



Representing soil landscapes from digital soil mapping products – let the map speak for itself


David G. Rossiter^{1,2} and Laura Poggio¹

¹ISRIC-World Soil Information, Wageningen (NL)

²Section of Soil & Crop Sciences, College of Agriculture & Life Sciences, Cornell University, Ithaca NY 14850 (USA)

Correspondence: David G. Rossiter (david.rossiter@isric.org)

* using vars + depths ... space vs. depth variance ...
↳ does this generalize across
"simple" + "complex" horizons?





Abstract. Since the earliest days of soil geography, it has been clear that soils occur in more-or-less clearly mappable bodies, within which soil forming factors have been fairly homogeneous or in a regular pattern, and between which there is usually a clear transition in one or more factors. This has been the basis for polygon-based soil mapping: make a concept map from landscape elements leading to a mental model of the landscape, confirm or modify it with strategically-placed observations, find the transitions, delineate the soil bodies, and characterise them. By contrast, common methods of Digital Soil Mapping (DSM) predict per pixel over a regular grid, from training observations at pedon support. Accuracy assessment of DSM products has been at this “point” support, ignoring the existence of spatial soil bodies and the relations between pixels. Different approaches to DSM – datasets, model forms, analyst choices – result in maps with distinctly different patterns of predicted soil properties or types. Techniques from landscape ecology have been used to characterize spatial patterns of DSM products. The question remains as to how well these products reproduce the actual soil patterns at a given cartographic scale and categorical level of detail. Our approach is to let DSM maps “speak for themselves” to reveal spatial patterns. We do this by grouping pixels, either (1) by aggregation based on property homogeneity using the `supercells` algorithm, or (2) by segmentation based on within-block property pattern similarity, using the GeoPAT suite of computer programs. Segments can be hierarchically clustered into groups of presumed soil landscape elements. Supercells and segments can be compared to existing soil maps, other land resource maps, and expert judgement. To the extent that presumed soilscape patterns are reproduced, this is evidence that DSM has identified the soil landscape at the chosen scale. Since map users perceive patterns, and most land use decisions are for areas rather than pixels, we propose that DSM products be evaluated by their patterns, as well as by pointwise evaluation statistics.



soil scientists have been doing
this since the 1960s

"statistical"
"probabilistic" } more to the point
simpler language

1 Introduction

20 Digital Soil Mapping (DSM) is a general term for the creation of digital maps of soil classes or properties by fitting geostatistical or machine-learning models between observations of soil classes or properties at known locations and a set of environmental covariates representing soil-forming factors. Since its formal introduction by McBratney et al. (2003) it has been extensively applied worldwide at a wide range of scales and target classes and properties, see for examples reviews by Mulder et al. (2023), Arrouays et al. (2020) and Nenkam et al. (2024).

25 DSM products are routinely and (almost) exclusively evaluated by point-based evaluation statistics, and these are almost never based on probability or even representative observations (Piikki et al., 2021). Point-based evaluation ignores the existence of soil bodies that form a pattern over the landscape. Maps with distinctly different patterns of predicted soil properties or types can result from different approaches to DSM, see for example Rossiter et al. (2022) and Poggio et al. (2010). We propose to also evaluate DSM products by their patterns, as revealed by segmentation of the gridded maps into areas with more or
30 less homogeneous internal composition of soil properties. As Vayssé and Lagacherie (2017) aptly state, "DSM products are simplified representations of more complex and partially unknown patterns of soil variations" (emphasis added).

Soil geographers conceive of the soilscape as a continuum in 3D, with the vertical dimension (soil profile) defining a *pedon* (Soil Survey Staff, 1999, p. 11). The pedon has a horizontal dimension sufficient to show the local variability of horizons and properties, e.g., cyclic or irregular horizons. Pedons are connected laterally into relatively homogeneous *polypedons* (Johnson,
35 1963), within which the soil-forming factors and hence the pedons are within some defined limits. The transition zones between polypedons are marked as borders between natural soil bodies (according to those limits). Figure 1 shows a typical conceptual model from a detailed soil survey in the USA, design scale 1:12 000 (minimum mappable area 0.576 ha). The transitions between polypedons in this scene are due to parent material, topography, and hydrology. "order 2"

Together, these make up the soilscape, i.e., distribution of polypedons on the landscape. These form a pattern. The classic
40 example is the catena of Milne: "a sequence of distinct but pedogenetically-related soils that are consistently located on specific facets down a slope, giving recurrent topographically associated soil pattern" (Borden et al., 2020). We would hope that a map of a catena would clearly show these elements and their transitions.] often abrupt. *

In traditional, expert-based soil class mapping (Hudson, 1992) the landscape is segmented according to the mapper's conceptual model of soil-landscape relations, and by examination of external clues, notably relief, vegetation, and land use, and
45 by augering or full profile examination. DSM replaces the conceptual model with correlative relations with digital coverages meant to represent, at least in part, one or more of the seven predictive SCORPAN factors of McBratney et al. (2003). Therefore, there is no longer an explicit relation with the soil landscape, but it is hoped that the implicit correlative relations can find these. ← searching for connections vs. identifying connections

The concept of areas with distinct patterns of contrasting soils goes back to the "soilscape fabrics" from the soilscape
50 analysis of Hole (1978) and the "soil combinations" of Fridland (1974). At increasingly detailed scales and with increasingly fine distinctions in the definition of soil bodies, increasingly finer patterns are revealed. Conversely, at coarser scales, patterns are based on less precise definitions of distinct soil bodies. As Fridland puts it, "Soil combinations consist of elementary soil

Thesis
*

Yes, but
we recognize
the importance
+ freq.
of abrupt
changes

description
of soil morphology
& limited
characterization

can we just say soil-forming
factors + processes ← (proxies)




Figure 1. Conceptual block diagram, Otsego County NY (USA)

Source: <https://www.nrcs.usda.gov/publications/NY-2010-09-28-14.png>

areas which are genetically linked to various degrees and which produce a definite pattern in the soil mantle ... Multiple spatial repetition of a certain soil combination or several soil combinations alternating in a definite order creates various forms of structures of the soil mantle." An example of a fine-scale soil pattern is the pit and mound topography found on a hillslope in southwest Poland by Pawlik et al. (2024). *The repeating nature / or not / of these patterns = "mappable"*

In traditional soil mapping, these areas with sufficiently homogeneous soils or patterns of them at a given cartographic scale are the units that are delineated on the map. However, as Fridland explains: "The structure of the soil mantle and soil combinations are in their essence not cartographic but genetic-geographic concepts, even though they constitute a basis for elaborating cartographic units." This implies that the resulting soil properties distributed vertically in the profile, as products of pedogenesis, can be the basis for map units. Therefore, if at each pixel DSM accurately predicts a sufficiently rich set of properties over the soil profile, these should be grouped on the DSM map as recognisable cartographic units.

Within a mappable soilscape segment, there will of course be variability, ranging from some smaller deviations from a central concept (typical soilscape position and pedon), to a mixture of contrasting pedons, in National Resource Conservation Service (NRCS) soil survey terms a *complex*. Since DSM predicts per pixel, it may be possible to resolve these complexes into their components at the pixel scale, if that is fine enough to match the pattern within the complex. If this is the case, our evaluation of the DSM product should identify this.

useful to restate into mn kinds

*• assoc. • complex
• concave • undiff group*

This objective should be more clearly defined

"conceptual precision"



" ~~that~~ ^{This sounds awkward} " ^{grid} "fine": ~~smaller~~ ^{smaller than target MRA...?} ^{combine} Digital Soil Mapping (DSM) predicts at each pixel of a regular, ~~more or less fine~~ grid, either as the centre point or a block average of the area covered by the pixel. DSM typically predicts multiple soil properties at a set of standard depth slices.

70 Although some DSM methods use covariates in areas around a pixel, they do not enforce any relation between adjacent pixels. ^{!!! *} These relations are particularly important in soil hydrology models. Thus, the question is to what degree the pixels of DSM products at various resolutions can be aggregated into groups to realistically represent a soil landscape, whether the soilscape segment is relatively homogeneous in its properties or represents a complex. Intuitively, if the soil forming factors responsible for a polypedon are also spatially associated in the covariates used in DSM, the relations between pixels should occur as a ^{* reasonable proxies}
75 byproduct of per-pixel DSM. More abrupt transitions in the covariates should be reflected in the predictions. The pattern of the pixels should therefore represent the soil landscape. The question is, does the DSM product show these relations?

One way to assess the success of DSM in reproducing a soil landscape is to aggregate the individual predictions from pixels into more or less homogeneous *superpixels*, following methods used in image processing, where these are called *superpixels* (Nowosad and Stepinski, 2022). This can be based on single properties and depth layers, or, more usefully, on the multivariate
80 collection of DSM-predicted properties at a pixel. We explain the aggregation algorithm in §2.1.

At coarser scales, homogeneity of properties within some larger area may not be possible or even desirable. This has led to the concept of landscape *segments*, defined by the co-occurrence pattern, referred to as a *signature*, of a group of contrasting pixels of a class map, within a pre-defined size of the segment. Segmentation was developed by geographers to find similar land cover patterns for ecoregionalization (Nowosad and Stepinski, 2018). In that case, the pixels represent land cover classes.
85 The aim is not homogeneity of land cover, rather, homogeneity of the land cover pattern ^{soil mu concept!} within some analyst-defined area. The relation to a soil cover pattern is obvious, and corresponds well to concepts such as the catena ^(MU) or soil associations. Depending on the scale of the analysis and the inherent scale of the soil landscape, we may expect to see homogeneity at the level of *consociation* (e.g., ^{single} soil series), or a heterogeneous pattern of contrasting soils at the level of soil *association*, or a fine-scale pattern of contrasting soils, the *complex* (Soil Science Division Staff, 2017). ^{* "concept scale" 100m}

^{Integrate}
90 Segmentation requires that DSM maps of continuous predictions be classified, i.e., sliced according to analyst-defined class limits. The classes can correspond to meaningful classes for soil management, or can be based on laboratory precision. They can be wider (more general) or narrower, roughly corresponding to cartographic detail. Clearly, the classification can greatly influence segmentation. This is also the case when segmenting land cover classes. We explain the segmentation algorithm in §2.2.

95 Once a segmentation has been performed, the segments can be clustered according to their similarity of internal pattern, i.e., the signature of the segment. We explain the clustering algorithm in §2.3.

Thus, the objective of this study is present methods to create presumed soil landscape units from DSM products, by both aggregation and segmentation, and then to cluster the segments to identify similar soil landscape units within the map. We first describe the methods and then apply them to three case studies corresponding to different DSM projects at various resolutions
100 and extents. Finally, we discuss how these methods can be used in routine evaluation of DSM products.

^{This could use a more detailed definition + target "scale": a spatial + conceptual}



2 Methods

We contrast two approaches to letting the map “speak for itself”: aggregation based on homogeneity of properties (§2.1), and segmentation based on patterns of classified properties within segments (§2.2).

2.1 Aggregation

105 Aggregation seeks to find contiguous groups of pixels with relatively homogeneous property values, either single or multivariate. This is implemented by the `supercells` R package (Nowosad, 2025), which uses the Simple Linear Iterative Clustering (SLIC) image-processing algorithm (Nowosad and Stepinski, 2022), with the improvement that an appropriate data distance measure and function for cluster averaging can be defined. For multivariate aggregation there must be a distance measure defined in in multivariate space. A common choice, used here, is the Jensen-Shannon divergence, (Lin, 1991). which quantifies the distance between two histograms by the deviation between the Shannon entropy of the combination of two uni- or

110 multivariate histograms and the mean of their individual entropies.

The `supercells` function is controlled by several parameters that have a large effect on the results. First and most important is *compactness*, which trades off internal homogeneity of the supercells with their geometric compactness. The absolute compactness value depends on the range of input pixel values and the selected distance measure. A large value prioritizes spatial

115 distances between pixels and superpixel centres (more geometric compactness), whereas a smaller value prioritizes distances in feature space (more property homogeneity). Second is the approximate number of supercells, k . This should correspond to the number of landscape segments expected in the study area, at the design scale of the corresponding polygon map. Third is the minimum supercell size, *minarea*. This should correspond to a minimum mappable area or a minimum size needed for an application, e.g., land management or stratified sampling.

120 The quality of the aggregation can be evaluated by the standard deviation or coefficient of variability of each property in the supercell. As supercells decrease in size, these measures will necessarily have smaller values.

2.2 Segmentation

Segmentation seeks to find contiguous groups of blocks of grid cells with similar internal patterns of pixels, which represent soil classes or properties, these either univariate or multivariate. Patterns are computed within blocks of at least 10 x 10 pixels,

125 as specified by the analyst. Unlike supercells, segments must have rectilinear borders.

Segmentation proceeds as follows. The first step is to select classified soil properties and their depth slices to represent soil individuals at each pixel. The second step is to find the co-occurrence pattern of the pixels within pre-defined grid cells. The third step is to aggregate grid cells with similar internal spatial patterns into larger units, sufficiently distinct from neighbouring units in terms of their internal spatial patterns. Finally, the result is evaluated by its segmentation statistics, namely,

130 inhomogeneity within the segment and isolation of the segment from its neighbours. The segmentation can be inspected by expert judgement, perhaps comparing with conventional soil maps, to evaluate how well it represents the soil landscape at the selected cartographic scale.

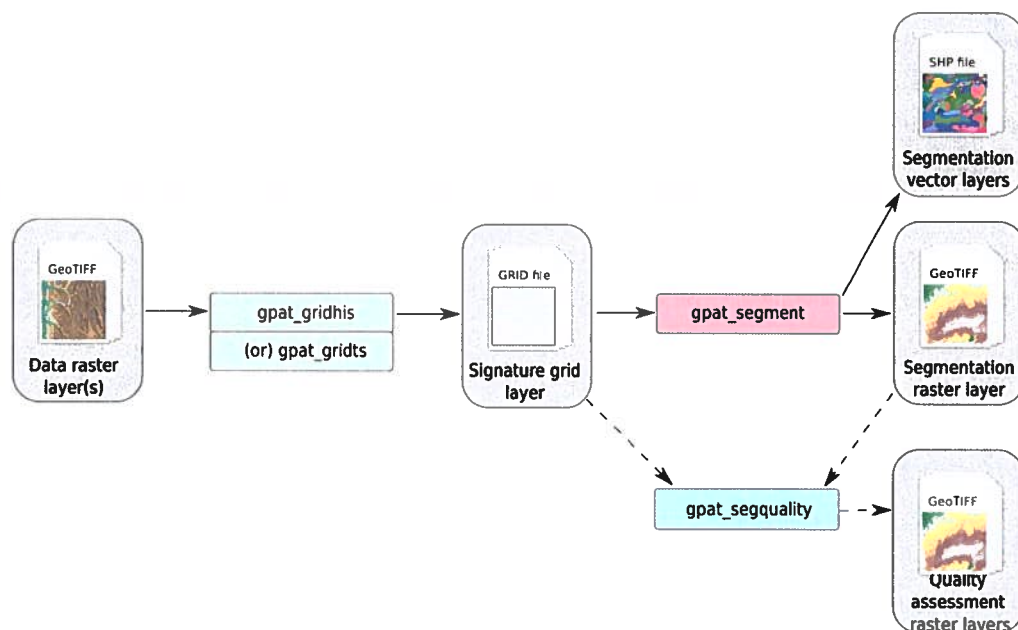


Figure 3.13: Workflow path for segmentation

Figure 2. GeoPAT segmentation workflow. Source: (Netzel et al., 2018)

For segmentation, we use the GeoPAT suite of standalone Unix programs (Jasiewicz et al., 2015). These are invoked in sequence, via the R system function, to obtain a segmentation and an evaluation of its quality. GeoPAT has been used successfully to segment categorical rasters such as land cover maps (Jasiewicz et al., 2018) and for global ecoregionalization based on multiple environmental factors (Nowosad and Stepinski, 2018). Figure (2) shows the segmentation workflow using GeoPAT.

Several parameters control the signature computation of the `gpat_gridhis` program. Two related parameters are `size` and `motifel`. The first is the size of the output grid cell of the segmented map. This must be at least 10 x 10 pixels of the source DSM. Thus, the segmentation is of similar patterns within an output grid cell and its neighbours. This dictates the largest equivalent map scale at which soilscape patterns (groups of output grid cells) can be discerned. The second is the “Motif Element”, referred to as the *motifel*, defined as the size of the window within which the pattern will be computed. This must be at least as large as the `size`, but could be larger to account for edge effects in the pattern. Also important are two thresholds for joining grid cells into segments: `lthreshold` to control the sizes of segments and `uthreshold` to prevent the growth of inhomogeneous segments.

Another important option for `gpat_gridhis` is the signature type within each grid cell, default `cooc`, “Spatial co-occurrence of categories”. This characterizes signatures with a “colour” co-occurrence histogram, a variant of the Gray-Level Co-occurrence Matrix (GLCM) used to characterise texture in greyscale images (Haralick et al., 1973; Hall-Beyer, 2017). In



Can any of this be related to existing ideas like survey order? ... Examples.

GeoPAT, discrete greyscale numbers, as in GLCM, are replaced by cell classes. A separation of one pixel is used to calculate the co-occurrence histogram, which then represents the spatial pattern within a grid cell. Related to this is the normalization type, default pdf “probability distribution function”, which is recommended for the cocc signature type. This harmonizes the signatures from different motifs.

Grid creation requires the selection of grid sizes. To evaluate DSM products we select these based on their correspondence to nominal map scales, using the Vink definition of a minimum legible delineation (MLD), i.e., the smallest area that can be displayed on a printed map, of 0.25 cm² at map scale, i.e., a grid cell side of 0.5 cm (Vink, 1963). To determine the Minimum Legible Area (MLA) and corresponding side on the ground, these are multiplied by the scale number (denominator of the scale ratio). For example, at 1:200 000 the MLA is 100 ha, with a side of 1 km. Signature computation requires at least 100 pixels from the DSM map in order to produce a reliable signature, i.e., the minimum edge of the segmentation grid (the “shift” parameter) must be 10 times the original DSM resolution. For example, a 25 x 25 m DSM product can only be segmented at 250 x 250 m or coarser (6.25 ha), corresponding to the MLA of a 1:50 000 scale map. To match a 1:200 000 map (MLA 100 ha), the 25 x 25 m pixel must be aggregated 40 times per side, i.e., 1 km x 1 km.

The segmentation phase in GeoPAT is implemented by the `gpat_segment` program. This groups grid cells based on their motif signatures computed by `gpat_gridhis`. Segments have a “brick” topology, in which square grid cells are arranged in alternating layers with each layer is shifted by one-half the size of the motif. Thus, the analysed area (i.e., the MLA) is four times the motif size.

Segment homogeneity is characterised by their normalised Shannon entropy H , defined as:

$$H = - \sum_{i=1}^{n_y} p_i \log_{n_z} p_i \quad (1)$$

where p_i is the proportion of the segment in class i , n_z is the number of possible classes, and these are summed over all n_y pixels in the grid cell. Using the logarithm to base n_z normalizes the entropy to the unit range regardless of the number of possible classes, so that 0 indicates complete homogeneity, i.e., one class for the entire segment. By contrast, 1 indicates maximum heterogeneity, i.e., all classes are equally represented in the segment. This only depends on class composition, not on pattern, even though the latter is the basis for segmentation.

Segmentation quality is measured with the `gpat_segquality` program. This produces two quality measures: (1) the inhomogeneity within each segment, and (2) the isolation of each segment from its neighbours. Inhomogeneity measures the degree of mutual dissimilarity between a segment’s motifs, on a [0...1] scale, where smaller values are better, i.e., more homogeneous, less internally diverse. Isolation is the average dissimilarity between a segment and its immediate neighbours, on a [0...1] scale, where larger values are better, i.e., more isolated. These measures depend on the pattern, not just the class composition, of segments.

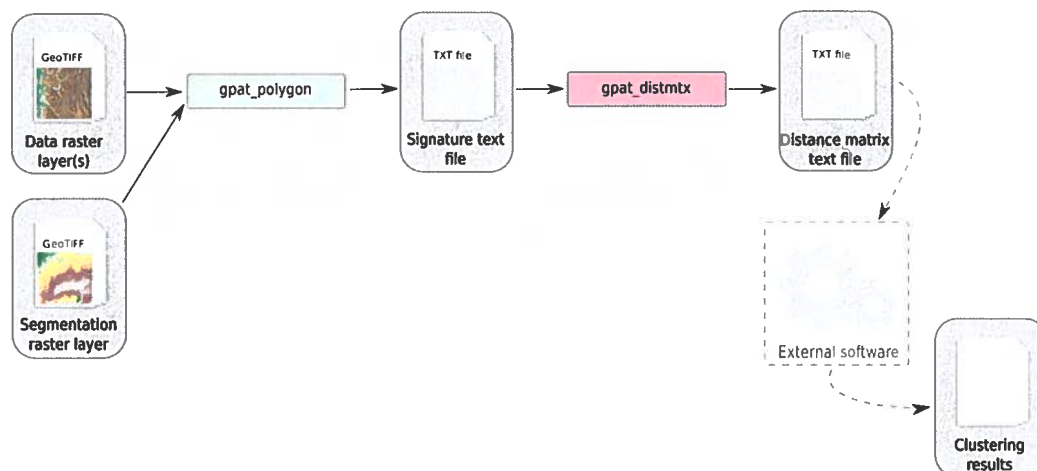


Figure 3.27: Workflow path for a clustering of segments (regions)

Figure 3. GeoPAT clustering workflow. Source: (Netzel et al., 2018)

2.3 Clustering

180 Once segments are created, their internal patterns can be characterised by the same signature methods used to perform the segmentation. Figure (3) shows the workflow for clustering in GeoPAT. The `gpat_polygons` program computes the signature within each segment. The distance between these signatures is then computed by the `gpat_distmtx` program. Here we used the default Jensen-Shannon divergence. The segments can then be clustered on the basis of their distance measures by many clustering algorithms. Here we use hierarchical clustering, as implemented by the R function `hclust` using Ward's linkage
 185 with squared distances to produce a dendrogram. This is cut at an analyst-determined number of classes to represent groups of internal homogeneity of segments. There are other choices in both the distance measurement and clustering linkage method. We chose Jensen-Shannon divergence because it is easily interpretable on a $[0 \dots 1]$ scale and is stable. We chose Ward's with squared distances to minimize within-cluster variance.

↑
 explain please

3 Case study 1 – BIS-4D (Netherlands)

190 BIS-4D (“Bodeninformatiesysteem 4-Dimensional”) (Helfenstein et al., 2024) is a high-resolution (25 m horizontal, six depth slices vertical) soil modelling and mapping platform for the Netherlands. The 3D are geographic space and depth along the soil profile. The fourth dimension is time, applied only to soil organic matter (SOM), which we ignore here by using only the most recent SOM map. Predicted properties are clay, silt, sand and SOM concentrations %, bulk density g cm^{-3} , pH in KCl, total N mg kg^{-1} , oxalate-extractable P mmol kg^{-1} , and cation exchange capacity mmol(c) kg^{-1} . Depth slices are
 195 the *GlobalSoilMap* standard 0–5, 5–15, 15–30, 30–60, 60–100 and 100–200 cm (Science Committee, 2015). Each map is




Figure 4. Semi-detailed soil map of the Netherlands, design scale 1:50 000 (part).

Source and detailed legend: Ministerie van Volkshuisvesting en Ruimtelijke Ordening (2024).

General legend: Dark and medium green: river clays with different clay concentrations; Light green: glacial depression sediments; Brown, pink: push moraines with varying sand and gravel sizes; Yellow: wind-blown sands; Purple: peat.

accompanied by uncertainties (quantiles and 90% prediction interval). We did not use these in this analysis, only the mean predictions. Coverages in the *GeoTIFF* format are free to download and use, and can be directly read into the `terra` R package (Hijmans et al., 2025).

BIS-4D is highly accurate at point support, as assessed by cross-validation, due to a very dense sampling network and the country-specific covariates used in the DSM. Visual inspection of layers agrees well with traditional 1:50 000 scale polygon soil maps (Steur and Heijink, 1980; Brouwer et al., 2021) and expert views of the soil landscape.

We selected a 40 x 40 km test area (Figure 4), because of its diverse soil-forming environments, including river clays of various ages and compositions, sandy push moraines, organic soils in glacial depressions, and coversands.



3.1 Aggregation

205 The `supercells` algorithm can work directly on raster stacks of the `terra` package. All 54 maps (7 properties, 6 layers) were combined in a `SpatRaster` raster stack. Since the values and ranges are not compatible, the Jensen-Shannon divergence was used to evaluate the distance in feature space between pixels and supercell centres. In this landscape there are non-compact (extended) features parallel to the river, in the fen areas and along the push moraines, so after some experimentation a low *compactness* value (0.2) was selected. We selected a minimum mappable area of 10 ha, equivalent to the 1:50 000 design scale
 210 of the Dutch conventional soil map, using the Cornell definition of 0.4 cm² minimum legible area on the map (Forbes et al., 1982). Thus we set the *minarea* was to 1,600 25 m x 25 m pixels.] *typo? missing comma?*

Figure 5 shows the supercells (outlined in black) with several properties as a background. Note that the supercells in all maps are the same, but of course the mean values of each property within the supercells are different. The median size of the 270 supercells was 433 ha, ranging from 104 to 5 044 ha, with a strongly right-skewed distribution. Aggregation clearly shows
 215 the differences between soil bodies, with some properties being more prominent in certain supercells.

To evaluate the quality of the aggregation, we computed the standard deviation of each property within each supercell (Figure 6). These are quite low for clay and SOM, and for pH with some small but notable exceptions. Bulk density is less successfully aggregated. The exceptions are where that property is not important in the computation of Jensen-Shannon divergence to that supercell.

Explain

220 3.2 Segmentation

Since `gpat_gridhis` requires class maps, we classified the soil property maps as follows: bulk density by 0.1 g cm⁻³, CEC by 25 mmol(c) kg⁻¹, clay, silt, sand concentrations by 5%, P_{ox} by 4 mmol kg⁻¹, pH by 0.1 units, SOM concentration by 4%, and total N by 1000 mg kg⁻¹.

The minimum grid size for segmentation (10 x 10 pixels) is 250 x 250 m (62.5 ha), corresponding to a 1:158 000 scale map,
 225 as explained in §2.2. Segmentation at this resolution is expected to more closely match the 1:200 000 generalised soil map (Haans, 1965) than the 1:50 000 semi-detailed map shown in Figure 3.

3.2.1 Univariate segmentation of individual maps

To examine the effect of grid size, we segmented all properties at all depths, individually, at the minimum possible grid cell size, i.e., 10 x 10 and at several multiples: 40 x 40 (1 000 ha) and 80 x 80 (4 000 ha), corresponding to nominal map scales
 230 1:400,000 and 1:800,000, respectively. *commas?*

The finest segmentation produced 4,393 (pH 100–200 cm) to 675 (P_{ox} 100–200 cm), median 2,678 segments, average area 0.597 km². Comparing this to the single grid cell at resolution, 0.625 km², we see that many segments were of one or two grid cells. The pattern was mostly very fine, with a few large segments for most single properties.

Segmentation at 1:400,000 equivalent scale produced 231 (sand 0–15 cm) to 41 (SOM 15–30 cm), median 181 segments,
 235 average area 8.84 km². Compared to the single grid cell at resolution, i.e., 1 km², there was significant grouping. Segmenta-

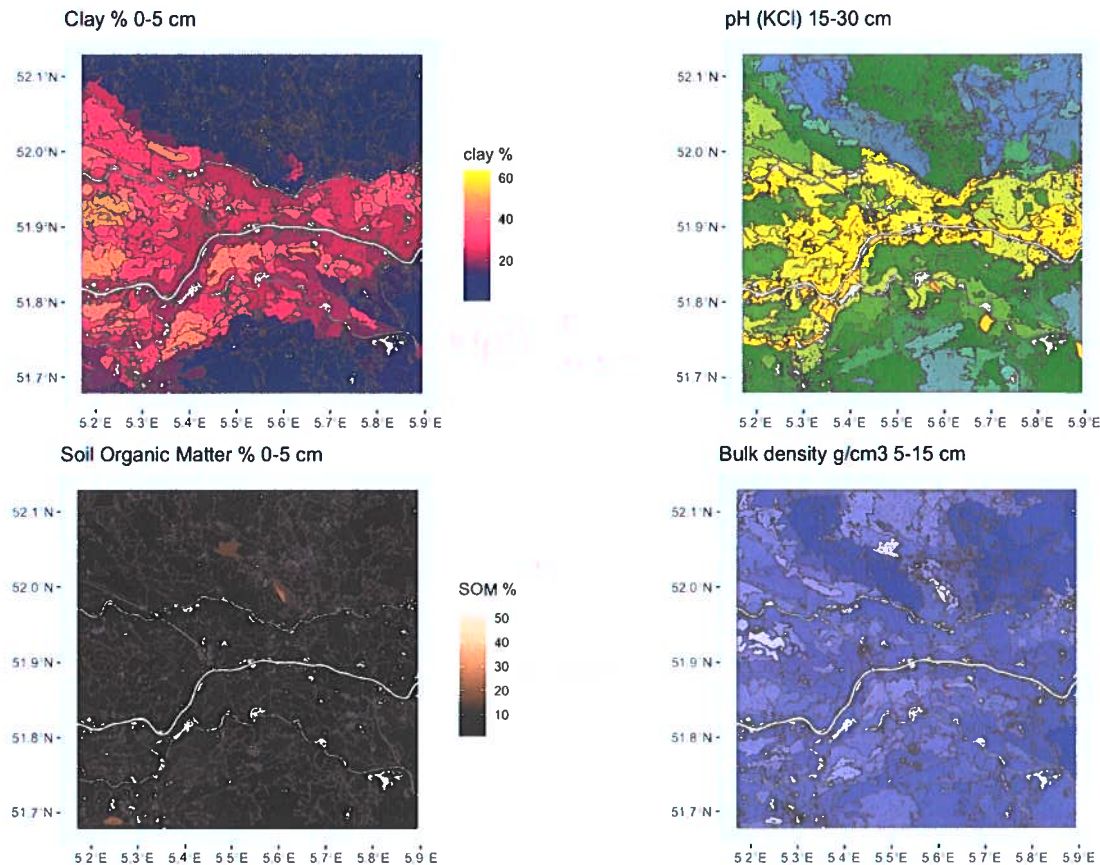


Figure 5. Results for selected properties of aggregation by supercells algorithm using all properties and layers

tion at 1:800 000 equivalent scale produced 66 (sand 0–15 cm) to 12 (Pox 60–100 cm), median 47 segments, average area 34.04 km². Again, compared to the single grid cell at resolution, i.e., 4 km², there was significant grouping.

3.2.2 Multivariate segmentation of individual properties, all depth slices

We then performed a multivariate segmentation using all depth slices of single properties. By default, GeoPAT normalizes each layer and by default weights them equally. In this mode, a motif must meet the threshold conditions for all input layers to be joined to a segment. In this way the segmentation is meaningful for each layer. Because of the different spatial structures of the properties at each depth slice, it was expected that the segmentation would be finer at each scale than for individual depth slices, i.e., it would be more difficult to merge grid cells.

The finest segmentation using all depth slices of a single property produced 3 316 (pH) to 168 (SOM) segments, median 1 873 segments, average area 0.854 km². Segmentation at 1:400 000 equivalent scale produced 190 (sand) to 13 (SOM) segments, median 127 segments, average area 12.59 km². Segmentation at 1:800 000 equivalent scale produced 55 (sand) to 6 (SOM)

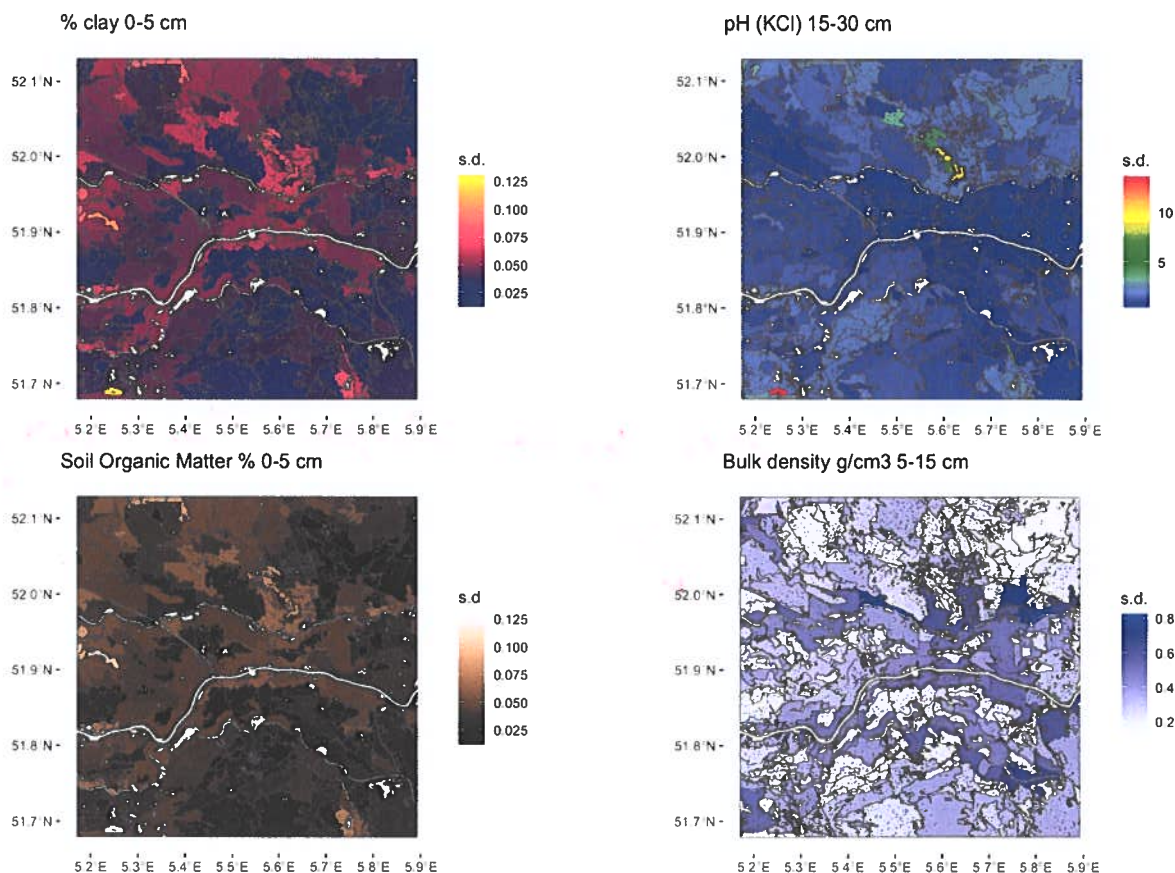


Figure 6. Standard deviations for selected properties of aggregation by supercells algorithm using all properties and layers

segments, median 36 segments, average area 44.44 km². Contrary to our expectations, the median number of segments were all smaller than those for the corresponding property's single depth slice segmentations.

Figure 7 shows the segment boundaries for this multivariate segmentation by bulk density over the whole profile, at the three resolutions overlaid on the Dutch 1:50 000 soil survey polygons. It is clear that the 1:800 000 segmentation misses important differences and that the 1:100 000 segmentation finds quite small areas, mostly just one grid cell, within soil bodies. The 1:400 000 segmentation (i.e., shift size 40, 1 km²) grid cells) matches well with many soil map boundaries.

Figure 8 shows the success of the segmentation based on bulk density over the whole profile at the 1:400 000 design scale: inhomogeneity of each segment and isolation from its neighbours. For example, the pixels in the large segment in the top-centre are quite similar in their bulk density profiles, but this segment is only moderately different from its neighbours. This shows the relative homogeneity of the bulk density profiles of the central Gelderse Vallei (Gelderland Valley) in the vicinity of Renswoude and Scherpenzeel. Note that this area also has large segments based on all properties and depth slices, as seen in Figure 5.



Can these be larger? Hard to "see" patterns



Figure 7. Segmentation based on bulk density over the whole profile (red lines), overlaid on soil map polygons (grey lines). Design scales left to right: 1:100 000, 1:400 000, 1:800 000

How can we interpret these figures?

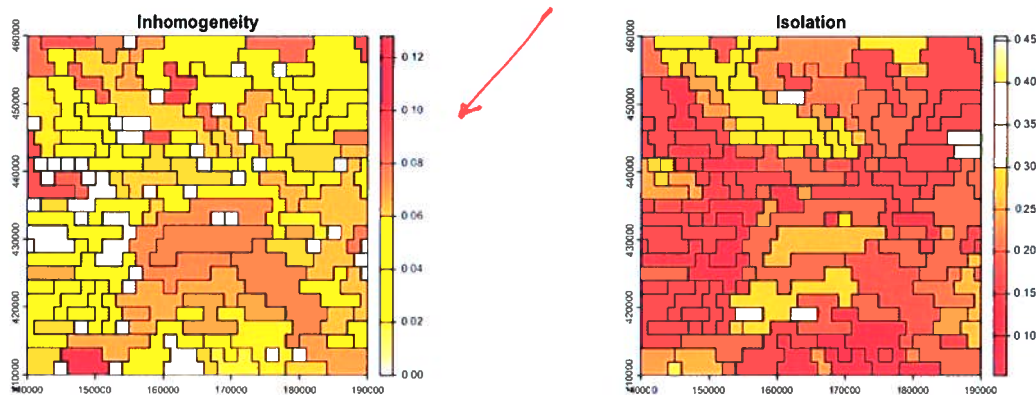


Figure 8. Evaluation of segmentation based on bulk density over the whole profile at the 1:400 000 design scale

As the segmentation becomes coarser the inhomogeneity and isolation both decrease, i.e., segments are internally more consistent in their patterns, and less isolated from their neighbours. For example, median inhomogeneity values from the segmentation based on whole-profile bulk density (1 266, 96, 28 segments) decreased from 0.108, 0.086, to 0.076. In parallel, median isolation values decreased from 0.288, 0.211, to 0.178.

3.2.3 Multivariate segmentation with selected properties and depth slices

Another segmentation is obtained by selecting properties and depth slices to represent the profile. Using all 56 layers results in an impractical Jensen-Shannon divergence, hence we selected key properties at key depths: pH, clay, silt, SOM 0-5 cm, clay, bulk density 15-30 cm, CEC 30-60 cm, sand, SOM 100-200. Figure 9 shows the segment boundaries from this segmentation at the 1:400 000 design scale, overlaid on several single soil properties and depth slices. Note that the segment boundaries are the same for all maps. This segmentation should best group soils considered holistically, not per-property. Many of the segments

Why?

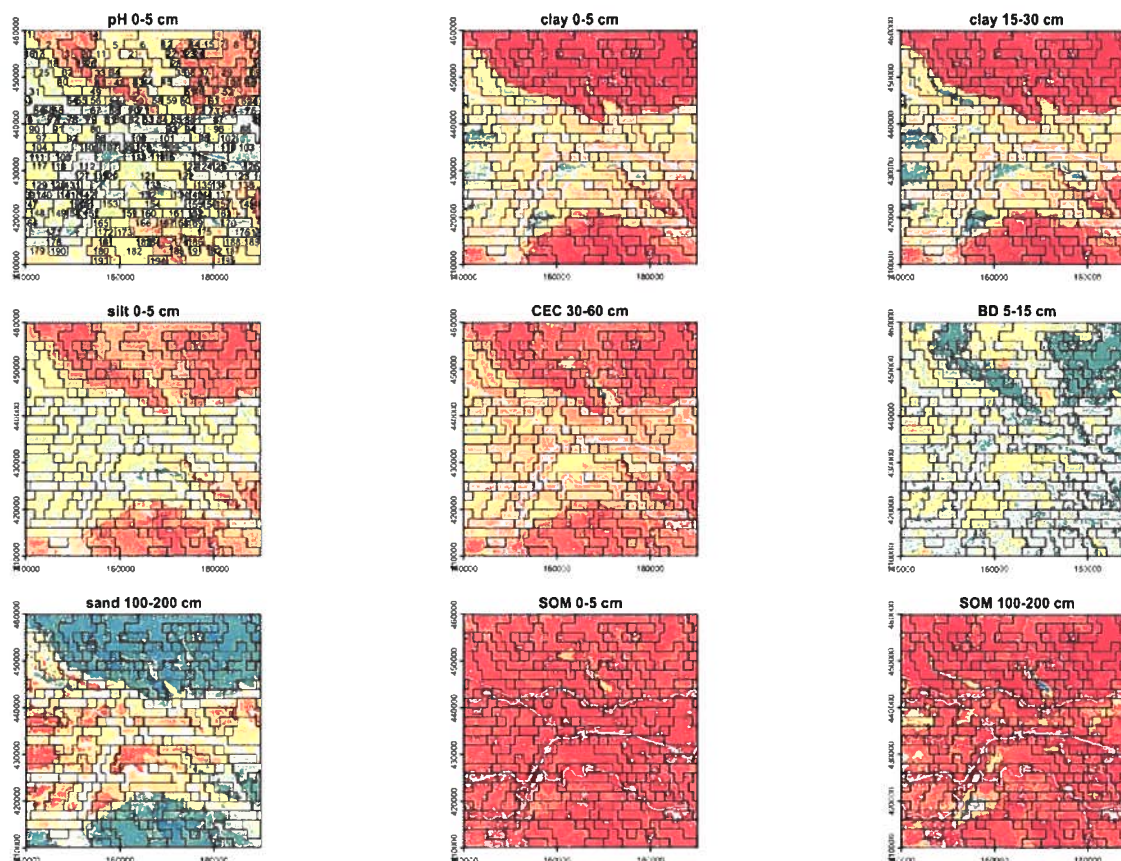


Figure 9. Segmentation based on selected properties and depth slices, overlaid on DSM of selected soil properties, 1:400 000 design scale. Legends not shown. Scale is from dark red (low values of the property) to dark green (high values). Top-left map includes segment numbers

correspond to landscape features shown in the conventional soil map of Figure 4, although constrained to the rectilinear shape and minimum grid cell size.

3.2.4 Scaling of segmentation

meaning? efficient? / "generalizes well"?

The segmentation method scales well. The land area of the Netherlands ($\approx 33\,240\text{ km}^2$) was segmented using all depth slices for several properties. At the nominal 1:400 000 design scale, this resulted in 2 535 (pH) and 1 547 (bulk density) segments; at 1:800 000 design scale 649 (pH) and 371 (bulk density). Figure 10 shows the segmentation by pH of the entire Netherlands at these two scales. For this extent the coarsest segmentation seems most useful for understanding the country-wide soil pattern.

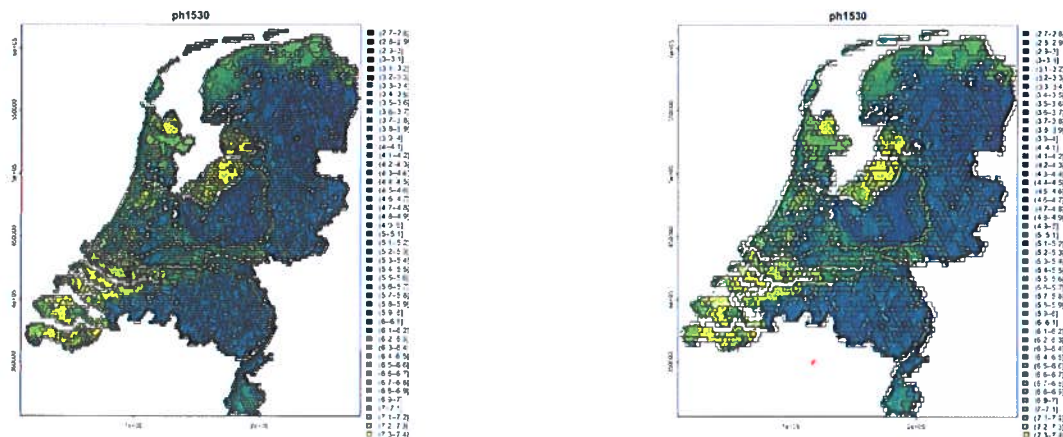


Figure 10. Segmentation by whole-profile pH of the Netherlands at 1:400 000 (left) and 1:800 000 (right) nominal scales, overlaid on the pH 15–30 cm DSM product

3.2.5 Segmentation parameters

Segmentation is greatly affected by the two thresholds. For example, segmenting the test area using all depth slices for clay using the default lower and upper thresholds (0.1 and 0.3, respectively) results in 1 932 (1:100 000) and 148 (1:400 000) segments, whereas using more liberal (easier segmentation) thresholds 0.3 and 0.8 the number of segments is reduced to 285 and 18. In effect, the more liberal segmentation at a finer scale is similar to the more conservative one at a coarser scale. Figure 11 shows the multivariate segmentation of the test area on the basis of clay concentration at all depth slices at nominal 1:400 000 scale with default thresholds, and the same for the 1:100 000 scale but with liberal thresholds. These maps are comparable.

3.3 Clustering

helpful interpretation

Hierarchical clustering was applied to the segments of Figure 9, i.e., based on selected properties and depth slices, to represent the profile. The resulting dendrogram is shown in Figure 12. Note the large separation in internal patterns between the two top-level branches (height 6). These represent the river clay landscape, Gelderse Vallei depression, and lower terraces (right branch, clusters 4–7) and the sandy uplands (left, clusters 1–3). At the second level for the right branch (height 3.5) the large separation is between the Gelderse Vallei depression and terraces (clusters 4 and 5) and the river clays (clusters 6–7). At the third level for the rightmost branch is the separation between the actively flooded zones (cluster 7) and the somewhat higher zones (cluster 6). While not a perfect separation, the clustering does separate the principal soil landscape components.

The seven generalised clusters identified in the dendrogram are shown on the landscape in Figure 13. These group similar segments well and could serve as landscape management units. For example, cluster 4 groups the mostly homogeneous seg-



*black lines hard to "see"
 consider alt. color ramp + lines*




Figure 11. Segmentation by whole-profile clay at 1:400 000 with default thresholds (left) and 1:100 000 (right) with liberal thresholds, overlaid on the clay 0–5 cm DSM product

worth comparing to multivariate segmentation

ments dominated by low pH, clay, SOM, CEC, high sand, and medium silt. Cluster 7 groups the heterogeneous segments along the rivers and large brooks.

3.4 Evaluation

The BIS-4D product can “speak for itself” quite well, to reveal both compact units of homogeneous soils and segments with similar heterogeneous patterns of soil classes. Aggregation based on properties and depths selected to represent the results of the principal soil forming factors delineates patches (Figures 5 and 9) that closely correspond to polygons of the 1:50 000 design scale conventional soil map with design scale 1:50 000 (Figure 4), generalized to about 1:158 000 design scale, although with some variations in form. Segmentation was most successful with grid cells of 1 000 ha, corresponding to nominal map scale 1:400 000. This grouped patterns of pixels with different internal patterns of classes. Hierarchical clustering of these segments found groups of similar patterns within the map. These represent separate segments of the same landscape component. These results increase confidence in the BIS-4D DSM product. This is perhaps a best case, due to the extremely high quality of the source data (training points and covariates), the conventional map which can be used for comparison with aggregation and segmentation, and sophisticated modelling approach specific to the Netherlands.

4 Case study 2 – SoilGrids v2.0 (Global)

At the other extreme from the country-specific DSM exercise based on a large quality-controlled and spatially complete training set (§3) is a global DSM exercise based on a heterogeneous and spatially-unbalanced training points, using only covariates with global coverage. For this case we selected SoilGrids v2.0 (Poggio et al., 2021) from ISRIC-World Soil Information. This is a



* suggestions:

- horizontal layout w/ arc package
- use height scale + annotations to follow text in 3.3
- ~~label~~ label clusters w/ numbers
- if not referencing segment ID, do not include
- use shaded symbols / filled symbols vs. colored outlines
→ place filled symbols @ terminal "tips" of dendrogram

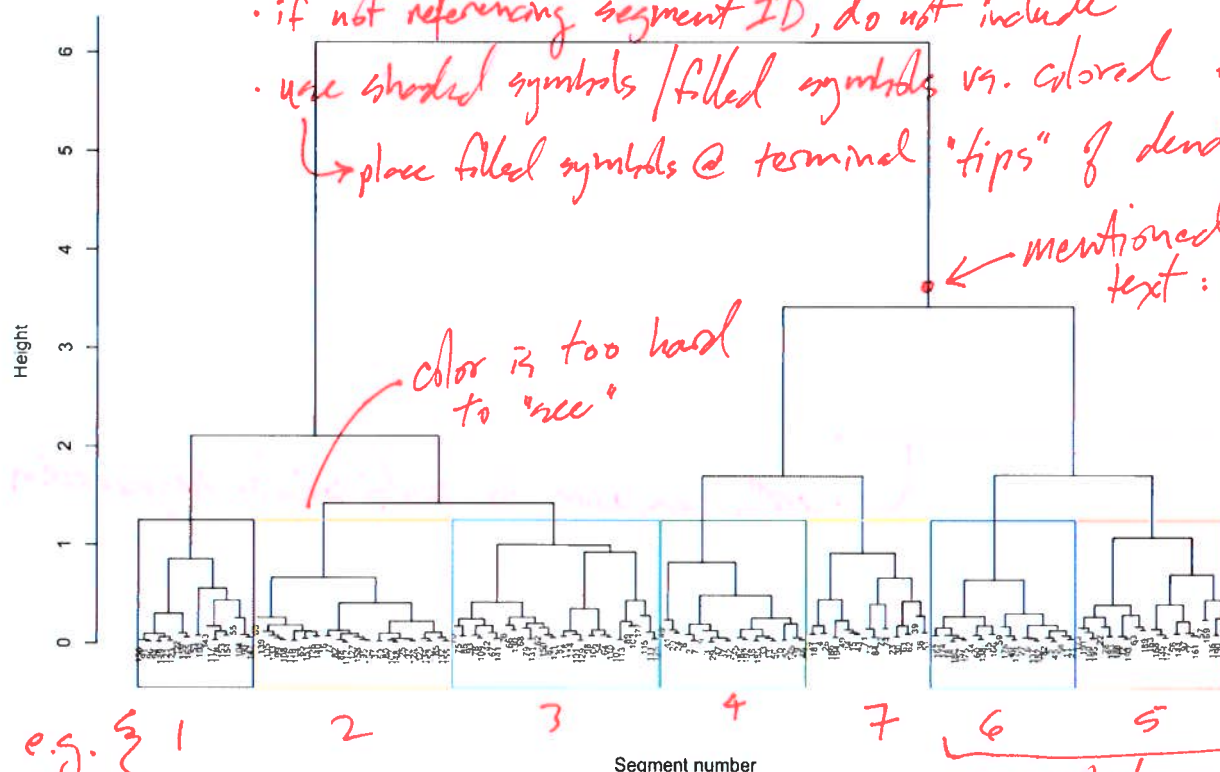


Figure 12. Hierarchical clustering of the segments shown in Figure 9

set of predictive maps of soil properties for the entire globe at 250 m nominal spatial resolution. Aggregations to 1 km and 5 km resolutions are provided for modelling at coarser scales. It is a globally-consistent product that uses all available point data from the World Soil Information Service (WoSIS) database (Batjes et al., 2024), also from ISRIC-World Soil Information, and covariates with global coverage. Political boundaries are nowhere visible, except where one or more covariates match these. ... *irrelevant*

315 SoilGrids provides both predictions and their uncertainty, via quantile random forest ~~machine-learning models~~. It closely follows the *GlobalSoilMap* specifications of properties and depth slices (Science Committee, 2015). It also predicts the derived property of SOC stocks from 0-30 cm, in $T\ ha^{-1}$, computed from SOC concentration and bulk density. We chose to evaluate this layer, in order to compare it with the FAO's Global Soil Organic Carbon Map (GSOCmap) project (FAO, 2018).

320 We selected a transnational study area with corners (-109.99, 27.86) E and (-100.03, 35.64) N. This covers most of Chihuahua and Coahuila and part of Sonora States (MX) and portions of Texas and New Mexico States (USA). Figure 14 shows this area,



Are all segmentations the same? ?
How can we better "see" the spatial structure of clustering (segments)? → fill clusters/segments w/ color, no



*outline,
 contour
 soil property
 over data
 overlay...?*

Figure 13. Generalised clusters of the segmentation of Figure 9, based on slicing the clustering dendrogram shown in Figure 12 for seven general clusters.

with the SOC stocks over the 0–30 cm depth slice. The higher stocks are in mountains and wetlands along the Rio Grande, the lower in high deserts.

Individual $2 \times 2^\circ$ tiles of the 250 m product were downloaded in the GeoTIFF format from the interactive SoilGrids site (ISRIC-World Soil Information, 2024b), imported into R with the `terra` package, mosaicked, projected from the original geographic coordinates to a local Albers Equal Area projection, and trimmed to $3\,270 \times 3\,610 \times 6.25 \text{ km}^2$ pixels, covering 737 793.8 km². The global map of the 1 km product was downloaded in the GeoTIFF format from the ISRIC WebDAV repository (ISRIC-World Soil Information, 2024a), projected from the original Homolosine coordinate reference system to the same local Albers Equal Area projection, and trimmed to $900 \times 900 \times 1 \text{ km}^2$ pixels, covering 810 000 km².

Predicted SOC stocks per pixel ranged from 0 to 83, median 28 T ha⁻¹ for the 250 m product, and 7 to 76, median 29 28 T ha⁻¹ for the 1 km product, showing the smoothing effect of upscaling.



* why only SOC stock?
 seems like a huge scale m3-metres
 units

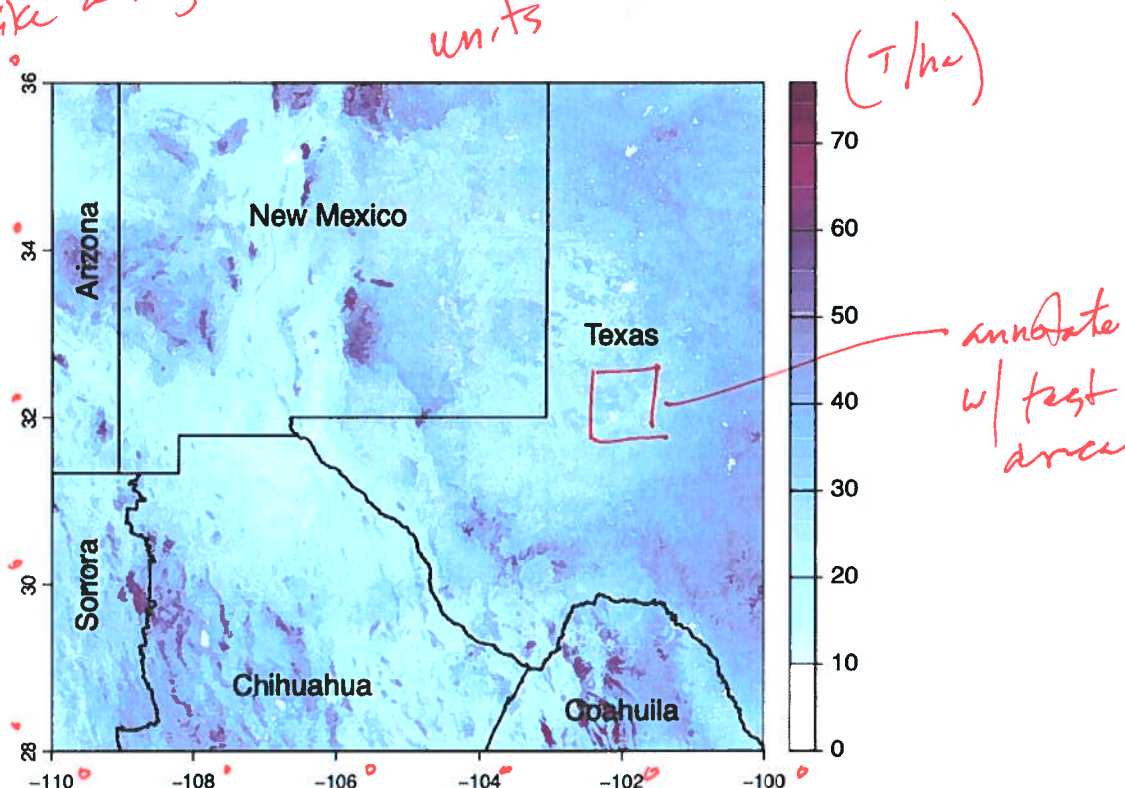


Figure 14. SoilGrids v2.0; SOC stock 0-30 cm, T ha⁻¹.

was this entire extent used?
 where is the test area?

4.1 Aggregation

We applied the `supercells` algorithm to the SOC stocks 250 m resolution layer. To limit processing time and memory requirements, we selected a small test area of 80 x 80 km, i.e., 640 000 ha, centred on (-105 E, 32 N) at the Texas (N) / New Mexico (S) border, near Dell City NM (Figure 15). The centre pivot irrigated fields at the centre-left are $\approx 800 \times 800$ m and should thus be resolvable on the SoilGrids map. This area includes a wide range of the SOC stocks (Figure 16 left), with high values in the Guadalupe Mountains to the east and very low values in the salt flats in the centre of the area.

After some experimentation, a medium value (0.5) for *compactness* was selected. We did not set a minimum mappable area *mineara*, rather a number of proposed supercells *k*. A choice of ≈ 400 supercells corresponds to an average area of 1 600 ha, corresponding to 1 cm² on a 1:400 000 printed map. This is much larger than the area of single centre-pivot irrigated fields, so we did not expect these to be individually resolved.




Figure 15. Test area for aggregation, centred on (-105 E, 32 N). Source: © Google Earth

Figure 16 (right) shows the computed supercells. Median size of the 412 supercells was 1 388 ha, ranging from 431 to 5 462 ha, with a strongly right-skewed distribution. This aggregation clearly groups the pixels with similar SOC concentrations. However, the shapes do not seem to correspond to natural landscape boundaries. We attempted other combinations of compactness and supercell numbers, with poorer results.

345 The quality of the aggregation can be measured by the standard deviation of the property within each supercell (Figure 17). These ranged from 0.34 to 6.08, median 1.18 $T\ ha^{-1}$, with corresponding coefficients of variation from 1.36 to 26.61, median 4.39%. The highest heterogeneity was in the pivot irrigation area, where the minimum supercell size forced pixels with a wide range of values together.

4.2 Segmentation

350 Segmentation was applied to the SOC stock map of the full study area, for both resolution SoilGrids DSM products. Since `gpat_gridhis` requires class maps, SOC stocks were classified in 19 (250 m) and 18 (1 km) equal intervals of 4 $T\ ha^{-1}$, with from 31 to 1'956 813 (250 m) and 14 to 128 549 (1 km) pixels per class. The minimum grid resolution for the 250 m product is here 2.5 × 2.5 km. The map was segmented at this resolution, and also four coarser resolutions: 5 × 5 km, 10 × 10 km,

→ could this be an artefact of spectral correlative used in by SoilGrids?

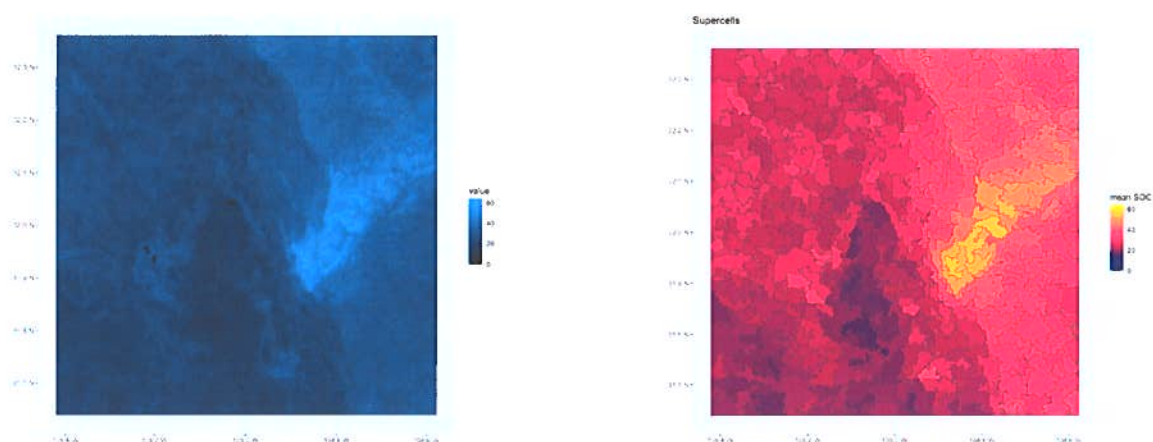


Figure 16. SoilGrids v2.0 SOC stock (left) and its aggregation into supercells (right); SOC stock 0-30 cm, T ha^{-1}




Figure 17. Standard deviation within supercells; SOC stock 0-30 cm, T ha^{-1}




Figure 18. Segmentation of 250 m resolution SoilGrids map (part) at 1:1M nominal resolutions (SOC)

** would ~~see~~ one expect meaningful generalization of the soil landscape based only on SOC stock ... I doubt it*

20 × 20 km, and 40 × 40 km, corresponding to map scales 1:1M, 1:2M, 1:4M, 1:8M, and 1:16M, respectively. These produced

355 7 600, 1 905, 491, 127, and 35 segments from the 250 m resolution map, respectively. Figure 18 shows the segmentation at the finest scale. The level of detail is apparent, but many segments seem to be of a single class, with no internal pattern. Broader landscape patterns are obscured by this level of detail.

From the 1 km resolution SoilGrids map the three coarsest resolutions resulted in 669, 165, and 43 segments. Figure 19 shows these three segmentations. As resolution decreases, broader landscape patterns are increasingly apparent. All segmentations
360 seem useful at their respective design scales.] *Explain*

For the 250 m SoilGrids segmentation, median standard deviation increased from 2.35, 3.08, 3.94, 4.62, to 6.15 T ha⁻¹, while the median normalized Shannon entropy increased from 0.311, 0.369, 0.433, 0.472, to 0.580, for the 1:1M, 1:2M, 1:4M, 1:8M, and 1:16M scales, respectively. Entropy and standard deviation increase with segment size, as expected. The comparable values for the 1 km SoilGrids segmentation are median standard deviation 3.37, 4.29, and 6.19 T ha⁻¹, and median normalized



can't read, can use single (larger) legend





Figure 19. Segmentation of 1 km resolution SoilGrids map (part) at (left to right) 1:4M, 1:8M, and 1:16M nominal resolutions



colors are spc stack?

Figure 20. Normalized Shannon Entropy of segments of the SoilGrids v.2 1 km map (part) at 1:16M nominal resolution.

365 Shannon entropy 0.417, 0.481 and 0.584 for the 1:4M, 1:8M, and 1:16M scales, respectively, similar to those from the 250 m segmentation.

Figure 20 shows the entropy for each segment of the 1:16M nominal resolution map from the 250 m product. This is a measure of the internal class homogeneity of each segment, although not the spatial pattern of the classes. The highest entropies are found in the segments with mixed high and low terrain, shown as contrasting purple and light blue colours.

370 4.3 Clustering

Figure 21 (left) shows the 39 segments signatures from the 1 km product, using motif size 40 cells, and Figure 21 (right) shows the assignment to seven generalised clusters. Figure 22 shows the dendrogram for the clustering of the 39 segment signatures.



make maps larger
use single legend

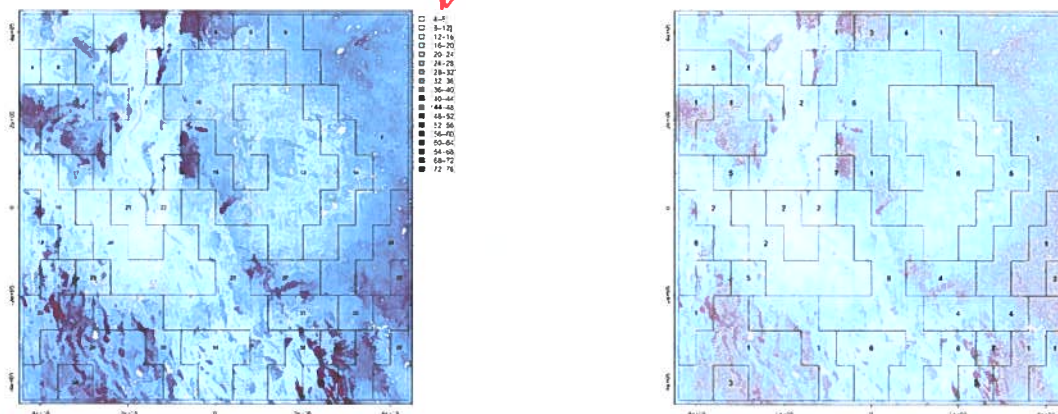


Figure 21. Left: Segmentation of SoilGrids v2.0 1 km SOC stocks (T ha^{-1}), motif size 40 cells; Right: Assignment of segments to seven generalised clusters

** same comments as Figure 12.*

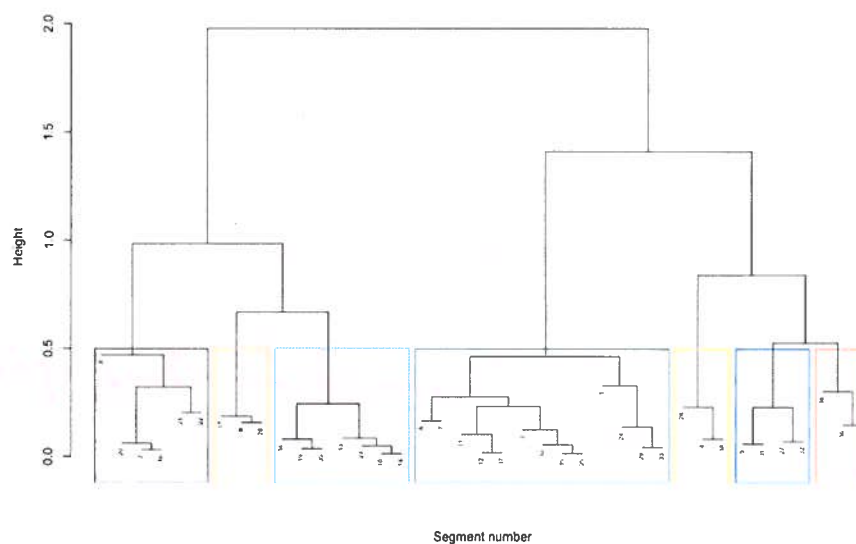


Figure 22. Dendrogram of segment signatures, SoilGrids v2.0 1 km, motif size 40 cells.

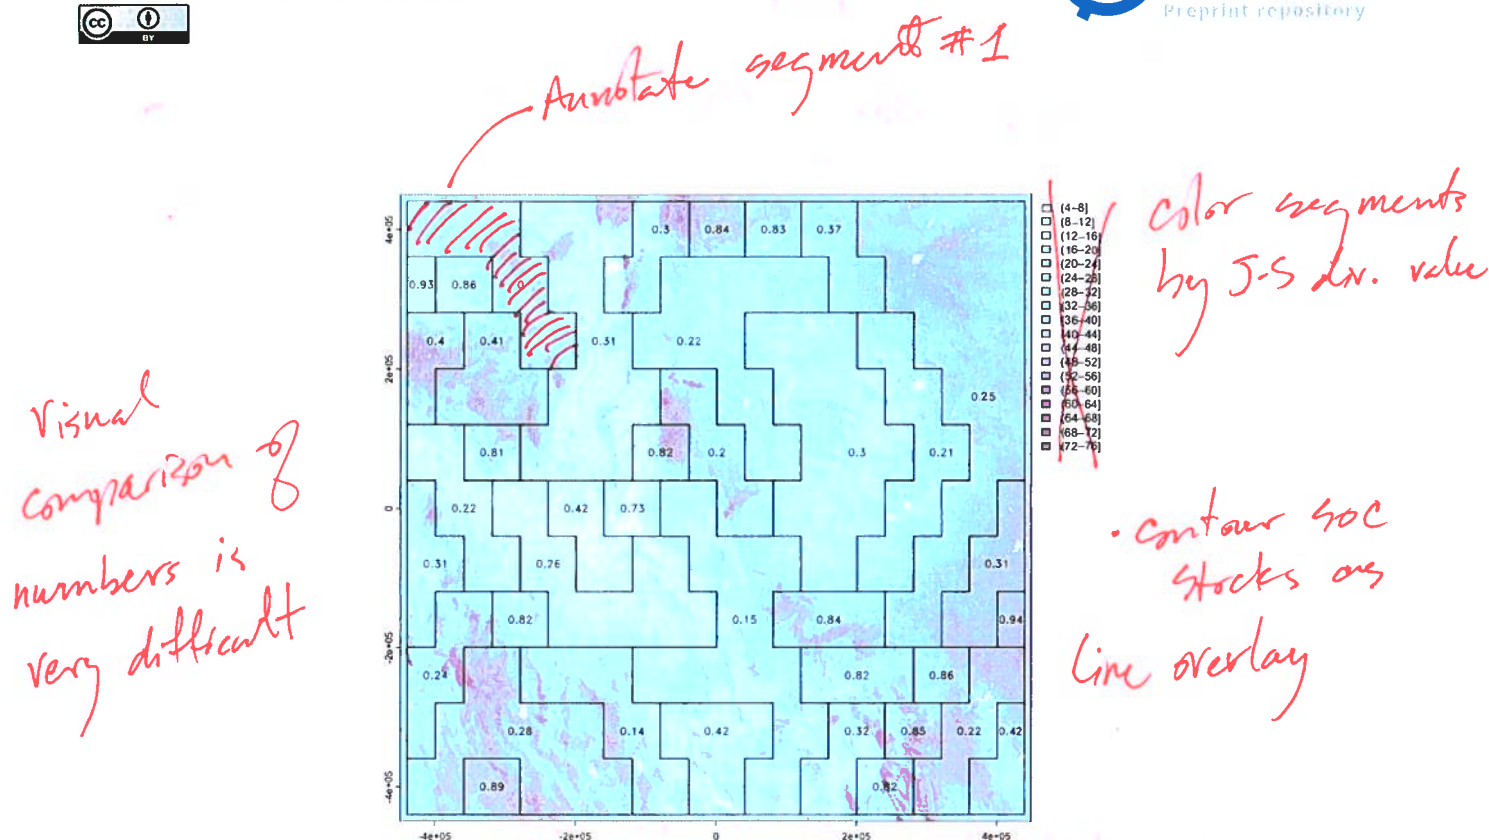


Figure 23. Jensen-Shannon divergence from Segment 1

The co-occurrence pattern of classes is similar within each general cluster. The clusters should group similar soil landscapes, at least with respect to the SOC concentration. For example, cluster 1 groups mountainous terrain with high SOC interspersed with basins with medium SOC in an intricate pattern.

Figure 23 shows the Jensen-Shannon divergence with the first segment, which necessarily has no divergence. These range from 0.14 (segment 30, in the same cluster 1 as the target segment, although on a different first branch at height 0.45) to 0.84 (segment 4, in widely-separated cluster 3, different at branch height 1.45). This can be used to find the soil patterns that are most similar to any segment, independently of cluster membership. The distance does not directly correspond to cluster distance in the dendrogram when linkages other than single are used, as in this case, Ward's D2.

4.4 Evaluation

Aggregation was able to form compact groups of pixels with similar SOC stocks, which could be useful for, e.g., stratified sampling. However the polygons did not seem to correspond well with landscape units. Segmentation was more successful. At several increasingly-general scales it grouped distinctive patterns of SOC stocks, corresponding to large landscape units. Clustering was then able to identify general groups of landscape units, and the Jensen-Shannon divergence identified the segments most similar to a selected segment.

annotate these points on Fig. 22

** Does the map generalize to recognized units?*



5 Case study 3 – SOLUS100 (USA)

The third case study is intermediate to the first two. Like the BIS-4D study it is of one country and with training points from one source, but (1) it covers a much wider area and so can't use covariates that are only available for part of the area, and (2) the product is based on numerous traditional soil surveys of varying age and quality control which can be used to some extent for evaluation.

SOLUS100 ("Soil Landscapes of the United States 100-meter") is a recent DSM product from the USDA-NRCS (Nauman et al., 2024). This contains predicted values, high and low estimates, and prediction intervals for soil properties at the *Global-SoilMap* standard depths, at 100 m horizontal resolution (i.e., 1 ha pixels) over the entire conterminous United States (CONUS). The maps are available in GeoTIFF format (Nauman, 2024). These can be compared to the Gridded Soil Survey Geographic Database (gSSURGO) digital product from the NRCS (NRCS Soils, 2022). This was created by digitising the polygons from traditional soil-landscape survey, with its linked relational database of polygons, map units, components, horizons, and soil properties. Thus aggregation and segmentation can be compared to a product based on expert judgement and field-based soil survey, although gSSURGO is also quite heterogeneous in the age and quality of the soil surveys on which it is based, and so must be used with caution as a ground truth.

We selected a 570 km² test area in Wayne County NY, mapped in 1978 on an unrectified airphoto base (Higgins, 1978), and later digitised by the NRCS and incorporated into gSSURGO. This area has a distinctive pattern of NNW-SSE orientated drumlins of various sizes and shapes, and inter-drumlin depressions. Some of these developed into peatlands, with drained areas used for agriculture and undrained areas used as wildlife reserves. The genesis of this soil landscape has been studied for more than a century (Menzies et al., 2016).

Figure 24 shows the predicted surface layer clay concentration for the original soil survey, as compiled in gSSURGO, and for SOLUS. Notice the different legend scales, otherwise the SOLUS map would not clearly show its pattern, since SOLUS predicts a narrower range of concentrations, as is typical of DSM products. It is obvious by visual inspection that SOLUS misses much of the fine pattern, and especially that it does not identify most of the organic soils with very low clay concentrations (dark blue on the gSSURGO map).

5.1 Aggregation

We aggregated the SOLUS map of surface layer clay concentration with the supercells algorithm. We set the minimum area parameter minarea to be comparable to MLD at original design scale. The source map in this area was at 1:24k design scale, so the MLD was set to 2.304 ha. The Optimal Legible Delineation (OLD) is 4 x MLD (Forbes et al., 1982), so here 9.216 ha, corresponding to nine SOLUS cells. Aggregation complexity is controlled by the number of supercells. This should be comparable to the number of gSSURGO polygons in this study area. In this way we can evaluate how well the DSM can match the traditional soil survey. In this area there are 14,949 gSSURGO polygons, with a median area of 2.43 ha, corresponding to 2 to 3 cells. The mean is area 5.30 ha, because of some large polygons, mainly organic soils, i.e. Histosols in US Soil Taxonomy (Soil Survey Staff, 1999).



- larger maps
 - less saturated color ramp
 - use the same CRS ...?

- histogram of each grid would be useful



Figure 24. Clay concentration % of the 0-5 cm layer, gSSURGO (left), SOLUS 100 m (right). Coordinate Reference System is an Albers Equal Area for CONUS.

use same grid

We aggregated with a range of compactness values from 0.2 to 2. Because of the long linear shape of the drumlins, we expected that the lower compactness would best match the landscape. Indeed, this parameter value produced the map with the least rounded features, but their orientation did not match the landscape pattern (Figure 25). From this we conclude that SOLUS in no way represents the actual soil pattern. This same result was obtained with other layers of clay concentration, and with several other soil properties.

what does it appear to represent?
 elevation?

5.2 Segmentation

SOLUS resolution is 100 m, so that the minimum shift is 10 i.e., 1 000 m = 1 km, corresponding to 1:250k nominal scale. Thus we did not expect to reproduce the fine pattern, but rather to group these into regions. We segmented with raster stacks of single properties at all depth slices, and with a raster stack of seven properties (clay, silt, and soil organic carbon weight concentrations, coarse fragments volume, pH measured at 1:1 in water, CEC, bulk density) at one depth slice. The continuous properties were converted to classes, as required by the GeoPAT segmentation algorithm: particle-size separates in units of 4%, pH in units of 0.2 pH, CEC in units of 10 meq (100 g)⁻¹, bulk density in units of 0.1 kg m⁻³, and SOC in units of 0.2% up to 6% and then in units of 5% to the maximum of 30%.


Figure 26 shows the segmentation based on all depth slices of SOC concentration, units 1000 x %, for three of the slices. Some segments are well-separated, notably the depressions with swamps and organic soils, as well as sections with different intensities of drumlins.

Figure 27 shows the segmentation based on all depth slices of clay %, for three of the slices. The segments are quite large and do not identify collections of the main landscape elements, i.e., drumlins and depressions.

Similar and even worse results were found with other properties, as well as with an attempt to use all properties at one depth slice.



SSURGO vs. supercells, compactness parameter 0.2



very hard to
 decipher...
 how about a
 zoomed / detail
 map?

Figure 25. Supercells of clay concentration % of the 0-5 cm layer from SOLUS 100 m overlaid by the polygon boundaries (in red) from gSSURGO

must be larger

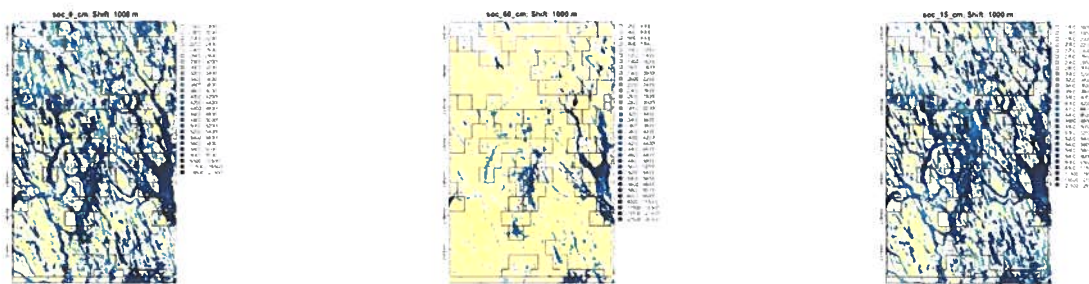


Figure 26. Segmentation based on all depth slices of SOC concentration

what do ~~these~~
 colors represent?
 source?



Figure 27. Segmentation based on all depth slices of clay concentration

is that what the colors represent?

5.3 Clustering

Because of the poor results of segmentation, we do not present the results of clustering for this case study.

5.4 Evaluation

this metric needs to be defined @ the start.

SOLUS100 was able to "speak for itself", but the message was not clear and even misleading. Notably, the attempts to aggregate and segment based on a representation of the profile resulted in unrealistic maps. In this area the landscape pattern is striking and easy to map by conventional methods. SOLUS was unable to approximate the conventional map, let alone improve its resolution. This is likely because SOLUS lacks locally-important covariates to represent this recently glaciated soil landscape with its characteristic drumlins.

*Also: global model — local interpretation
 fraught w/ peril*

6 Discussion

The supercells algorithm was able to delineate relatively homogeneous soils, based on all soil properties and layers in the BIS-4D example and the SoilGrids SOC example, but failed completely with SOLUS. A limitation of this approach is that there is no objective way to adjust the compactness and supercell number parameters, other than the expert opinion on which choice looks most "realistic". However, the minimum size parameter can be set to match a minimum legible delineation corresponding to a desired map scale.

The GeoPAT algorithm was able to segment DSM products into objectively-defined areas made up of fixed-size blocks, each relatively homogeneous in its pattern internally and relatively isolated from its neighbours. Segmentation was quite successful on appropriate scales for BIS-4D and the test area and property of SoilGrids v2.0, but much less successful for the test area of SOLUS100. The class composition of segments, although not their internal spatial pattern, were well-characterised by normalized Shannon Entropy.

A limitation of the GeoPAT approach is the requirement for relatively large numbers of pixels per grid cell, and the rectangular shape of the grid cells that are combined into segments. Thus, the segment boundaries can not follow complex natural boundaries. Also, the landscape segments are at much more general scale than the source map.



The question remains as to the relation of the supercells or segments with the actual soil landscape at the several scales. There are two related questions. (1) For aggregation, do the relatively homogeneous (according to the supercells algorithm) groups of pixels correspond to landscape elements? These would correspond to polypedons or consociations. (2) For segmentation, do the patterns of pixels within the segment correspond to finer-scale patterns at the design scale of the segmentation? These would correspond to associations or complexes.

In the case of BIS-4D and the detailed traditional Dutch soil survey, the degree to which the aggregation matches the published map (Figure 3) is likely sufficient. The success of segmentation was discussed in §3.2. It is not clear which segmentation scale is the most appropriate.

In the case of SoilGrids, the “true” soil landscape pattern in the test area is not so clear. When comparing SoilGrids with the USA, a problem is that the detailed gSSURGO map (NRCS Soils, 2022) has been compiled from multiple survey areas, mapped over many years, and with imperfect correlation between areas. This is compiled from traditional surveys at design scales from 1:12'000 to 1:24'000 in most areas, but somewhat coarser in less populated areas in the western USA. The INEGI map in México is a consistent 1:250'000 national product (Instituto Nacional de Estadística, Geografía e Informática (INEGI), 2024), which can show a minimum delineation of 250 ha. Figure 28 shows a SOC stock maps of the study area, compiled from the above-mentioned USA and Mexican sources by the FAO as part of the Global Soil Organic Carbon Map (GSOCmap) project (FAO, 2018). Version 1.6.1 of this product was downloaded from the FAO's Global Soil Information System (GloSIS) (FAO, 2024). The inconsistency in values and pattern between México and the USA is obvious, as are several sharp boundaries between survey areas in the USA. So it is difficult to evaluate how well SoilGrids identifies supercells or segments.

In the case of SOLUS, the geomorphology and soil pattern of the test area is well understood and has been mapped in detail. Of the SOLUS layers only soil organic carbon and coarse fragment volume showed a relation with known patterns in the test area. Aggregation based on multiple properties completely failed to find landscape units. Segmentation based on multiple properties failed to find more general units with consistent internal patterns.


7 Conclusions

The methods presented in this paper are part of an effort to evaluate DSM products based on how well they represent the soil landscape. The approach taken here complements pattern analysis of the DSM product, which characterises the map without attempting aggregation or segmentation, as in Rossiter et al. (2022). Both the aggregation and segmentation approaches were able to allow the DSM product “speak for itself”. Individual predictions in pixels were combined into possible soil-landscape elements, which could be evaluated statistically and by expert judgment. Both of these approaches require the intervention of the analyst to select scales and parameters, often with large differences in resulting patterns. This has the advantage that the analyst can match desired scales of landscape analysis, and indeed can perform a multi-resolution evaluation. The analysis of the resulting maps is a significant addition to the commonly-used “point”-based evaluation statistics, which (1) do not evaluate the full map, (2) even at point support, do not take into account the spatial relation between evaluation points. We hope that

*
SOC
pattern
are too
“fine”
to use
here



how is this
figure helpful?



units

Figure 28. Global Soil Organic Carbon (GSOC) map (part). Boundary is between México (south) and the USA (north). Source and legend: FAO (2024).

495 this will stimulate digital soil mappers to evaluate their own products in this light. This should lead to clearer communication with DSM users, so that digital soil maps become more widely accepted and properly used.

* What about eval (pred.) beyond point-based methods?

Code and data availability. The Geopat modules are available at its GitHub repository¹. The *superpixels* R package is available at CRAN² and must be installed from within the R environment. The analysis code for this paper is available in a GitLab repository³. The datasets used in case studies can be obtained from the websites referenced in the text.

500 *Author contributions.* DGR conceived of the approach, wrote the code, carried out the test cases, and wrote the initial draft. LP contributed advice at every step, especially the concepts and interpretations, and contributed detailed knowledge of SoilGrids.

¹<https://github.com/Nowosad/geopat2>

²<https://cran.r-project.org/web/packages/superpixels/index.html>

³https://git.wur.nl/isric/scientific-publications/Rossiter-2025-Soil_landscapes_from_DSM

<https://doi.org/10.5194/egusphere-2025-1896>

Preprint. Discussion started: 20 May 2025

© Author(s) 2025. CC BY 4.0 License.



Competing interests. The authors have no competing interests



References

- Arrouays, D., McBratney, A., Bouma, J., Libohova, Z., Richer-de-Forges, A. C., Morgan, C. L., Roudier, P., Poggio, L., and Mulder, V. L.: Impressions of Digital Soil Maps: The Good, the Not so Good, and Making Them Ever Better, *Geoderma Regional*, 20, e00 255, <https://doi.org/10.1016/j.geodrs.2020.e00255>, 2020.
- Batjes, N. H., Calisto, L., and de Sousa, L. M.: Providing Quality-Assessed and Standardised Soil Data to Support Global Mapping and Modelling (WoSIS Snapshot 2023), *Earth System Science Data Discussions*, pp. 1–46, <https://doi.org/10.5194/essd-2024-14>, 2024.
- Borden, R. W., Baillie, I. C., and Hallett, S. H.: The East African Contribution to the Formalisation of the Soil Catena Concept, *CATENA*, 185, 104 291, <https://doi.org/10.1016/j.catena.2019.104291>, 2020.
- Brouwer, F., Maas, G., Teuling, K., Harkema, T., and Verzandvoort, S.: Bodemkaart en Geomorfologische Kaart van Nederland: actualisatie 2020-2021 en toepassing : deelgebieden Gelderse Vallei-Zuid en -West en Veluwe-Zuid, Tech. rep., Wettelijke Onderzoekstaken Natuur & Milieu, <https://doi.org/10.18174/557455>, 2021.
- FAO: Global Soil Organic Carbon Map (GSOCmap) : Technical Report, FAO, Rome, Italy, ISBN 978-92-5-130439-6, 2018.
- FAO: GLOSIS: Global Soil Information System, <https://data.apps.fao.org/glosis/?lang=en>, 2024.
- Forbes, T., Rossiter, D., and Van Wambeke, A.: Guidelines for Evaluating the Adequacy of Soil Resource Inventories, Cornell University Department of Agronomy, Ithaca, NY, ISBN 978-0-932865-07-6, 1982.
- Fridland, V. M.: Structure of the Soil Mantle, *Geoderma*, 12, 35–42, [https://doi.org/10.1016/0016-7061\(74\)90036-6](https://doi.org/10.1016/0016-7061(74)90036-6), 1974.
- Haans, J. C. F. M.: De bodem van Nederland : toelichting bij de Bodemkaart van Nederland, schaal 1 : 200 000. Stichting voor Bodemkartering, Wageningen, 1965.
- Hall-Beyer, M.: GLCM Texture: A Tutorial v. 3.0 March 2017, <http://hdl.handle.net/1880/51900>, 2017.
- Haralick, R. M., Shanmugam, K., and Dinstein, I.: Textural Features for Image Classification. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3, 610–621, <https://doi.org/10.1109/TSMC.1973.4309314>, 1973.
- Helfenstein, A., Mulder, V. L., Hack-ten Broeke, M. J. D., van Doorn, M., Teuling, K., Walvoort, D. J. J., and Heuvelink, G. B. M.: BIS-4D: Mapping Soil Properties and Their Uncertainties at 25 m Resolution in the Netherlands, *Earth System Science Data*, 16, 2941–2970, <https://doi.org/10.5194/essd-16-2941-2024>, 2024.
- Higgins, B. A.: Soil Survey of Wayne County, New York, U.S. Soil Conservation Service, [Washington], 1978.
- Hijmans, R. J., Barbosa, M., Bivand, R., Brown, A., Chirico, M., Cordano, E., Dyba, K., Pebesma, E., Rowlingson, B., and Sumner, M. D.: Terra: Spatial Data Analysis, <https://cran.r-project.org/web/packages/terra/index.html>, 2025.
- Hole, F. D.: An Approach to Landscape Analysis with Emphasis on Soils, *Geoderma*, 21, 1–23, [https://doi.org/10.1016/0016-7061\(78\)90002-2](https://doi.org/10.1016/0016-7061(78)90002-2), 1978.
- Hudson, B. D.: The Soil Survey as Paradigm-Based Science, *Soil Science Society of America Journal*, 56, 836–841, <https://doi.org/10.2136/sssaj1992.03615995005600030027x>, 1992.
- Instituto Nacional de Estadística. Geografía e Informática (INEGI): Edafología, <https://www.inegi.org.mx/temas/edafologia/>, 2024.
- ISRIC-World Soil Information: Index of /Soilgrids/Latest/Data_aggregated/1000m/, https://files.isric.org/soilgrids/latest/data_aggregated/1000m/, 2024a.
- ISRIC-World Soil Information: SoilGrids Web Portal, <https://soilgrids.org>, 2024b.
- Jasiewicz, J., Netzel, P., and Stepinski, T.: GeoPAT: A Toolbox for Pattern-Based Information Retrieval from Large Geospatial Databases, *Computers & Geosciences*, 80, 62–73, <https://doi.org/10.1016/j.cageo.2015.04.002>, 2015.



- 540 Jasiewicz, J., Stepinski, T., and Niesterowicz, J.: Multi-Scale Segmentation Algorithm for Pattern-Based Partitioning of Large Categorical
 Rasters, *Computers & Geosciences*, 118, 122–130, <https://doi.org/10.1016/j.cageo.2018.06.003>, 2018.
- Johnson, W. M.: The Pedon and the Polypedon, *Soil Science Society of America Journal*, 27, 212–215,
<https://doi.org/10.2136/sssaj1963.03615995002700020034x>, 1963.
- Lin, J.: Divergence Measures Based on the Shannon Entropy, *IEEE Transactions on Information Theory*, 37, 145–151,
 545 <https://doi.org/10.1109/18.61115>, 1991.
- McBratney, A. B., Mendonça Santos, M. L., and Minasny, B.: On Digital Soil Mapping, *Geoderma*, 117, 3–52, [https://doi.org/10.1016/S0016-7061\(03\)00223-4](https://doi.org/10.1016/S0016-7061(03)00223-4), 2003.
- Menzies, J., Hess, D. P., Rice, J. M., Wagner, K. G., and Ravier, E.: A Case Study in the New York Drumlin Field, an Investigation
 Using Microsedimentology, Resulting in the Refinement of a Theory of Drumlin Formation, *Sedimentary Geology*, 338, 84–96,
 550 <https://doi.org/10.1016/j.sedgeo.2016.01.017>, 2016.
- Ministerie van Volkshuisvesting en Ruimtelijke Ordening: Bodemkaart (SGM). <https://basisregistratieondergrond.nl/inhoud-bro/registratieobjecten/modellen/bodemkaart-sgm/>, 2024.
- Mulder, V. L., Roudier, P., and Arrouays, D.: Editorial: Digital Soil Mapping - Advancing the Knowledge Frontiers, *Frontiers in Soil Science*,
 3, <https://doi.org/10.3389/fsoil.2023.1225672>, 2023.
- 555 Nauman, T.: Data from: Soil Landscapes of the United States 100-Meter (SOLUS100) Soil Property Maps Project Repository.
<https://doi.org/10.15482/USDA.ADC/25033856.v1>, 2024.
- Nauman, T. W., Kienast-Brown, S., Roecker, S. M., Brungard, C., White, D., Philippe, J., and Thompson, J. A.: Soil Landscapes of the
 United States (SOLUS): Developing Predictive Soil Property Maps of the Conterminous United States Using Hybrid Training Sets, *Soil
 Science Society of America Journal*, 88, <https://doi.org/10.1002/saj2.20769>, 2024.
- 560 Nenkam, A. M., Wadoux, A. M. J.-C., Minasny, B., Silatsa, F. B. T., Yemefack, M., Ugbaje, S. U., Akpa, S., Zijl, G. V., Bouasria, A.,
 Bouslih, Y., Chabala, L. M., Ali, A., and McBratney, A. B.: Applications and Challenges of Digital Soil Mapping in Africa, *Geoderma*,
 449, 117 007, <https://doi.org/10.1016/j.geoderma.2024.117007>, 2024.
- Netzel, P., Nowosad, J., Jasiewicz, J., Niesterowicz, J., and Stepinski, T.: GeoPAT 2: User's Manual, <https://zenodo.org/records/1291123>,
 2018.
- 565 Nowosad, J.: Supercells 1.0.3: Superpixels of Spatial Data, <https://jakubnowosad.com/supercells/index.html>, 2025.
- Nowosad, J. and Stepinski, T. F.: Towards Machine Ecoregionalization of Earth's Landmass Using Pattern Segmentation Method, *International
 Journal of Applied Earth Observation and Geoinformation*, 69, 110–118, <https://doi.org/10.1016/j.jag.2018.03.004>, 2018.
- Nowosad, J. and Stepinski, T. F.: Extended SLIC Superpixels Algorithm for Applications to Non-Imagery Geospatial Rasters, *International
 Journal of Applied Earth Observation and Geoinformation*, 112, 102 935, <https://doi.org/10.1016/j.jag.2022.102935>, 2022.
- 570 NRCS Soils: Description of Gridded Soil Survey Geographic (gSSURGO) Database.
https://www.nrcs.usda.gov/wps/portal/nrcs/detail/soils/home/?cid=nrcs142p2_053628, 2022.
- Pawlik, Ł., Kasprzak, M., Ignatiuk, D., Glowacki, T., Milczarek, W., and Kajdas, J.: Evaluation of the Hillslope Fine-Scale Morphology under
 Forest Cover with Pit-Mound Topography - Integration of Geomorphometry, Geophysical Methods, and Soil Features, *Geomorphology*,
 460, 109 283, <https://doi.org/10.1016/j.geomorph.2024.109283>, 2024.
- 575 Piikki, K., Wetterlind, J., Söderström, M., and Stenberg, B.: Perspectives on Validation in Digital Soil Mapping of Continuous Attributes – a
 Review, *Soil Use and Management*, 37, 7–21, <https://doi.org/10.1111/sum.12694>, 2021.



- Poggio, L., Gimona, A., Brown, I., and Castellazzi, M.: Soil Available Water Capacity Interpolation and Spatial Uncertainty Modelling at Multiple Geographical Extents, *Geoderma*, 160, 175–188, <https://doi.org/10.1016/j.geoderma.2010.09.015>, 2010.
- 580 Poggio, L., de Sousa, L. M., Batjes, N. H., Heuvelink, G. B. M., Kempen, B., Ribeiro, E., and Rossiter, D.: SoilGrids 2.0: Producing Soil Information for the Globe with Quantified Spatial Uncertainty, *SOIL*, 7, 217–240, <https://doi.org/10.5194/soil-7-217-2021>, 2021.
- Rossiter, D. G., Poggio, L., Beaudette, D., and Libohova, Z.: How Well Does Digital Soil Mapping Represent Soil Geography? An Investigation from the USA, *SOIL*, 8, 559–586, <https://doi.org/10.5194/soil-8-559-2022>, 2022.
- Science Committee: Specifications: Tiered GlobalSoilMap.Net Products; Release 2.4 [07/12/2015]. Appendix C: Correlations of Soil Properties Derived from Different Soil Analytical Methods, Tech. rep., GlobalSoilMap.net, 2015.
- 585 Soil Science Division Staff: Soil Survey Manual, no. 18 in USDA Handbook, Government Printing Office, Washington, DC, <https://www.nrcs.usda.gov/resources/guides-and-instructions/soil-survey-manual>, 2017.
- Soil Survey Staff: Soil Taxonomy: A Basic System of Soil Classification for Making and Interpreting Soil Surveys, US Department of Agriculture Soil Conservation Service, Washington, DC, 2nd edn., 1999.
- Steur, G. G. L. and Heijink, W.: Bodemkaart van Nederland : Algemene Begrippen En Indelingen : Schaal 1 : 50.000. Stichting voor
 590 Bodemkartering (STIBOKA), Wageningen, ISBN 90-220-0754-5, 1980.
- Vaysse, K. and Lagacherie, P.: Using Quantile Regression Forest to Estimate Uncertainty of Digital Soil Mapping Products, *Geoderma*, 291, 55–64, <https://doi.org/10.1016/j.geoderma.2016.12.017>, 2017.
- Vink, A. P. A.: Planning of Soil Surveys in Land Development, no. 10 in International Institute for Land Reclamation and Improvement, Veenman & Zonen, Wageningen, 1963.