contributions and insights on the material presented in the current paper, L. Longchamps for his time labelling the test images and Prof. C. Gagné for his feedback and comments on the paper.

References

1. Aitkenhead MJ, Dalgetty IA, Mullins CE, McDonald AJS, Strachan NJC (2003) Weed and crop discrimination using image analysis and artificial intelligence methods. Computers and Electronics in Agriculture 39(3):157–171
2. Alpaydin E (2010) Introduction to machine learning, 2nd edn. MIT Press, USA
3. Asif M, Amir S, Israr A, Faraz M (2010) A vision system for autonomous weed detection robot. Int J Comput Electr Eng 2(3):486-491
4. Bowman A, Azzalini A (1997) Applied Smoothing Techniques for Data Analysis. Clarendon Press, Oxford

5. DeLorenzo ME, Scott GI, Ross PE (2001) Toxicity of pesticides to aquatic microorganisms: a review. Environmental Toxicology and Chemistry 20(1):84–98
6. Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood estimation from incomplete data via the EM algorithm. Journal of the Royal Statistical Society Series B (Methodological) 39:1–38
7. Figueiredo MAT, Jain AK (2002) Unsupervised learning of finite mixture models. IEEE Transaction on Pattern Analysis and Machine Intelligence 24(3):381–396
8. Freemark K, Boutin C (1995) Impacts of agricultural herbicide use on terrestrial wildlife in temperate landscapes: A review with special reference to north america. Agriculture, Ecosystems & Environment 52(2):67–91
9. Hemming J, Rath T (2001) Precision Agriculture Computer-Vision-based Weed Identification under Field Conditions using Controlled Lighting. Journal of Agricultural Engineering Research 78(3):233–243
10. Hough PVC (1962) A Method and Means for Recognizing Complex Patterns. U.S. Patent 3,069,654
11. Jones G, Gée C, Truchetet F (2009) Assessment of an inter-row weed infestation rate on simulated agronomic images. Computers and Electronics in Agriculture 67(1-2):43–50
12. Longchamps L, Panneton B, Samson G, Leroux G, Thériault R (2010) Discrimination of corn, grasses and dicot weeds by their uv-induced fluorescence spectral signature. Precision Agric 11:181–197
13. Ngouajio M, Lemieux C, Fortier JJ, Careau D, Leroux GD (1998) Validation of an operator-assisted module to measure weed and crop leaf cover by digital image analysis. Weed Technol 12(3):446–453
14. Otsu N (1979) A threshold selection method from gray-level histograms. IEEE Transaction on Systems, Man, and Cybernetics 9(1):62–66
15. Panneton B (2009) Initiative des stratégies de réduction des risques liés aux pesticides. Rapport annuel PRR07–10, Centre pour la lutte antiparasitaire – Agriculture et Agroalimentaire Canada
16. Panneton B, Brouillard M (2009) Colour representation methods for segmentation of vegetation in photographs. Biosyst Eng 102:365–378
17. Panneton B, Simard MJ, Leroux GD, Légère A (2010) Mise au point et impact sur la distribution spatio-temporelle des adventices d'un système d'aide à la décision pour l'application des herbicides en maïs-soya. Rapport final PRR07–10, Centre pour la lutte antiparasitaire – Agriculture et Agroalimentaire Canada
18. Parzen E (1962) On estimation of a probability density function and mode. Annals of Mathematical Statistics 33:1065–1076
19. Rosenblatt M (1956) Remarks on some nonparametric estimates of a density function. Annals of Mathematical Statistics 27:832–837
20. Slaughter DC, Giles DK, Downey D (2008) Autonomous robotic weed control systems: A review. Computers and Electronics in Agriculture 61(1):63–78, DOI 10.1016/j.compag.2007.05.008
21. Stafford JV (2000) Implementing precision agriculture in the 21th century. J Agric Eng Res 76:267–275
22. Tellaeche A, Pajares G, Burgos-Artizzu XP, Ribeiro A (2011) A computer vision approach for weeds identification through Support Vector Machines. Applied Soft Computing 11(1):908–915
23. Timmermann C, Gerhards R, Kuhbauch W (2003) The economic impact of site-specific weed control. Precision Agric 4(3):249–260
24. Yang CC, Prasher SO, Landry JA, Ramaswamy HS (2003) Development of an image processing system and a fuzzy algorithm for site-specific herbicide applications. Precision Agric 4(1):5–18
25. Åstrand B, Baerveldt AJ (2005) A vision based row-following system for agricultural field machinery. Mechatronics 15(2):251–269