Computers and Electronics in Agriculture 70 (2010) 199–208

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

A segmentation algorithm for the delineation of agricultural management zones

Moacir Pedroso[a,*,1], James Taylor[b], Bruno Tisseyre[c], Brigitte Charnomordic[d], Serge Guillaume[c]

[a] EMBRAPA, SGE, Av W3 Norte (final), 70770-901 Brasilia, DF, Brazil
[b] INRA, UMR LISAH and UMR ITAP, 2 Place Pierre Viala, Montpellier, France
[c] Cemagref, UMR ITAP F-34196 Montpellier, France
[d] INRA, UMR ASB, 2 Place Pierre Viala, Montpellier, France

ARTICLE INFO

Article history:
Received 26 March 2009
Received in revised form 1 October 2009
Accepted 13 October 2009

Keywords:
Management classes
NDVI
k-means clustering
Viticulture
Irregular grid
Effective zones
Voronoi tessellation

ABSTRACT

In this paper we present a segmentation algorithm, inspired from an image-processing region-merging algorithm, for the delineation of discrete contiguous management zones in agriculture. The algorithm is unique in that it is applicable to high- or low-density irregular data sets, such as yield data. The algorithm is described and a brief example presented using unprocessed sensor-derived grain yield data. A comparison between the segmentation algorithm and a common classification algorithm (*k*-means clustering) was done using an aerial normalised differences vegetation index (NDVI) image collected on a 200 ha vineyard in Olite, Southern Navarre, Spain. Classification was performed as a univariate (NDVI) analysis and a spatially constrained analysis. Segmentation and classification were run to find 2, 4, 6, . . ., 24 levels and the effectiveness of the outputs determined by how well it explained the variance in vine trunk circumference, a correlated but independent measurement. The results obtained demonstrated that for a given number of manageable (effective) zones the segmentation outputs were equivalent or superior to the classification outputs for partitioning vine circumference variance. The segmentation output also generated more coherent management units that should facilitate differential management. The algorithm presented is a first generation segmentation algorithm and several aspects still need to be developed, in particular methods for eliminating edge effects and converting management zones into management (treatment) classes. The results of the segmentation algorithm presented here would indicate that with further development, segmentation might provide an alternative and possibly preferable approach to delineating management zones.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Management units are a common intermediate step of moving from a uniform 'field average' management system to a site-specific crop management (SSCM) system. The concept of management units was proposed in the mid 1990s (Lark and Stafford, 1997) and there have been many different delineation methods recorded in the precision agriculture (PA) literature (see Taylor et al., 2007; Roudier et al., 2008 for example). Although Lark and Stafford (1997) originally proposed the term 'management units', the term 'management zones' is generally used. This can be somewhat ambiguous because the majority of statistical methods for management unit delineation, including the original *k*-means approach by Lark and Stafford (1997), are based on classification algorithms. These approaches produce 'management classes'. The difference is important.

A *management zone* (MZ) is a spatially contiguous area to which a particular treatment may be applied.

A *management class* (MC) is the area over which a particular treatment may be applied. This may constitute more than one zone.

A *management unit* (MU) is a generic term that relates to both MCs and MZs.

There is no doubt that classification algorithms are effective at explaining the variance in both univariate and multivariate cases. However, classification algorithms do have some limitations. Firstly, as the name suggests, they produce management classes not management zones. There is no constraint within the algorithm to form contiguous classes, i.e. zones. Differential management and profitability in agricultural systems is dependent on both the magnitude of variation, how it is partitioned, and the spatial structure of this partitioning (or the morphology of the classes/zones) (Pringle et al., 2003; Tozer and Isbister, 2007). Classification algorithms generally address the first issue well but not the second. There have been a few attempts (McBratney et al., 2000; Shatar and McBratney, 2001; Ping and Dobermann, 2003; Frogbrook and Oliver, 2007) to spatially constrain the clustering algorithm to produce management zones instead of classes, but these have not been widely

* Corresponding author. Tel.: +55 61 34484395; fax: +55 61 34484319.
E-mail address: moacir.pedroso@embrapa.br (M. Pedroso).
[1] On leave at INRA, UMR ASB, 2 Place Pierre Viala, Montpellier, France.

0168-1699/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.compag.2009.10.007

adopted. Most approaches and software in precision agriculture (PA) still rely on unconstrained clustering (e.g. Fridgen et al., 2004; Taylor et al., 2007).

Another major issue with the use of classification algorithms for management class delineation is an absence of a clear statistic or indication of the optimum number of classes (Cupitt and Whelan, 2001). This usually requires some expert intervention and is further compounded by the presence of small, discrete, unmanageable zones within the classes. Finally, it is important to remember that multivariate classification relies on co-located data, although these do not need to be on a regular grid.

An alternative to classification for variance partitioning is the introduction of segmentation algorithms that are routinely used in image processing. However, despite the use of imagery in PA, segmentation algorithms have not been widely applied (Roudier et al., 2008). Segmentation algorithms differ from classification algorithms in that they are object-oriented (note: the term "object-oriented" here is used in its image analysis context, not a software engineering context). This focus leads to the production of discrete zones rather than classes and the output is spatially structured. This overcomes one of the main issues with classification algorithms. Segmentation algorithms aim to build homogenous regions with respect to a given criterion. They either start with the whole image (data) and split it in an iterative process or commence with a large number of regions that are iteratively merged.

One of the disadvantages with many object-oriented segmentation algorithms is a reliance on regular, grid data for determining segment morphology. This is probably an artefact from their primary application in image analysis and has restricted the use of these algorithms on irregular PA data sets, for example raw yield data or site-directed soil sample data. Another drawback to their application in PA has been a lack of research into methods of optimising variance partitioning during the aggregation stage. It is important to realise that segmentation produces zones, however from a management perspective two or more discrete zones may be managed identically (i.e. as a management class). There is very little literature on how to aggregate agronomic data (using agronomic knowledge) and convert management zones into management classes.

The aims of this paper are therefore (a) to present a segmentation algorithm that is able to process irregular grid data, (b) illustrate the algorithm on an irregular dense data set, and (c) compare and contrast this algorithm with the more common *k*-means classification algorithm. The paper is divided into several sections. The next section introduces the algorithm. Section 3 demonstrates the application of the algorithm on an irregular yield data set. Section 4 provides a comparison and discussion on variance partitioning between the segmentation algorithm and *k*-means classification. Section 5 is a general discussion and, finally, Section 6 presents the conclusions.

2. The segmentation algorithm

In agricultural system analysis, it is important to take into account spatial structures. Spatial data analysis is usually based on clustering. If clustering is only done in the attribute space, spatial continuity can be achieved by *a posteriori* filtering. Alternatively, a common method to impose a spatial dimension on the output is to include point coordinates (x,y) as variables in the clustering process along with attribute data (McBratney et al., 2000; Davis et al., 2007; Castrignano et al., 2009). This is a simplistic approach to spatial constraint and while more elegant and geo-statistically robust models exist (Webster and Oliver, 1989; Lark, 1998; Frogbrook and Oliver, 2007) these are still based on the weighting of coordinates in the clustering.

Besides this classical data analysis approach, image analysis provides completely different methods where the spatial properties are of first concern. The process of defining zones inside an image is called segmentation. The segmentation methods (Coquerez and Philipp, 1995) can be classified into two main families: the contour-based ones and the region-based ones. The first family is more suitable for object recognition, and the second one is useful when there are no well definite borders. This last case may well correspond to agricultural data zoning.

The segmentation algorithm presented here is inspired from a region-merging algorithm (Lardon et al., 2006), which belongs to the second family described above (Chassery and Garbay, 1984).

The outline is given below:

```
- Start: one point = one zone
- Iterate on:
   Merge the pair of neighbouring zones that are
   closest in the attribute space.
   Update zone list and zone neighbours
- Until any stop criterion is met
```

2.1. Use of spatial coordinates in this algorithm

A fundamental point is the way the spatial coordinates are used here. They are not involved in any distance calculation, but are only used to define point and zone neighbourhood. The attribute proximity criterion used for zone merging is only calculated within a given neighbourhood.

2.2. Details

Finding zone neighbours on an irregular grid requires a specific process, which is a major difference when compared with segmentation methods in image analysis.

To start, a Voronoi tessellation is used to convert each data point to a zone and to define the initial neighbourhood, as illustrated in Fig. 1. Each point is associated with a unique polygon, and the list of its neighbours is built, including all adjacent polygons. Initially there are as many zones as data points, and the zone neighbourhood is equivalent to the point neighbourhood.

At each step of the zone merging process, the zone neighbourhood is updated by considering all zones that share a vertex in the tessellation as neighbours.

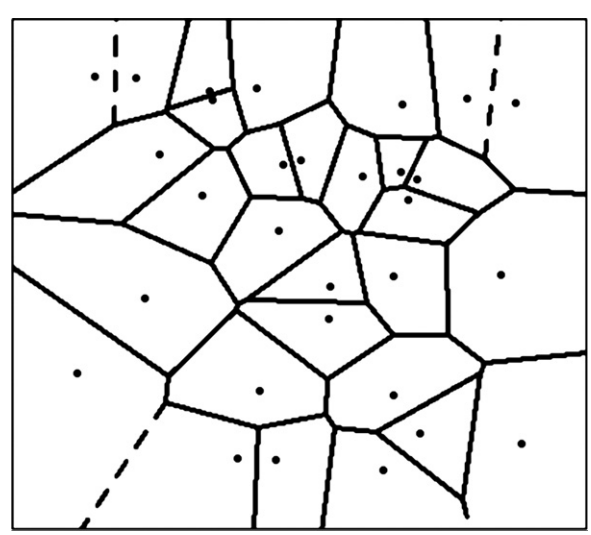

Fig. 1. Illustration of the Voronoi tessellation.

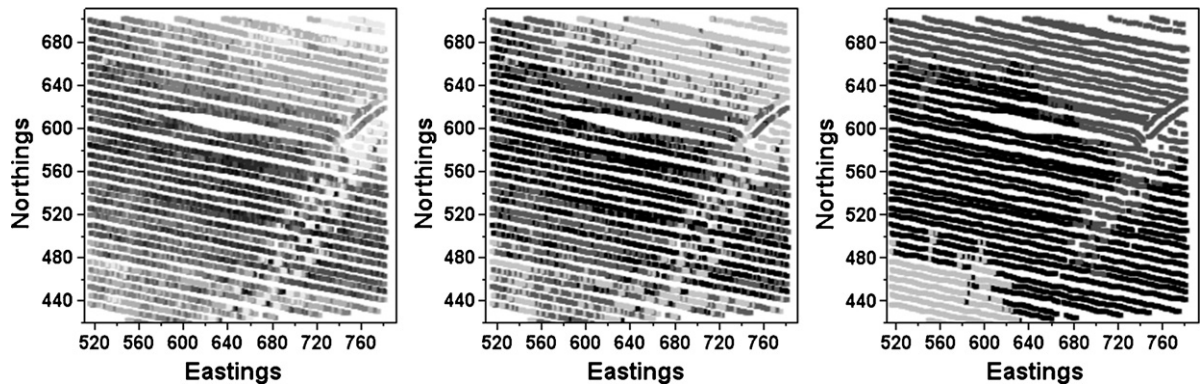

Fig. 2. An illustrated example of *k*-means clustering and the segmentation algorithm applied to irregularly collected sorghum yield data. From left to right the maps show (a) the raw data divided into 0.5 ton/ha increments (grading from low (light grey) to high (black) yield), (b) a 3-class *k*-means cluster map based on the yield data and (c) the 3-zone map from the segmentation algorithm. Different classes/zones are shown as different colours (light grey, grey or black).

The attribute is a one-dimensional numerical value. Closeness in the attribute space is evaluated according to the furthest distance. For each pair of zones, Z_1 and Z_2, this distance corresponds to the maximal difference in absolute value computed over all pairs of points (p_i, p_j), p_i in Z_1, p_j in Z_2: $d(Z_1, Z_2) = \max_{i,j} |\mathrm{attr}(p_i) - \mathrm{attr}(p_j)|$.

Several stop criteria can be used and the algorithm stops when any of them is met. Some are related to zone intra-variability: they are computed from the distribution of attributes within each zone, such as range, mean and variance. The Moran coefficient can also be used to describe the overall spatial variability. Finally the stop criteria can be simply set at a certain number of zones. In this paper, only the number of zones is used as a stopping criterion to illustrate the segmentation approach.

The algorithm can be considered a first generation PA segmentation algorithm. It is capable of handling large irregular data sets, which overcomes a major limitation from previous approaches in this area (Shatar and McBratney, 2001). It is presented here in its simplest form as a general framework for segmentation algorithms to be applied to regular or irregular grids. As stated above, the algorithm could be parameterised by a range of different criteria for the distance metric and the stopping criteria. It could also be applied in a univariate or multivariate context. The univariate approach taken to demonstrate the algorithm in the following case studies is only one example from the family of possible algorithms. We do not claim that it is necessarily the best possible one.

3. Application to a dense irregular grid data

The primary advantage to the proposed algorithm is its ability to segment dense data presented on an irregular grid. This functionality is illustrated in Fig. 2. The data is a 9 ha square section of sorghum (*Sorghum* sp.) yield data taken from a larger (~100 ha) field in north-western NSW, Australia.

The raw yield data (Fig. 2, left) has been trimmed to remove excessive outliers (Taylor et al., 2007) however it still retains irregularities associated with GPS drift and changes in directions. The data is also denser along the transects than across them. The raw yield map illustrates the noise associated with individual point data. In yield monitoring systems there is a general trade off between the density of data collected and the accuracy of individual point measurements. The theory is that even if a value at an individual point is erroneous, averaging points around a site provides a good estimation. Despite the noise, a broad trend of low-high-low yield from the north-east to the south-west corner is clear.

Two management unit maps are presented together with the raw data. The first (Fig. 2, centre) shows a univariate *k*-means clustering of the yield data and the second (Fig. 2, right) the output from the proposed segmentation algorithm. The univariate *k*-means highlights the trend however retains the noise associated with the point data giving a 'salt and pepper' effect, particularly in the south-west corner. In contrast, the segmentation output presents three relatively smooth zones and again highlights the trend that is observable in the raw data. There is some noise at the boundary of the zones associated with the neighbourhood analysis on the irregular grid generating thin incursions. These are unlikely to be manageable, however, this effect is much less problematic for management than those observed in the *k*-means maps.

The objective in this section is to illustrate the difference between the 3-level classification and segmentation, even if the 3-level choice is not the optimal solution for classification or segmentation. A more detailed analysis and discussion of variance partitioning between methods is presented in the next section.

4. Illustration and discussion on variance partitioning for classification and segmentation

The previous section illustrated the spatial output of the segmentation algorithm compared to the widely used *k*-means algorithm. It did not consider how well the variance was partitioned between the two approaches or discuss in any depth the practical implications of adopting the outputs as the basis for a management unit map.

Intuitively, a classification algorithm with no spatial constraints will provide the optimum partitioning of the variance in the data, as it delineates classes solely in the attribute space. However, as can be seen in Fig. 2, delineation in the attribute space does not always produce a result that can be readily applied in practice (the physical space).

In this section we consider the trade off between spatial organisation of the management units and the partitioning of variance from the classification and segmentation algorithms. This was achieved by analysis of a normalised differences vegetation index (NDVI) image and determining how well the different algorithms partitioned variance in vine trunk circumference, a related vine parameter. The algorithm outputs were not directly compared to the NDVI response for two reasons. Firstly, given that a univariate *k*-means classification has no spatial constraints, this will always produce the best result as the classes are only partitioned on the variance of the attribute (NDVI), which may not necessarily prove practical for management. Secondly, high-density ancillary data (soil sensor data, imagery, etc.) are often used as surrogates for actual crop parameters. Therefore, growers are interested in how ancillary data partitions a crop response, in this case trunk

circumference. The exception to this rule is when harvest sensor data is directly analysed. The data and methodology for this comparison are explained below.

4.1. Data

The data to be classified/segmented was an aerial NDVI image of a 200 ha vineyard located in Olite, Southern Navarre, Spain, acquired in July 2007 by GEOSYS SL (Madrid, Spain) using a Leica ADS40 sensor. All the fields were planted with the variety Tempranillo in an east-west orientation and a plant spacing of 2.5 m × 1.1 m. The image was originally collected at 0.3 m pixels and post-processed into 5 m pixels using a moving average window. The data was then interpolated onto a continuous 5 m^2 grid using punctual kriging to fill in the inter-field sections (roads). This was done purely to simplify the analysis and comparison. The presence of roads (gaps) in the image will impose more divisions within classes and zones. The continuous NDVI map will make the zones and classes as contiguous as possible. We wish to stress that interpolation is not necessary for the algorithm to run, as shown in the previous section. The final processed NDVI image is shown in Fig. 3.

4.2. Classification and segmentation analysis

Classification was performed with the commonly used *k*-means clustering algorithm using the R Statistical Package (R Development Core Team, 2008). The clustering was first run as a univariate analysis. This effectively partitions the NDVI histogram, similar to a map legend, although the interval between classes may not be constant and will be dependant on the shape of the histogram.

The coordinate points were then included in the *k*-means clustering analysis to spatially constrain the clusters (McBratney et al., 2000; Castrignano et al., 2009). The spatially constrained *k*-means classification is partitioning the variance of both the coordinate and NDVI data, i.e. it is trivariate. The clustering algorithm was run to generate output from 2–24 clusters in increments of 2 (i.e. 2, 4, 6, . . ., 22, 24 clusters) for both the univariate and spatially constrained approach.

The segmentation algorithm was run on the univariate NDVI data to generate from 2 to 24 zones in steps of 2 zones, as for the classification. The number of final zones was the only stopping criteria used. Bear in mind that although the relationship between neighbours is considered in the algorithm, the segmentation is univariate and does not directly incorporate the coordinate data.

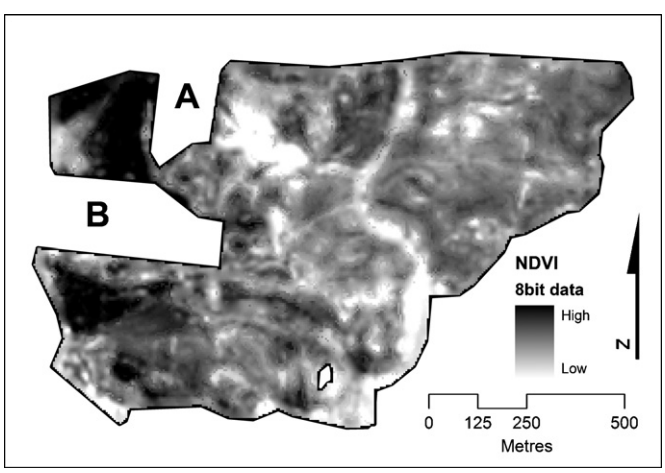

Fig. 3. Normalised differences vegetation index image of a 200 ha established vineyard in Southern Navarre, Spain. Missing data in the bottom centre is the location of farm buildings. Letters A and B relate to areas where neighbouring points occur across gaps.

4.3. Validation of output

The output from the classification and segmentation analysis was imported into ArcMap® v9.2 (ESRI, Redlands, CA, USA) and classes/zones were converted into polygon data. The number of polygons and area of each polygon for each level (2–24) for each approach (classification and segmentation) were recorded. The number of polygons indicates the number of zones (discrete areas) in the output. With a regular shaped field the number of zones in the segmentation algorithm output should equal the number of polygons generated. With non-convex area fields, neighbours, and thus aggregation, can occur across gaps. This is a difference between aggregation in the attribute space rather than the physical space. The result is that it is possible to create several discrete regions in real space that equate to one discrete area in the attribute space.

From an agronomic perspective small zones are undesirable and may be impossible to manage due to technological limitations (Tisseyre and McBratney, 2008). For this analysis, zones less than 0.1 ha were considered to fall into this category. For other production systems this threshold area may be higher or lower depending on the resolution of management possible with current technology. For each level in each approach, the number of polygons with an area below the threshold (0.1 ha) was subtracted from the total number of polygons to give the number of effective (manageable) zones (Z_E) resulting from the analysis. The summed area of the small polygons (<0.1 ha) was also recorded as the unmanageable area (A_U). Since these polygons are below the threshold level for management, they can be considered as polygons that will receive the incorrect treatment, and areas of lost opportunity.

A crop measurement correlated with ($r = 0.61$), but independently collected from, the NDVI image is vine trunk circumference. In April 2008 there were 476 vine circumference measurements taken within the target area on a rough square grid. The vine point data were intersected with the analysis output in ArcGIS (ESRI, Redlands, CA, USA). Thus each vine measurement was allocated to a class or zone from each of the classification and segmentation outputs. The effectiveness of the management unit delineations for all methods was determined by the adjusted r^2 from ANOVA using the management units as the independent variable and vine circumference data as the dependent variable, i.e. how well do the different methods of delineation explain a key growth parameter. The adjusted r^2 was used due to irregular sample sizes. Analysis was performed in JMP® v8.0.1. (SAS Institute Inc., Cary, NC, USA). The results were plotted first using the management unit level on the abscissa and then using Z_E.

5. Results and discussion

5.1. Evolution of delineation using classification and segmentation—visual

The spatial evolution of classes and zones, from the classifications and segmentation for 2, 4, 6, 10, 14 and 24 levels is illustrated in Fig. 4. All images grade from white to black for the lowest to highest NDVI response. As expected, the higher order univariate classification (>6 classes) produces maps visually similar to Fig. 3, i.e. they appear to have a continuous scale. These maps are presented without defining the border for the classes/zones. When a border was imposed, the large number of small zones in the images (see Table 1) generated a map that was illegible. Zone borders are defined in the 2, 4 and 6-class maps.

The influence of the spatial coordinates on the clustering can be clearly seen in Fig. 4b. These maps have more coherent zones than the univariate classification maps as classes are generally spatially constrained to defined regions of the vineyard. However, they are not contiguous and within a region there may be 2 or more

Table 1
Comparison of the total number of zones, number of effective zones (Z_E) and unmanageable area (A_U) derived from the three approaches (univariate clustering, spatially constrained clustering and the proposed segmentation algorithm) for a range of 2–24 levels. Bold numbers identify levels at which the three approaches produce a similar Z_E result.

Level	Total zones			Effective zones (Z_E)			Area unmanageable (A_U) (%)		
	Univariate clustering	Spatial clustering	Segmentation	Univariate clustering	Spatial clustering	Segmentation	Univariate clustering	Spatial clustering	Segmentation
2	181	50	7	**28**	6	4	2.14	0.54	0.18
4	468	244	10	63	**30**	6	5.34	2.44	0.24
6	955	275	14	129	37	8	12.60	3.10	0.24
8	1989	268	16	159	38	10	27.74	2.54	0.26
10	3252	373	19	171	43	12	41.17	3.92	0.34
12	4488	411	21	155	47	13	52.48	4.42	0.45
14	5919	401	26	145	48	17	62.42	4.35	0.45
16	7223	407	28	132	52	19	69.76	4.28	0.45
18	8497	475	30	95	48	21	77.53	5.01	0.45
20	9688	494	33	89	51	24	81.76	5.01	0.45
22	11093	582	35	74	53	26	85.63	6.27	0.45
24	12089	612	37	56	58	**28**	89.12	7.77	0.18

classes interspersed. Small and likely unmanageable zones are still present in the maps. In contrast, the segmentation maps (Fig. 4c) presented much clearer zone delineation. The effect of segmenting based on neighbours, regardless of distance between neighbouring points, is evident in the western section of the vineyard. The dark segment in the 2-zone map is itself split into two zones separated by the gap in the data (labelled B in Fig. 3). Likewise the data gap along the northern boundary (labelled A in Fig. 3) causes some

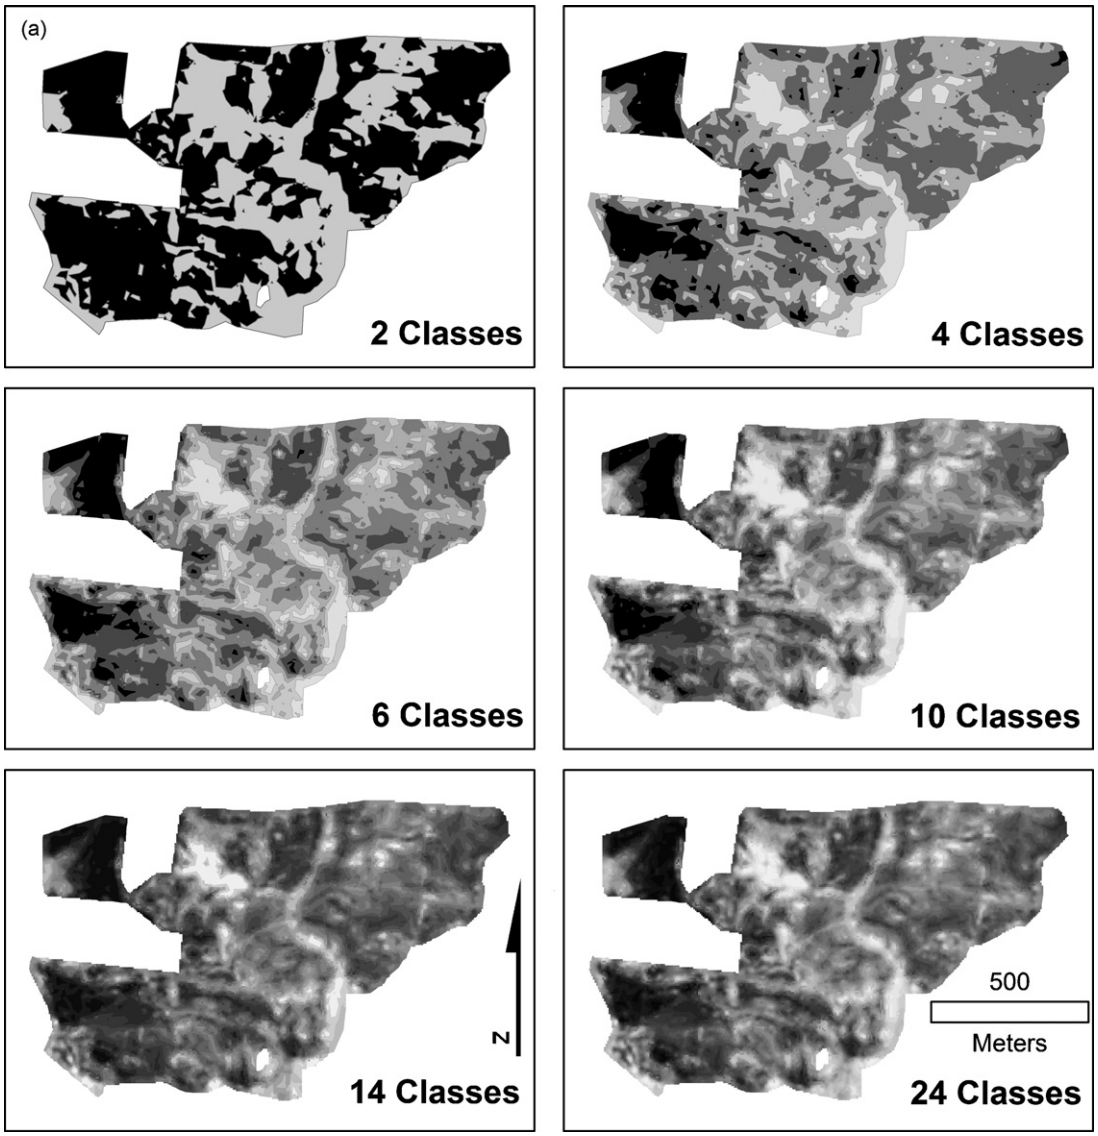

Fig. 4. Evolution of management units using the univariate *k*-means clustering algorithm (a), spatially constrained clustering (b) and the segmentation algorithm (c). 2, 4, 6, 10, 14 and 24 level maps shown.

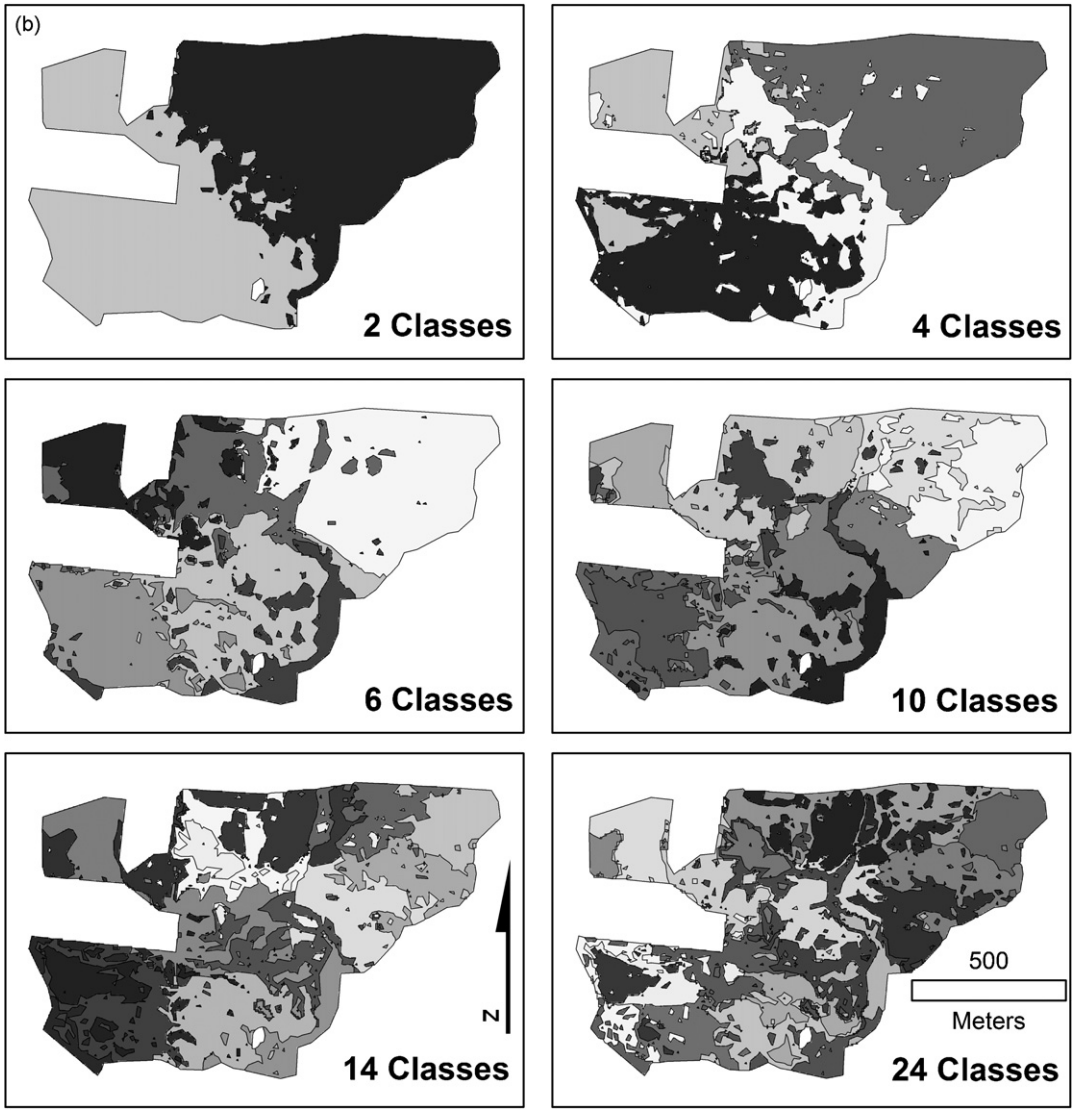

Fig. 4. (*Continued*)

small edge-effect zones that are present in all the segmentation results.

5.2. Validation of classification and segmentation using trunk circumference data

The amount of variation in trunk circumference measurements explained by the different levels of classification and segmentation models is shown in Fig. 5. As expected the classification procedures explained more of the variation in trunk circumference for a given level. The univariate classification becomes asymptotic after 6 classes (at $r^2 \sim 0.62$). The spatially constrained classification achieves the same asymptote but after 14 classes. The segmentation analysis lags behind the two classification approaches and in this range of levels only explains 46% of the variance in trunk circumference.

However, Fig. 5 compares classes with zones. The abscissa is the level of classification or segmentation. The relationship between classification levels (classes), segmentation levels (zones) and Z_E is shown in Table 1.

The first thing that is notable in Table 1 is the large number of zones generated by the univariate (spatially unconstrained) *k*-means clustering. The total and effective number of zones generated is significantly larger than the spatially constrained clustering. In Fig. 5 there appears to be no benefit to univariate clustering beyond 6 classes, however for comparison the univariate results up to 24 classes have been shown in Table 1.

Fig. 5 also indicated that the 6-class univariate and 14-class spatially constrained classifications explained the same level of variance in the trunk circumference data ($r^2 \sim 0.62$). An analysis of these two responses shows that the spatially constrained classification has fewer Z_Es (48 c.f. 129). Applying the principle of parsimony, the spatially constrained clustering would be preferred. Visually, the 14-class spatially constrained classification map (Fig. 4b) is more coherent than the 6-class univariate classification map (Fig. 4a). To add to this, the unmanageable area in the spatially constrained classification output (4.35%) is approximately a third of that generated in the univariate classification (12.60%), again indicating a preference for the spatially constrained clustering. For the univariate clustering, the percentage of area assigned to zones <0.1 ha rises rapidly after 6 classes and by the 12-means clustering accounts for more than half the total area.

Following the same line of reasoning, the 4-class univariate and 8-class spatially constrained clustering output can be compared as they have similar adjusted r^2 values ($r^2 \sim 0.51$). Again the spatially constrained output has fewer Z_Es and a lower A_U.

Fig. 4. (*Continued*).

Alternatively, the two classification procedures can be compared at equivalent Z_Es. The 2-class univariate and 4-class spatially constrained output produce similar results (28 and 30 Z_Es and 2.14 and 2.44% A_U, respectively) for management. However, the spatially constrained output explains a significantly greater percent of the trunk circumference variance (42% c.f. 32%), thus for approximately equal Z_Es and A_U, the spatially constrained clustering appears preferable for management. Given these results, the spatially constrained classification model is used as a reference against which to compare the segmentation approach.

The segmentation model explained less of the trunk circumference variance than that explained by the classification for a given level (Fig. 5). However, the maximum number of Z_Es attained by segmentation was 28 (24-segment analysis). Comparing this output to the 4-class spatially constrained clustering (30 Z_Es) the segmentation explains slightly more of the trunk circumference variance (46%) and has only 0.18% of the field assigned to zones <0.1 ha. The segmentation algorithm is therefore performing as well if not slightly better than the spatially constrained clustering algorithm.

Higher order segmentation analyses (26–70 segments in increments of 2) were run to generate a segmentation output equivalent to the 14-class (48 Z_Es) spatially constrained clustering. This occurred with a 52-segment analysis, which generated 66 segments in total and 48 Z_Es. The amount of variance in the trunk circumference data explained by this 52-segment model was not significantly different from the 14-class spatially constrained model (62% and 61%, respectively). The A_U statistic was lower for the segmentation output (0.90%). To visualise the response of the classification and segmentation in terms of Z_Es, Fig. 5 has been redrawn with Z_E as the abscissa (Fig. 6) and with the results from the higher order segmentation analysis included.

Fig. 6 presents a more accurate comparison between the classification and segmentation approaches. Beyond 30 Z_Es the spatially constrained classification and segmentation outputs yield a similar response. The perceived advantage in variance partitioning with classification is not present in these results. The main difference between classification and segmentation is the range of Z_Es generated by segmentation. The univariate classification only provides two options in this range of Z_Es, and the first result (2-class or 28-Z_E) is inferior to the equivalent spatially constrained and segmentation results (as discussed earlier). The spatially constrained classification generates multiple results >30 Z_Es but only one <30 Z_Es (6 Z_E). The segmentation results present many more

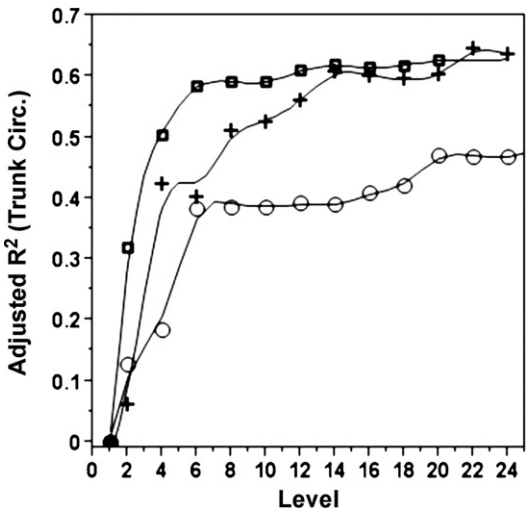

Fig. 5. Fits (adjusted r^2) of the MU models to the trunk circumference data. The abscissa shows the number of management units (levels of classes or segments). Response from the univariate *k*-means clustering (classification) shown by (□), the spatially constrained clustering by (+) and segmentation results shown by (○). A spline has been fitted through the data to show the trend.

options for potential management maps, without compromising on the amount of attribute variance explained. Interestingly, the segmentation result appears nested, with a different response pre- and post-6-zones but the reason for this is unclear without more detailed analysis.

While the percentage of variance in trunk circumference explained is not high ($r^2 \sim 0.4$–0.6), from an agronomic perspective the ability to explain (or manage) half the variation in a production system can make a significant improvement in production efficiency.

6. General discussion on the segmentation algorithm and effective zones

The concept of an effective zone (Z_E) has been introduced to help compare outputs. This concept was derived to use the principle of parsimony as applied to management units (Lark, 2001). It is possible to apply this principle of parsimony in two ways–to zones (or Z_Es) or to classes (treatments). This analysis has concentrated on the former, which focuses on simplifying the mechanical aspects of differential treatment, not necessarily simplifying agronomic decision-making (i.e. determining differential input rates). A current limitation to this segmentation analysis is that the zones have not been assigned into treatment classes to achieve effective agronomic decision-making, thus limiting our ability to apply the principle of parsimony to treatment classes. Local agronomic will define this re-assignment and production conditions will likely be crop and region/site-specific. Nevertheless, developing support tools to transform zones into classes is an important next step for segmentation analysis. Since this is an early foray into this area, the analysis has focused on the parsimony of Z_Es. In this context segmentation outperforms classification.

A quick test of the parsimony of treatment classes can be achieved by observing the distribution of the mean trunk circumference response for the Z_Es derived from the three approaches. The ideal response would have the mean Z_E responses arranged into a multinomial distribution with discrete peaks. A treatment response would then be based on the peaks in the distribution. Given the methods of variance partitioning used, any methodological advantage should favour the classification algorithms.

Outputs that generated 28–30 Z_Es for each method (the 2-class univariate, 4-class spatially constrained and 24-segment analysis) were chosen for comparison. The mean trunk circumference for each Z_E was determined and the histograms of Z_E mean values plotted (Fig. 7). The univariate classification produces the most compact range of mean Z_E values (55–95 mm), but this is derived from only a 2-class analysis (see Fig. 4a for the map). Both the spatially constrained clustering and segmentation algorithm histogram contain an outlier, which correspond in both cases to the Z_E located in the north-west corner of the vineyard (Fig. 4b: 4-class and Fig. 4c: 24-zones). The 2-class univariate analysis was unable to identify this area as a unique zone, despite generating the same number of Z_Es, as the Z_E that covers this area is contiguous with a large area within the main section of the vineyard. This unique zone in the north-west corner is also delineated in the lower-order segmentation maps (Fig. 4). For the rest of the spatially constrained classification and segmentation histograms, the shape is similar to the univariate

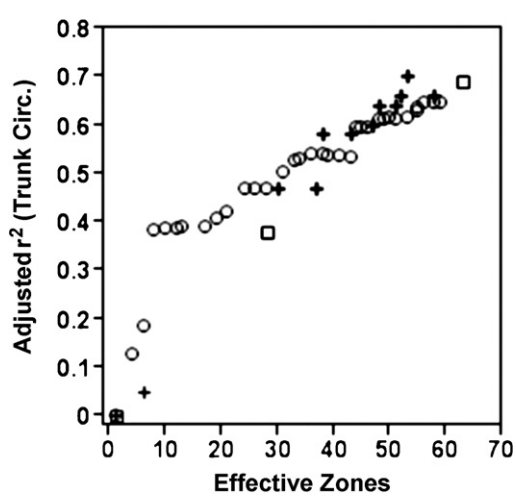

Fig. 6. Fits (adjusted r^2) of the MU models to the trunk circumference data plotted against Z_E. Legend is the same as Fig. 5.

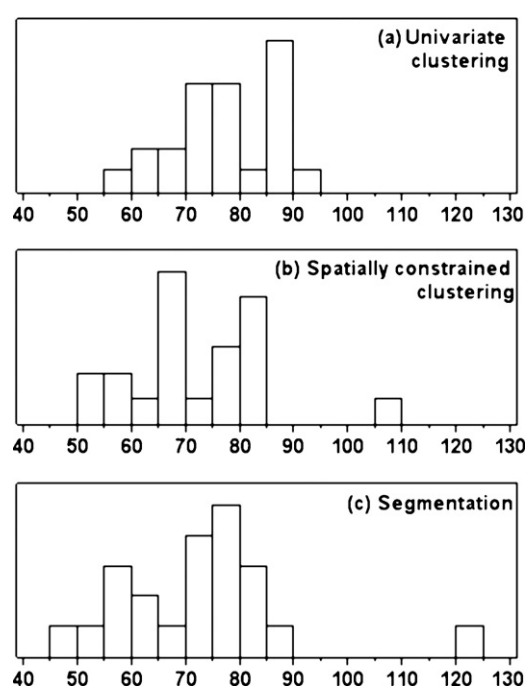

Fig. 7. Histogram of mean trunk circumference (mm) within each effective zone for (a) univariate clustering (2-class), (b) spatially constrained clustering (4-class) and (c) segmentation (24-zones), all of which generate between 28 and 30 Z_Es.

histogram, although the values are slightly lower as the effect of the vines with a larger trunk circumference are contained in the outlier Z_E.

This simple comparison showed that with a 2-class clustering approach, the resulting Z_Es exhibit a range of mean values similar to a higher order spatially constrained clustering or segmentation approach, i.e. the mean of an individual Z_E derived from clustering does not necessarily equal the class mean to which it belongs.

Therefore, if the principle of parsimony is applied to the agronomic decision-making process, the shape of the segmentation histogram does not appear to provide any disadvantage to agronomic decision-making. The benefit of the univariate analysis lies in the lower-order treatment classes that have been defined. However, if the segmentation algorithm can be effectively aggregated into fewer treatment classes this advantage will be likely lost while retaining the advantage of the spatial structure of the segmentation zones for practical application. The analysis here is very superficial but an indicator of the hypothesis that needs further testing once zones have been properly aggregated into treatment classes.

The A_U (unmanageable area) statistic has also been used to identify differences in outputs of classification and segmentation. This is not a perfect statistic but does provide some insight into the practical application of the output. It is statistically obvious (and demonstrated in Table 1) that classification will produce smaller zones (particularly for the univariate classification) than segmentation when the stopping criteria for segmentation is based on the number of zones. Our contention is that these small zones to some extent degrade the quality of the output, thus A_U can be used to compare models. This argument contains some flaws. Small zones are present in all models, and the difference lies in whether the small zones are smoothed over within the algorithm, as in segmentation, or are retained as discrete units in the output, as in *k*-means clustering.

The smoothing of these small zones within the algorithm may not yield a better output, although the A_U will be lower. In this sense the A_U must be considered in context with other model indicators. However, merging small unmanageable zones into neighbours within the delineating algorithm does present some advantages. Firstly, it is automatic and requires no further expertise or post-processing and secondly, the merged zone will contribute to the response of the final zone. This analysis indicated that the 'smoothing' of small zones within the segmentation algorithm did not diminish the ability of the output to partition variance for a given level of Z_Es compared to classification.

In contrast, the clustering output generates many small polygons that require some form of filtering. This can be achieved by merging small zones into larger neighbouring zones, which requires another layer of data analysis, and possibly expertise. Alternatively, the small zones could be ignored and the application machinery left to 'smooth' them into neighbouring zones. This requires no further processing but does assume that the small polygon responses are not included in the larger zone response when determining input rates. Given these restrictions, we believe that A_U is a relevant, though not definitive, statistic and provides some indication of more effective zoning.

The segmentation algorithm still produces some minor artefacts. In particular edge effects are common if there is a large variance between neighbours near an edge. The ability of the algorithm to recognise neighbours across physical gaps also contributes to these edge effects. Edge effects can be clearly seen in the 4 and 6-zone maps in Fig. 4c. The 4-zone map generates a small zone (white polygon) on the centre bottom edge due to a local high point. Further development of the algorithm intends to address this issue. Whilst edge effects are an issue, these effects are much smaller and easier to manage than the small polygon 'noise' generated in the *k*-means classifications.

The aggregation component of the algorithm is currently based on merging of neighbours that are closest in the attribute space and the stopping criteria is defined as a number of zones. There are other options for aggregation and the stopping criteria. An interesting perspective of the present and on-going work is to determine optimal criteria for different agricultural cases. Aggregation is a decision-making process and it is possible and desirable to incorporate agronomic knowledge into this process. This could be, for example, achieved by incorporating a fuzzy inference system into the aggregation process. A system such as this would make zoning more targeted and hopefully more applicable for a desired management operation.

The use of zones also opens the way to the introduction of approached borders. This corresponds to a realistic scenario in many agricultural applications, where the attributes take continuous values. Contrary to object recognition, there is no unique way to define homogeneous meaningful zones, mainly because of biological variability and continuous soil variation. The use of expert knowledge, as mentioned above, can easily be incorporated in the algorithm to help to define the transitions between zones.

7. Conclusions

A segmentation algorithm has been presented that can deal with large irregular data sets. The algorithm partitions the data into zones, not classes. As the zones are contiguous, they are manageable units for agricultural practices. Conversely, as the algorithm does not produce classes, the zones have to be labelled according to a treatment class. Using the Effective Zone concept introduced here, the application of the segmentation algorithm to an auxiliary variable (NDVI) proved more effective than univariate *k*-means clustering and *k*-means clustering with a spatial constraint at partitioning the variance in a correlated but independent variable of agronomic interest (trunk circumference).

This work can be considered as a first generation segmentation algorithm for precision agriculture. Opportunities exist to adapt and apply the algorithm to different production systems through the incorporation of an agronomic as well as statistical decision process during the aggregation phase.

Acknowledgements

Dr James Taylor's Post-doctoral position in France is funded through the Agropolis Foundation. The authors are grateful to the vineyard Pagos de Araiz, Navarra, Spain, the Verdtech company (www.verdtech.es) and the Australian Centre for Precision Agriculture, The University of Sydney, for providing us with the data used in this paper.

References

Castrignano, A., Guastaferro, F., De Benedetto, D., Moneta, A., Basso, B., Troccoli, A., Pisante, M., 2009. Delineation of site-specific management zones using geostatistics and fuzzy clustering analysis. In: van Henton, E.J., Goense, D., Lokhorst, C. (Eds.), Precision Agriculture '09. Proceedings of the 7th ECPA. Wageningen, the Netherlands, pp. 537–544.

Chassery, J.M., Garbay, C., 1984. An iterative segmentation method based on a contextual color and shape criterion. IEEE Trans. on PAMI 6 (6), 794–800.

Coquerez, J-P., Philipp, S., 1995. Analyse d'images: filtrage et segmentation. Masson, Paris, ISBN: 2-225-84923-4, ISSN: 0992-5538.

Cupitt, J., Whelan, B.M., 2001. Determining potential within-field crop management zones. In: Grenier, G., Blackmore, S. (Eds.), ECPA 2001. Proceedings of the 3rd European Conference on Precision Agriculture. Montpellier, France, agro-Montpellier, pp. 7–12.

Davis, D., Fountas, S., Aggelopoulou, T., Gemtos, A., Velidis, G., 2007. Management zones for high valued crops in Greece. In: Aggelopoulou, A., Blackmore, S., Gemtos, T., Fountas, S. (Eds.), Poster proceedings of the 6th European Conference on Precision Agriculture. Skiathos, Greece, 6 pp.

Fridgen, J.L., Kitchen, N.R., Sudduth, K.A., Drummond, S.T., Wiebold, W.J., Fraisse, C.W., 2004. Management zone analyst (MZA): software for sub-field management zone delineation. Agron. J. 96, 100–108.

Frogbrook, Z.L., Oliver, M.A., 2007. Identifying management zones in agricultural fields using spatially constrained classification of soil and ancillary data. Soil Use Manage. 23 (1), 40–51, doi:10.1111/j.1475-2743.2006.00065.x.

Lardon, L., Guillaume, S., Charnomordic, B., 2006. Extraction de zones floues sur des cartes. In: Actes des Rencontres francophones sur la logique floue et ses applications LFA 06, Toulouse, October 2006, pp. 321–328.

Lark, R.M., 1998. Forming spatially coherent regions by classification of multivariate data. International Journal of Geographical Information Science 12, 83–98.

Lark, R.M., 2001. Some tools for parsimonious modelling and interpretation of within-field variation of soil and crop systems. Soil Till. Res. 58, 99–111.

Lark, R.M., Stafford, J.V., 1997. Classification as a first step in the interpretation of temporal and spatial variability of crop yield. Ann. Appl. Biol. 130, 111–121.

McBratney, A.B., Whelan, B.M., Taylor, J.A., Pringle, M., 2000. A management opportunity index for precision agriculture. In: Robert, P.C., Rust, R.H., Larson, W.E. (Eds.), Precision Agriculture, Proceedings of the 5th International Conference on Precision Agriculture. ASA/CSSA/SSSA, Madison, Wisconsin (CD publication).

Ping, J.L., Dobermann, A., 2003. Creating spatially contiguous yield classes for site-specific management. Agron. J. 95, 1121–1131.

Pringle, M.J., McBratney, A.B., Whelan, B.M., Taylor, J.A., 2003. A preliminary approach to assessing the opportunity for site-specific crop management in a field, using a yield monitor. Agric. Syst. 76, 273–292.

R Development Core Team, 2008. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, http://www.R-project.org.

Roudier, P., Tisseyre, B., Poilvé, H., Roger, J.-M., 2008. Management zone delineation using a modified watershed algorithm. Precis. Agric. 9 (5), 233–250, doi:10.1007/s11119-008-9067-z.

Shatar, T.M., McBratney, A.B., 2001. Subdividing a field into contiguous management zones using a K-zones algorithm. In: Grenier, G., Blackmore, S. (Eds.), ECPA 2001, Proceedings of the 3rd European Conference on Precision Agriculture. Montpellier, France, agro-Montpellier, pp. 115–120.

Taylor, J.A., McBratney, A.B., Whelan, B.M., 2007. Establishing management classes for broadacre grain production. Agron. J. 99, 1366–1376.

Tisseyre, B., McBratney, A.B., 2008. A technical opportunity index based on mathematical morphology for site-specific management: an application to viticulture. Precis. Agric. 9 (1–2), 101–113.

Tozer, P.R., Isbister, B.J., 2007. Is it economically feasible to harvest by management zone? Precis. Agric. 8, 151–159.

Webster, R., Oliver, M., 1989. A geostatistical basis for spatial weighting in multivariate classification. Math. Geol. 21 (1), 15–35.