

Appendices to: Ultrahigh Dimensional Variable Selection for Interpolation of Point Referenced Spatial Data: A Digital Soil Mapping Case Study

Benjamin R. Fitzpatrick^{1,2,3,*}, David W. Lamb^{4,2}, Kerrie Mengersen^{5,1,2,3}

1 Mathematical Sciences School, Queensland University of Technology (QUT), Australia

2 Cooperative Research Centre for Spatial Information (CRCSI), Australia

3 Institute for Future Environments, Queensland University of Technology (QUT), Australia

4 School of Science and Technology, University of New England (UNE), Australia

5 ARC Centre of Excellence for Mathematical and Statistical Frontiers, Queensland University of Technology (QUT), Australia

*** E-mail: Corresponding b1.fitzpatrick@qut.edu.au**

Appendix A and Appendix B contain additional context for our case study while Appendix C, Appendix D and Appendix E contain expanded considerations of methodology relevant to our case study. Appendix A contains a description of the study site and field methodology. In Appendix B we explain our choice of covariates in terms of data availability and previous soil carbon modelling studies. In Appendix C we compare and contrast modelling methods of various degrees of appropriateness for data like those of our case in order to make the case for using modelling method we have adopted. In Appendix D we outline our method for filtering design matrices to enforce a maximum correlation coefficient magnitude between remaining covariates. In Appendix E we explain our choice of design matrix filtering austerity and training set size for our cross validation scheme.

1 Appendix A: Study Site and Field Methodology

1.1 Study Site

The study site was a 137ha area of land on the Sustainable, Manageable, Accessible, Rural Technology (SMART) Farm of the University of New England near Armidale, New South Wales (NSW), Australia. The north-west corner of the SMART farm had coordinates 30°22'59"S 151°35'23"E and the south-east corner of the SMART farm had coordinates 30°27'26"S 151°39'52"E. The study site was situated at the base of Mount Duval (1393m [1]) and formed a part of the Uralla Plutonic Suite/Mount Duval Adamellite (acid porphyritic, hornblende-biotite monzogranite) and was characterized by yellow and brown

chromosols [2] upon the hills with alluvial soils and siliceous sand complexes distributed along the drainage routes. The study site typically received 790mm of annual, summer dominant rainfall [3]. The maximum elevation within the study site was 1120m and the elevation range across the study site was less than 110m. The study site consisted of selectively cleared native pasture containing some remnant vegetation and regrowth and had historically been grazed by sheep and cattle. The south-east corner of the study site was situated at grid reference 371434E, 6632499N MGA GDA 94 Zone 56. Being used to grow pasture and receiving in excess of 450mm of annual rainfall, soils in this area fell within the class of agricultural soils deemed to have the highest potential of any agricultural soils in NSW for sequestration of atmospheric carbon as Soil Organic Carbon (SOC) [4]. A hill-shaded plot of a digital elevation model cropped to the approximate boundaries of the study site and an aerial photograph of the study site have been presented as the panels of Fig. 1.

1.2 Proximal Data Collection

1.2.1 Covariates

The study area was surveyed three times in 2009 for soil apparent electrical conductivity (EC_a) and the reflectance of the top of the pasture canopy under active illumination. The first survey was conducted in the warm summer month of February, following a prolonged dry period (0mm of precipitation in the preceding seven days) in what was otherwise the second wettest month of the year. The second survey was conducted in May, immediately after a significant rainfall event (84mm in the preceding seven days) in what was otherwise the cooler middle of the year when less rain fell than in summer. The third survey was conducted in November, which marked the end of the winter growing season and was a month in which less rain fell than did in the wet month of October and the very wet month of December. These surveys were conducted by an all terrain vehicle (ATV) towing a sensory array that consisted of a specially configured Geonics EM38 unit (Geonics Ontario Canada), an LED illumination array, near-infrared and visible light reflectance sensors (Crop CircleTM, Holland Scientific, USA) and a differential global positioning system (dGPS) (Trimble, Sunnyvale CA USA).

The Geonics EM38 instrument measured soil EC_a which may be understood as the integral of the electrical conductivity response recorded across soil depths. To collect these data the EM38 instrument emitted a varying magnetic field which induced an electric current in the soil underneath the instrument. The electric current induced in the soil created a magnetic field, the strength of which was proportional to the amount of electric current induced. The strength of the magnetic field that resulted from the current induced in the soil was taken as indicative of the strength of this induced electric current and recorded by the instrument. The strength of the electric current induced in the soil by the instrument-emitted electric field varied as a function of depth. With a Geonics

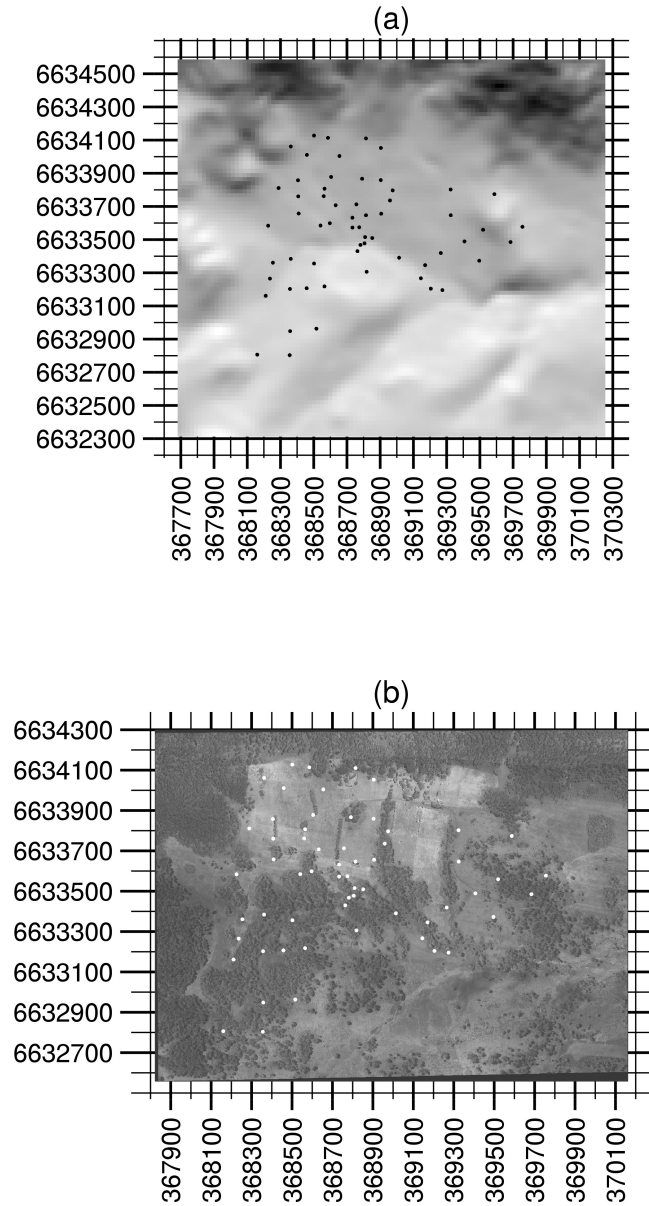


Figure 1: a) A hill shaded terrain surface for the study site calculated from a digital elevation model. b) An aerial photograph of the study site. The locations at which soil core samples were collected have been depicted as filled circles in both panels.

EM38 instrument operated in the vertical dipole orientation, as Garraway et al. orientated the instrument in their survey, the relative signal response, S_v , varied with depth, z , as follows $S_v(z) = \frac{4z}{(4z^2+1)^{\frac{3}{2}}}$ [5]. Due to air essentially not conducting electricity at these strengths of inductive magnetic fields, the conductivity signal commenced at the soil surface. Thus the EC_a measured by the EM38 instrument in vertical dipole mode was $\int_H^\infty \sigma(z)S_v(z)dz$ where $\sigma(z)$ was the electrical conductivity of the soil at depth z and $S_v(z)$ was the relative signal response of the soil at depth z [5]. Thus, when operated in vertical dipole mode, the Geonics EM38 instrument had a peak relative signal response to the electrical conductivity of the soil 350mm below the instrument and was essential insensitive to the electrical conductivity of any medium immediately below the instrument. Garraway et al. mounted a Geonics EM38 instrument on a rubber sled that held the instrument approximately 15mm above the ground [6] and towed this sled around the study area. Thus the EC_a data we have for the case study would have been dominated by the soil electrical conductivity at depths around the relative signal response peak at 335mm below the soil surface with very little contribution from the electrical conductivity of the soil surface.

The reflectance sensors measured top of pasture reflectance of active illumination in the Near InfraRed (NIR) and visible Red (RED) regions of the electromagnetic spectrum. This style of proximal sensing of the reflectance of active illumination has been applied to both crop and soil mapping [7]. The electro-optical principles governing the effectiveness of such sensors have been discussed in Holland et al. [2012]. The EC_a , RED reflectance and NIR reflectance were recorded simultaneously at regular time intervals as the ATV traversed the study area. The east-west and north-south coordinates that accompanied each of these covariate observations were also recorded along with the associated Position Dilution Of Precision (PDOP) and Horizontal Dilution of Precision (HDOP) values. The data from each of the ATV surveys were cleaned of observations with large inaccuracies in positioning (as assessed by HDOP and PDOP measurements). The number of point observations that remained from each ATV survey after this cleaning had been conducted have been included in Table 1.

1.2.2 Response

In 2009, the study area was divided into five strata by means of k-means clustering [8] the red, green and blue channels from aerial imagery and the EC_a data from the February survey [3]. At least six locations for soil core sampling were randomly selected within each strata with additional locations manually selected to improve the representation of landscape attributes. This stratified random sampling approach to choosing locations for soil core samples was similar to the process outlined in [9]. As a result, soil samples were collected to a depth of 200mm at 60 locations across the study area with locations georeferenced using a differential Global Positioning System (dGPS) instrument. At each of the 60 dGPS coordinates, three soil cores were collected from within a

Table 1: A summary of the types of data used in the case study, including for each type of data the: date of collection/publication, type of variable and number of observations or resolution (pixel dimensions).

Variable	Collected or Published	Type	Observations or Resolution
SOC	2009	Geostatistical	60
EC_a & VI	Feb 2009	Geostatistical	16179
	May 2009	Geostatistical	16094
	Nov 2009	Geostatistical	14059
FPC	2011 & 2012	Raster	10m ²
DEM & DEM Products	2004	Raster	25m ²
Radiometric & Electromagnetic Survey	2002 & 2003	Raster	50m ²

abbreviations: VI = Vegetation Indices, DEM = Digital Elevation Model, FPC = Foliar Projective Cover, Feb = February and Nov = November.

1m radius of the coordinates and aggregated to form a single soil sample that was laboratory analyzed for percentage SOC (hereafter %SOC). Garraway et al. detailed the preparation of soil samples for assessment of the total organic carbon with a Carlo Erba NA 1500 solid sample analyzer (Carlo Erba Instruments, Milan, Italy).

1.3 Remotely Sensed Data

A 25m² resolution Digital Elevation Model (DEM) (sourced from the Department of Lands, New South Wales State Government, Australia) for the Armidale-Dumaresq region which contained the catchment in which the study area was situated was read into the System for Automated Geoscientific Analyses (SAGA v2.1.0 [10]) software to calculate the terrain topographic and hydrological attributes listed in Table 2. The ready availability of the attributes listed in Table 2 and potential relevance of topography and hydrology to SOC distributions generally, lead us to include the full suite of such attributes that may be calculated with SAGA as potential covariates in this analysis. The resulting GIS layers were then read into R [11] with the ‘RSAGA’ [12] package. Full cover layers for Foliar Projective Cover (FPC) produced by applying the Statewide Landcover and Trees Study (<https://www.qld.gov.au/environment/land/vegetation/mapping/slats-methodology/>) method to imagery from the SPOT5 satellite (10m² resolution) were acquired for the study region from 2012 and 2011 from the New South Wales State Government Department of Environment. These layers were read into R [11] with the ‘raster’ package [13]. Potassium (K), Uranium (U) and Thorium (Th) count layers from an airborne γ ray radiometric survey [14] and similar layers from six channels of electromagnetic imagery [15] were read into R [11] with the ‘raster’ [13] package.

Table 2: The 63 potential covariates from which models were built for soil organic carbon in this study.

Source	Covariate Name	Acronym
ATV Top of Pasture	Soil Apparent Electrical Conductivity	ECA
Surveys	Near InfraRed Reflectance	NIR
12 covariates	Red Reflectance	RED
from each of February,	Simple Ratio	SR
May & November	Difference Vegetation Index	DVI
= 36 covariates	Normalized Difference Vegetation Index	NDVI
	Soil Adjusted Vegetation Index	SAVI
	Non-Linear Vegetation Index	NLVI
	Modified Non-Linear Vegetation Index	MNLVI
	Modified Simple Ratio	MSR
	Transformed Vegetation Index	TVI
	Re-normalised Difference Vegetation Index	RDVI
Terrain &	Catchment Area	CatAr
Hyrdology Metrics	Catchment Height	CatHe
Calculated from a	Catchment Slope	CatSl
25m ² resolution	Cosine(Aspect)	CosAsp
DEM	Elevation	Elev
= 16 Covariates	Slope Length Factor	LSF
	Plan Curvature	PlanC
	Profile Curvature	ProfC
	Sky View Factor	SVF
	Slope	Slp
	Stream Power Index	SPI
	Terrain Ruggedness Index	TRI
	Topographic Position Index	TPI
	Vector Terrain Ruggedness	VTR
	Visible Sky	VS
	Wetness Index	WI
Foliar Projective Cover	2011	FPCI
Layers = 2 Covariates	2012	FPCI
Electromagnetic	1	MagI
Imagery Channels	2	MagII
= 6 Covariates	3	MagIII
	4	MagIV
	5	MagV
	6	MagVI
γ Radiometric Layers	Potassium	K
= 3 Covariates	Thorium	Th
	Uranium	U

Cartographic projection systems are used to project latitude and longitude coordinates from a particular region of the surface of the Earth onto a two dimensional plane. In this analysis all spatial coordinates were treated as coordinates on a two dimensional plane and as such it was important to ensure that all data layers utilised the same projection system. This was accomplished with the R [11] package ‘raster’ [13] through the re-projection of all data layers that did not already use the most common projection system among the data layers to this projection system, namely a UTM projection for zone 56 South using the WGS84 ellipse.

2 Appendix B: Choice of Covariates

In this Appendix we explain our choice of environmental characteristics considered as potential covariates for modelling soil carbon. The first stage in collating this set of potential covariates was to identify the common covariates used in soil carbon modelling via a review of the available literature. Once we had this list of potential covariates, the second stage was to identify which of these we could obtain for our case study site.

2.1 Soil Organic Carbon and Soil Apparent Electrical Conductivity

Soil apparent electrical conductivity (EC_a) has variously been found indicative of some or all of soil: moisture content, pore distribution, pore size, salinity, clay content, mineralogy, cation exchange capacity and temperature [16]. Discretizing soil EC_a values into four classes produced significant factors in separate ANOVAs for each of soil total organic matter, soil particulate organic matter and total soil carbon in the top 30cm of soil sampled across 250 ha of farmland used for wheat, corn and millet in the state of Colorado in the United States of America (USA) [17]. A negative correlation was detected between soil EC_a and SOC ($r = -0.42$) in the top 90cm of soil in 9ha of farmland that had a long history of cotton row cropping in Alabama, USA but no such correlation was detected in the top 30cm [18]. Soil EC_a was more correlated with soil carbon ($r = -0.65$ to -0.76) than were any of the local relative elevation, local slope and satellite measured soil surface reflectance of near bare fields at three sites (48.7 ha, 52.4 ha and 65.4 ha in area respectively) in farmland used for maize and soybean cropping in Nebraska, USA [19]. Studies have also been performed where it seemed likely that influences other than SOC were dominating the soil EC_a signal. Principal Component Analysis (PCA) of soil properties in an 8ha agricultural field in Flanders, Belgium yielded largely independent spatial patterns in soil pH, EC_a and SOC [20]. Furthermore, in the soils of the world’s oldest continuous cotton experiment (Alabama, USA) little correlation was detected between SOC in the top 15cm of soil and the EC_a of the top 30cm of soil across 0.4ha of land [21].

2.2 Soil Organic Carbon and Spectral Vegetation Indices

Land plant biomass in any location will have been influenced by many soil conditions. Through direct and indirect effects on soil: structural stability, water and nutrient retention, faunal activity and diversity, and elemental recycling [22] SOC levels may have influenced land plant biomass in many situations. Conversely, plants will have also provided an input of carbon to SOC via litter fall and root turn over. Thus empirical correlations between %SOC and plant biomass are plausible. The amount of green plant biomass present in a location may be indicated by the density of green leaves present above the soil there. The density of green leaves present above the land surface has often been estimated from the reflection spectra of the land surface when observed from above canopy height. Three spectral signatures of particular interest for these considerations have been identified as: those of healthy green leaves, those of stressed or senescent green leaves and those of agricultural soils. The marked difference in the intensities of light reflected from green leaves in the visible red (RED) and near infra red (NIR) wavelengths and the general weakness or absence of such a ‘red edge’ in the spectral reflectance signatures of stressed or senescent leaves and agricultural soils has formed the basis of many spectral vegetation indices used for monitoring vegetation [23]. Healthy leaves have typically exhibited a high reflectance of NIR light due to scattering of these wavelengths at the interface between the mesophyll and cell walls and low absorption of these wavelengths by photosynthetic pigments and organelles [23]. Whereas, healthy green leaves have typically displayed a low reflectance of visible wavelengths due to the high absorption of light in this region of the spectrum by photosynthetic and accessory pigments [23]. Green plant stress and senescence have often manifested in the form of depressed chlorophyll concentrations and the expression of accessory leaf pigments which together lower the absorption of visible wavelengths by leaves. Where stress or senescence has lowered the absorption of visible wavelengths by leaves the reflectance peak of such leaves will have widened correspondingly from the green region of the spectrum typical of healthy green leaves through towards redder wavelengths. Where stress or senescence has manifested in this broadened reflectance of visible spectra by leaves a simultaneous decrease in the NIR reflectance of these leaves will have also occurred. Thus where it has occurred, stress or senescence would have resulted in a loss or weakening of the abrupt ‘red edge’ typical of the reflectance spectra of healthy green leaves [23]. Similarly, agricultural soils have been characterised by a lack of any such sharp contrast in the reflectance intensities of different wavelengths [23]. Most spectral vegetation indices have been constructed as functions of NIR and RED reflectance designed to quantify some aspect of the expected differences between the reflectance spectra of healthy green leaves and those of soils and stressed or senescent leaves. Thus vegetation indices have been designed to provide an estimate of the quantity of green plant material that contributed to the reflectance spectra of a land surface by quantifying the extent to which such differences in NIR and RED intensities occurred within this spectra [23]. The spectral signature obtained from an entire canopy may have

differed markedly from that obtained from an individual green leaf and furthermore may have varied across the growing season as the canopy geometry altered with plant growth [23]. Thus differences between vegetation indices calculated from reflectance surveys conducted at the same location at different stages in the year may have yielded information about how green plant biomass there changed across the growing season. Changes in biomass across a growing season could, in turn, have been indicative of soil attributes bearing other stronger influences on plant biomass lost or accrued. Thus the comparison of vegetation index values when collected in the same location at different points in the plant growth cycle could have been indicative of soil properties. For instance, plants that grew in poorer soils and experienced otherwise similar conditions could reasonably be expected to have produced less biomass in a growing season than the same plants in more favourable soils. Furthermore, the presence of SOC in surface soils has been documented as increasing soil aggregation thereby creating larger lacunae (also referred to as interstitial spaces) into which water may drain from the soil surface. Thus SOC has been broadly classified as beneficial to the infiltration of soil by water and the retention of water by soil [24]. Thus increased %SOC at a location could in turn aid water retention there and allow plants that undergo seasonal curing (e.g. grasses like those at our case study site) to remain green longer into the dry part of the year at that location.

A review of studies of the correlations between soil organic matter and crop reflectance in visible and NIR regions of the electromagnetic spectrum and the vegetation indices thereby derived formed a section of the review paper [16]. In certain situations, above ground plant biomass may have been related to SOC concentrations. Where this has been the case vegetation indices may have held relevance to %SOC concentrations. This seems to have been the case in the following studies. The Normalised Difference Vegetation Index (NDVI [25]) and the Soil Adjusted Vegetation Index (SAVI [26]) have both achieved considerable popularity as vegetation indices. A positive correlation between canopy NDVI and biomass was detected in a 7ha cotton field in Larissa, Greece [27]. Furthermore, the pasture canopy SAVI was found to be the best of a range of vegetation spectral indices for predicting pasture green dry mass across four 50ha paddocks in New South Wales, Australia [28]. Much like soil EC_a , plant visible and NIR reflectance have variously been found correlated with soil properties other than SOC such as soil moisture and Cation Exchange Capacity (CEC) [16] along with prevailing climate, ecosystem, terrain and physical soil properties [29]. We have summarised studies that found correlations between any of a selection of vegetation indices that may be calculated from reflectance intensities in the NIR and RED bands and soil carbon in Table 3. Since our study site consisted of native pasture with remnant woody vegetation we restricted this summary to one of studies that were conducted in pastures, grasslands, prairies and steppes.

Table 3: Vegetation Indices reported to have been correlated with grass biomass (which were thus possibly also correlated with soil carbon) in pastures, grasslands, prairies and steppes.

VI	Full Name	Formula	Correlation with Grass Biomass
SR	Simple Ratio	$\frac{NIR}{RED}$	[28]
DVI	Difference Vegetation Index	$NIR - RED$	[30, 31]
NDVI	Normalized Difference Vegetation Index	$\frac{NIR-RED}{NIR+RED}$	[28, 30]
SAVI	Soil Adjusted Vegetation Index*	$\frac{(NIR-RED)(1+L)}{NIR+RED+L}$	[28]
NLI	Non-Linear Vegetation Index	$\frac{NIR^2-RED}{NIR^2+RED}$	[28]
MNLI	Modified Non-Linear Vegetation Index	$\frac{(NIR^2-RED)(1+L)}{NIR^2+RED+L}$	[28]
MSR	Modified Simple Ratio	$\frac{\frac{NIR}{RED}-1}{(\frac{NIR}{RED})^{1/2}+1}$	[28]
TVI	Transformed Vegetation Index	$(NDVI + 0.5)^{1/2}$	[31]
RDVI	Re-Normalised Difference Vegetation Index	$\frac{NIR-VIS}{(NIR+VIS)^{1/2}}$	[31]

* L = 0.5 recommended for wide range of leaf area index values

2.3 Soil Organic Carbon and Scattered Paddock Trees

Below ground root turnover and above ground litter fall have both been recognised as sources of detrital carbon to topsoil. Thus in native pastures with remnant woody vegetation in the form of scattered paddock trees, such as the pastures from which our case study data were collected, the locations of these trees may have influenced the spatial distribution of SOC. Elevated concentrations of organic matter in soils beneath and around trees and shrubs relative to surrounding soils have been observed across a variety of environments and ecosystems [32, 33]. In the Northern Tablelands of New South Wales (the region from which the data analysed here were collected) an ANOVA detected significantly elevated ($P < 0.001$) organic carbon content in the top 5cm of soils underneath the canopies of scattered paddock trees compared to soils beyond these canopy margins [33].

2.4 Soil Organic Carbon and Digital Elevation Model Derived Terrain Descriptors

Climate, parent material, topography and biotic factors may all have influenced pedogenesis to varying degrees in different ecological and geographic contexts. Topography may have influenced soil characteristics to a greater or lesser extent by having influenced hydrologic and erosional processes (e.g. soil water content, runoff and sedimentation) along with soil temperature (via aspect, exposure etc.) which together form and alter soils through mineral weathering, erosion, leaching, decomposition, horizontal zonation and sedimentation [34]. Topography may also have affected the process of SOC loss that accompanied the conversion of natural land into agricultural land by having influenced SOC: leaching, movement as dissolved organic carbon or particulate organic carbon suspended in water flowing over or through the soil, and erosion by wind or water runoff moving soil and the constituent SOC [35]. For the purposes of geo-statistical modelling, topography has often been quantified via terrain metrics (e.g. elevation, slope, aspect, curvature, etc.) and hydrological metrics (e.g. catchment area, soil wetness index, stream force index, etc.) calculated for each of the pixels in a digital elevation model (DEM) of the land surface.

In 460ha of cropping and pastoral land in north-west New South Wales, Australia regression modelling identified elevation and plan curvature along with EC_a , γ -ray radiometric potassium and thorium related emissions as useful predictors of total soil carbon [9]. In a 9ha field with a long history of row crop monoculture subject to conventional tillage in central Alabama, USA SOC was found correlated with a compound topographic index (metric of potential for water pooling on the land surface) ($r = 0.48$) and with land slope ($r = -0.42$) which lead the authors to postulate that erosion and field scale hydrodynamics were likely responsible for a large portion of the variability detected there in soil carbon [36]. In mapping soil carbon in a 12.5ha field with a history of crop rotation between corn and soy beans in central Michigan, USA models that

utilised terrain slope, aspect, plan curvature, profile curvature and tangential curvature were generally found to perform better than those that did not [37]. In 9ha of cropping soil typically used for cotton in Alabama, USA models with combined topographic index, elevation, slope, silt content and EC_a as covariates were found to account for up to 50% of the SOC variability leading the authors to conclude that the spatial distribution of SOC had been affected prominently by topography and historical erosion [38]. In this same farmland in 2006 SOC concentrations in the top 30cm of soil were found to be correlated with the composite terrain index ($r = 0.48$) and terrain slope ($r = -0.41$) [18]. In a 4.2ha catchment used as agricultural land in North Rhine-Westphalia, Germany correlations between SOC and profile curvature, plan curvature, catchment area, stream power index [34] and predictions from water and tillage erosion models [39, 40, 41] of soil redistribution patterns have been detected [42]. From a study of 5.4ha of a dryland agroecosystem with a long history of winter wheat in North-Eastern Colorado, slope and wetness index were identified as the terrain attributes most correlated with soil organic matter [34]. Variable selection in this same study returned a linear model that explained 48% of soil organic matter variation with the covariates wetness index, stream power index and aspect. Studies that detected correlations between soil carbon and a selection of topography and hydrology metrics that may be calculated with the SAGA [10] have been summarised in Table 4.

2.5 Soil Organic Carbon and Radiometric Imaging of the Earth Surface

Recording γ radiation naturally emitted from the surface of the Earth has been established as a means to detect geochemical anomalies in particular those associated with ore bodies [43]. Collecting aerial images of such γ radiation emissions has been termed γ radiometry and the spectral signatures most frequently observed have been those associated with the production of ^{238}U , ^{232}Th and ^{40}K daughter radionuclides [43]. In addition to detecting minerals rich in Uranium and Thorium and mapping geology based on prior knowledge of associations between the above radionuclides and geological materials, γ radiometry of a landscape has also facilitated the tracking of geochemical anomalies and inference regarding erosional processes therein [43]. Such links to pedological processes have enabled γ radiometry to be used for soil mapping [43]. On a broad spatial scale airborne radiometric data (particularly the K band) was found to improve the mean square error of predictions of soil organic carbon across Northern Ireland ($\sim 13,843 \text{ km}^2$) when coupled with elevation data to 30.6% [44]. Similarly, digital elevation model derived soil properties and γ radiometric survey data when used to build regression trees were found to account for 54% of the total soil carbon variation across 50, 000 ha of state forest in south eastern Australia [45]. Furthermore, γ radiometry has also been found useful for predicting the spatial distribution of soil carbon on scales closer to that across which our data were collected. Over a 5625ha square area of cropping land on the lower plains of the Macquarie River, New South Wales (Australia) percentage soil or-

Table 4: Topographic and hydrological metrics reported to have been correlated with soil carbon in agricultural land.

Metric	Correlation with Soil Carbon
Aspect	[34, 37]
Catchment Area	[42]
Elevation	[38, 9]
Plan Curvature	[34, 37, 42, 9]
Profile Curvature	[34, 37, 42]
Slope	[34, 37, 38, 36, 18]
Stream Power Index	[34, 42]
Tangential Curvature	[37]
Topographic Indices	[38, 36, 18]
Wetness Index	[34]

ganic carbon in the top soil was found to be weakly negatively correlated with the concentrations of Potassium and Uranium in the ground as measured by a γ radiometric survey [46]. Furthermore, regression modelling identified elevation and plan curvature along with EC_a , γ -ray radiometric potassium and thorium reflectances as useful predictors of total soil carbon in 460ha of cropping and pastoral land in New South Wales [9].

3 Appendix C: Choice of Modelling Method

In this appendix we compare and contrast a selection of Multiple Linear Regression (hereafter MLR) and Binary Tree (hereafter BT) based techniques in the context of data and objectives akin to those of our case study. We consider the defining characteristics of our case study data to be: (1) more potential covariate terms than observations (the $p > n$ or ultrahigh dimensional situation for variable selection) (2) a high degree of collinearity among the potential covariate terms and (3) suspected importance of non-linear effects of covariates and interactions of covariate effects. The primary objective of our case study analysis was covariate assisted spatial interpolation of the response. Our case study also had the additional context of the modest computational resources

provided by one mid-range laptop and our desire for an easily interpretable predictive mechanism. In Section 3.1 we introduce a selection of BT based models and in Section 3.2 we introduce a selection of MLR based models. In Section 3.3 we compare the relative merits of the model introduced in Section 3.1 and Section 3.2 for modelling data like those of our case study with our objective of covariate assisted spatial interpolation of the response.

3.1 An Introduction to the Binary Tree Based Models Considered in this Work

3.1.1 CART

Classification And Regression Trees (CART) [47] are two closely related techniques. Both utilise a binary tree based model structure but classification trees predict a categorical response while regression trees predict a continuous response. Our interest here is in techniques for modeling a continuous response and thus methods for solving regression problems. Subsequently we focus on regression trees. The regression trees of Breiman et al. 1984 partition the response observations $y_i, i = 1, \dots, n$, into M mutually exclusive sets by recursively partitioning the associated covariate space into M mutually exclusive regions R_1, \dots, R_M through binary divisions along covariate axes. Each associated subset of the response is then modelled by the single parameter, maximum likelihood estimator for those observations, the group mean \hat{c}_m [48]. Thus, regression trees model a continuous response as per Equation 1.

$$\hat{y}_i = \sum_{m=1}^M \hat{c}_m I(\mathbf{x}_{i,\cdot} \in R_m) \quad (1)$$

Each y_i will be modelled by one and only one \hat{c}_m since the R_m are disjoint and thus for any particular i , $\mathbf{x}_{i,\cdot} \in R_m$ for exactly one unique R_m .

Thus when continuous covariates are supplied to regression trees the predictions produced vary in a stepwise manner across the range of the covariates whereas the predictions from an MLR supplied with these same covariates would vary in a continuous manner across this same range. Furthermore, the recursive nature of the binary partitions of regression trees enable far more complex interactions to be modelled than those permitted by taking products of pairs of covariates as may be done in MLR. However, deep trees (trees with many binary partitions) would be required to form good approximations to even simple linear relationships between covariates and the response whereas such relationships may be modelled naturally as components of the structure of MLR. Deep trees rapidly become a concern with a limited number of response observations since they result in terminal node parameters being estimated from fewer observations (and thus less reliably).

Frequentist CART Regression trees, as proposed by Breiman et al. 1984, are constructed via recursive binary partitions of the observations based on whether

particular covariate values of these observations exceeded threshold values with the range of these covariates. This results in the M mutually exclusive regions of covariate space R_1, \dots, R_M referred to in Equation 1. Within each particular R_m the \hat{c}_m that minimises the residual sum of squares for predicting the response in that region is simply the group mean for those response observations $\hat{c}_m = E(y_i | x_i \in R_m)$. The challenge when fitting regression trees is identifying a sequence of binary partitions that define a set of regions such that the residual sum of squares from the entire tree is low. Since an exhaustive search of the potential regression trees that may be constructed from a particular set of data is computationally infeasible in all but the most trivial cases, regression trees are typically constructed via a greedy algorithm in the hope of identifying a regression tree that fits the data well via a computationally tractable procedure [48]. Greedy algorithms are so named because at each step in the iterations of such an algorithm the decision that yields the best improvement in the decision metric (e.g. fit of the model etc.) between the current state and the next is the decisions that is taken. As such there is no long term planning or stochasticity involved in such algorithms and thus no guarantee of identifying the optimal solution. Indeed, such algorithms can be seen to be highly sensitive to local maxima.

For fitting regression trees using the residual sum of squares as the criterion for decisions to partition the data, the greedy algorithm is as follows. The algorithm commences with all data in a single set termed the root node of the tree. All possible binary divisions of this root node set along all possible covariate axes are then constructed in turn, and the residual sums of squares resulting from prediction of the response with the pairs of associated group means are computed. With a finite number of observations there is a finite number of ways to divide the data into two subsets based on the covariate values of these observations. This complete but finite set of possible divisions of the data may be obtained by considering threshold values equidistant between each of the pairs of observed values of a covariate when these values are arranged in an ascending (or descending sequence) for each covariate in turn. The binary partition that yields the best improvement in the residual sum of squares for the entire tree relative to that at the previous step is then selected as the partition to use at this step in the algorithm and the process is repeated for each of the resulting subsets (also referred to as child nodes). This recursive partitioning process is then continued until some stopping criterion is met. A simple and popular choice of stopping criterion is a minimum number of observations per terminal node from which the practitioner considers it is still reliable to estimate a mean. Once this tree growing algorithm is halted by the satisfaction of the selected stopping criterion the resulting tree may then be ‘pruned’ by sequentially examining the effects of collapsing the parent nodes of the current terminal nodes on the basis of some *cost-complexity* criterion and taking this action where it is judged meritorious. More formally this growing and subsequent pruning procedure may be described as follows.

The algorithm commences at the root node which contains all observations. Given the full set of covariates as the design matrix, X , the algorithm takes each covariate, X_j , in turn and computes the set of threshold values that would each produce different subsets of the response should they be used to define a binary partition of the data based on this covariate. For each unique pairing of a particular covariate, X_j , and a particular threshold value for that covariate, s , the regions defined by a binary partition based on X_j and s are two disjoint regions of covariate space: $R_1(j, s) = \{X|X_j \leq s\}$ and $R_2(j, s) = \{X|X_j > s\}$. The algorithm compares all such regions that may be constructed at this step to identify the choice of covariate X_j and threshold value s that solves

$$\arg \min_{j,s} \left(\arg \min_{c_1} \left(\sum_{x_{i,j} \in R_1(j,s)} (y_i - \hat{c}_1)^2 \right) + \arg \min_{c_2} \left(\sum_{x_{i,j} \in R_2(j,s)} (y_i - \hat{c}_2)^2 \right) \right) \quad (2)$$

where \hat{c}_1 and \hat{c}_2 are calculated for each pairs (j, s) as the respective response group means: $\hat{c}_1 = E(y_i|x_{i,j} \in R_1(j, s))$ and $\hat{c}_2 = E(y_i|x_{i,j} \in R_2(j, s))$. The resulting binary partition of the data is then made and the above process is repeated for each of the resulting child nodes until some stopping criterion, such as a threshold minimum number of observations per terminal node, is satisfied at which point the algorithm is halted. The resulting tree may then be pruned recursively subject to the changes in some *cost-complexity* criterion that result from collapsing the parent nodes of the various current terminal nodes. More formally, let: $T \subset T_0$ be defined as any tree that may be obtained by collapsing some non-terminal node(s) of T_0 , $|T|$ be defined as the number of terminal nodes in tree T and let m index the terminal nodes of T each with the associated subset of the data R_m . A *cost-complexity* criterion, $C_\alpha(T)$, may be defined as per Equation 6.

$$N_m = \text{count}\{x_i \in R_m\}, \quad (3)$$

$$\hat{c}_m = \frac{1}{N_m} \sum_{x_i \in R_m} y_i, \quad (4)$$

$$Q_m(T) = \frac{1}{N_m} \sum_{x_i \in R_m} (y_i - \hat{c}_m)^2 \quad (5)$$

$$C_\alpha(T) = \sum_{m=1}^{|T|} N_m Q_m(T) + \alpha |T|, \alpha \geq 0 \quad (6)$$

. This *cost-complexity* criterion is the sum of the residual sum of squares for the predictions from the whole tree and a multiple, α , of the number of terminal nodes in the tree. The tuning parameter α controls the trade off between the fit of the tree to the data and the complexity of the tree as quantified by the number of terminal nodes of the the tree. Increases in α will yield smaller trees with larger residual sums of squares. Each pair of original tree T_0 , grown as above, and tuning parameter value α will have some smallest sub-tree T_α that minimizes $C_\alpha(T)$. This T_α may be identified by *weakest link pruning* [48]

whereby the internal nodes that yield the smallest per-node increase in the residual sum of squares $\sum_{m=1}^{|T|} N_m Q_m(T)$ when collapsed are sequentially collapsed until there are no longer any binary partitions to collapse as the entire data are once again contained in the original ‘root’ node associated set. It has been shown that the sequence of sub-trees thus obtained dependably contains T_α [48, 47, 49]. For the purposes of building a regression tree for interpolation, α may be estimated via cross validation [48]. This is traditionally how regression trees have been fitted however there is also an approach extant for fitting regression trees under the Bayesian paradigm.

Bayesian CART The Bayesian approach to fitting CART models was developed by Chipman et al. (1998) and involves the use of particular prior specifications for the terminal node parameters and the tree structure itself along with a stochastic search. The Bayesian CART [50] is a CART model, it is the manner in which the model is fitted and the underlying assumptions that are Bayesian. Continuing our focus on models for a continuous response variable with continuous covariates we will describe the Bayesian approach to fitting a regression tree. The Bayesian regression tree model consists of a binary tree T with b terminal nodes and the associated parameter vector $\Theta = (\theta_1, \theta_2, \dots, \theta_b)$, each θ_i being associated with the i th terminal node of the tree. If $\mathbf{x}_{i,\cdot} = [x_{i,1}, x_{i,2}, \dots, x_{i,p}]$ falls within the region defined by the i th terminal node, the associated response variable $y_i|\mathbf{x}_{i,\cdot}$ is modeled by the distribution $f(y_i|\theta_i)$ with f representing some parametric family controlled by parameter(s) θ_i . Chipman et al do not specify a closed form prior for the binary tree structure but instead specify it implicitly by generating trees from a tree growing stochastic process. In this manner each tree grown by the stochastic process forms a randomly drawn observation from this tree prior. Such specification of the tree prior allows for simple evaluation of prior probability $p(T)$ for any tree T which in turn may be employed within a Metropolis Hasting (MH) algorithm. To draw an observation from the tree prior a new tree is propagated from the tree consisting of a single ‘root’ node by stochastically dividing terminal nodes in an iterative process. This tree propagating process is governed by a function that controls the selection of terminal nodes for division and another function that controls the assignment of a division rule to a terminal node that has been selected to be divided. The function $p_{DIVIDE}(\eta, T)$ generates the probability for tree T that terminal node η is divided. If terminal node η of tree T is chosen to be divided, the function $p_{RULE}(\rho|\eta, T)$ generates the probability of assigning the division rule ρ to this terminal node. A division rule for a binary partition specifies the covariate that will be used to define the binary partition and the threshold value of this covariate which will determine the division of the observations into two groups based on the values of this covariate associated with each of the observations. The form and parameter values assigned to these functions collectively control the frequency with which particular tree depths and geometries are generated and thus the eventual weighting of such trees in the prior distribution. As such,

via influence on the prior, these functions may be used to guide the posterior towards identifying trees of the desired depths and geometries (for instance to emphasise trees with a minimum number of observations in each terminal node). The priors for the parameters of the distribution functions used to model the response observations in the terminal nodes may be taken as standard conjugate forms. A convenient option for priors on the parameters $\Theta = (\theta_1, \theta_2, \dots, \theta_b)$ for the terminal nodes $1, \dots, b$ is to use mean shifted normal distributions for each θ_i with $(\mu_1, \mu_2, \dots, \mu_b) | \sigma, T \text{ iid} \sim N(\bar{\mu}, \frac{\sigma^2}{\alpha})$ and $\sigma^2 | T \sim IG(\frac{\nu}{2}, \frac{\nu\lambda}{2})$. Such a formulation permits each terminal node an individual mean parameter and models all node mean parameters with independent and identical Normal distributions. Should a more flexible formulation be desired an individual variance parameter may be introduced for each terminal node mean via mean-variance shifted normal distributions as follows i.e. for $i = 1, 2, \dots, b$, $\mu_i | \sigma_i, T \text{ iid} \sim N(\bar{\mu}, \frac{\sigma_i^2}{\alpha})$ and $\sigma_i^2 | T \sim IG(\frac{\nu}{2}, \frac{\nu\lambda}{2})$. Both the above the priors $p(\Theta|T)$ facilitate closed form solutions for $p(Y|X, T) = \int p(Y|X, \Theta, T)p(\Theta|T)d\Theta$ to be obtained analytically. Utilization of one of these closed form solutions along with a CART tree prior $P(T)$ enables the posterior of T to be obtained subject to a normalizing constant: $p(T|X, Y) \propto p(Y|X, T)p(T)$. The enormous number of trees that may be constructed from all but the smallest of data will render an exhaustive search of $p(T|X, Y)$ computationally intractable. Subsequently, the normalization constant will not be obtainable nor will it be possible to identify which trees have with the highest posterior probability. However, the posterior may still be explored using an MH algorithm to conduct a stochastic search. Such a stochastic search will result in a Markov chain sequence of trees that converges towards trees with higher posterior probabilities and converges in distribution to the posterior $p(T|Y, X)$. Chipman et al. (1998) note that their MH algorithm has a tendency to move rapidly towards a group of similar trees with high posterior probability proximate to the starting tree then remain in that vicinity exploring that group of trees with small local steps for many subsequent iterations of the algorithm. Given a large enough number of iterations MH algorithms will move between posterior modes and explore the entirety of the trees possible but there is no guarantee about how many iterations a MH algorithm must be run for in order to achieve such a complete exploration. In light of these considerations, Chipman et al. (1998) recommend comparing the results of multiple runs of their MH algorithm each originating from different starting values (origin trees). Chipman et al. (1998) recommend both multiple restarts of their MH algorithm from the single root node tree, citing high initial variability in the direction in which their MH algorithm will proceed often leading such restarts to converge on quite different trees, and multiple restarts from start values selected from interesting intermediate trees from previous runs of their algorithm or trees identified by other methods (e.g. bootstrap bumping [51]). Such multiple restarts of the MH algorithm for fitting Bayesian CARTS will result in a range of selected trees which can either be model averaged or chosen between depending on the goals of the particular analysis. Aids for selecting the trees to include in the model averaging or for selecting a single tree could include the

residual sum of squares of the resulting tree (either alone or constituent in a cost complexity metric) or plots of the observed likelihood of the trees $p(Y|X, T)$ against the number of terminal nodes of the same trees as a guide to the cost - complexity compromising being struck.

3.1.2 Bagged Trees

Perhaps the simplest elaboration of CART comes from applying it within a bootstrapping procedure. The term ‘bagging’ was created by compressing the term ‘bootstrap aggregation’ and refers to taking an average of the set of predictions obtained from applying the same model fitting procedure to a collection of bootstrap samples of some data. Bagged Trees are reviewed in Hastie et al. (2009) who make the following observations regarding properties of Bagged Trees that would be pertinent to the application this technique to data in which linear and non-linear effects of covariates are expected to be important and collinearity is extant among covariates. The average of many regression trees fitted to bootstrap resamples can better approximate linear and non-linear trends in the data than a single regression tree. This stems from how the average of many such stepwise approximations to a linear (or non-linear) relationship, many of which will differ slightly having been fitted to different bootstrap resamples of the data, the will form a much better approximation to this relationship than any of the individual constituent stepwise approximations. Collinearity among the covariates can lead to high variance among regression trees fitted to replicate data which bagging can smooth out in the hope of thereby obtaining a more generally applicable model. However bagging regression trees will not improve the bias in estimation relative to that associated with a single regression tree fit. This improvement in prediction due to reduction in variance from bagging comes at the cost of the interpretability of a single tree. This is occurs since bagged regression trees are an average of the predictions of many such binary trees that have different geometries and as such the simple binary, branching nature of a single regression tree is sacrificed.

3.1.3 Random Forests

Like bagging, random forests [52] also involve averaging the predictions of a set of binary tree based models such as regression trees. Furthermore, both bagging and random forests build this set of tree based models from a set of bootstrapped samples of the data. Random forests elaborate on bagged trees by building ‘de-correlated’ trees. These de-correlated trees are propagated by choosing each binary division based on a randomly selected subset of the potential covariates. Repeating this tree propagating process many times, each on a different bootstrap resample of the data results in a set of ‘de-correlated’ trees that are then averaged to obtain a final prediction. As such random forests share the advantage of bagged trees over a single regression tree in that they may better approximate linear and non-linear trends in the data while still being able to model complex interactions between covariates. Random forests frequently

perform synonymously to boosted trees and are easier to train and tune [48].

3.1.4 Boruta All Relevant Variable Selection

Random forests have also been taken as a starting point for further methodological elaborations and refinements. One such technique is Boruta all relevant variable selection [53], hereafter BARVS (our acronym). The essence of the BARVS method is the recursive process of fitting a random forest then assessing which of the covariates utilised in this random forest made a sufficient contribution to the predictive performance of this random forest to warrant retention. The random forest is then refitted, this time using only the covariates deemed worth retaining in the previous iteration. This process is then repeated until a random forest is fitted in which all covariates utilised are deemed to have made sufficient contribution to the predictive performance of the random forest to warrant retention. Kursa et al. assess the contribution of a covariate to the predictive performance of the random forest via a Z score (though note that $Z \sim N(0, 1)$). The Boruta Z score for a covariate in a particular random forest relates to the loss of accuracy of prediction resulting from the random permutation of the values of that covariate among observations. The Boruta Z score for a covariate in a random forest is calculated from these losses of predictive accuracy from all the constituent regression trees that utilized that covariate. In particular, the Boruta Z score is calculated by dividing the mean of these losses in predictive accuracy by the standard deviation of these losses (hence the choice of name). The BARVS algorithm runs approximately as follows. Firstly, permuted copies of all covariates currently under consideration for inclusion in the random forest are added to this set of considered covariates and a random forest is fitted to these composite data. The Z scores are then calculated for all of covariates, including the permuted copies of the actual covariates. The maximum Z score among the permuted covariates is then identified and two sided t tests are performed to test the equality of the Z score for each actual covariate against this maximum Z score among the permuted covariates. Next, all actual covariates that have Z scores significantly less than this maximum Z score among the permuted covariates are discarded from the set of considered covariates and the current set of permuted covariates is also discarded. This procedure is then repeated until all remaining actual covariates have Z scores significantly greater than the maximum Z score among the permuted covariates created at that iteration. The random forest constructed from these exclusively relevant covariates is then retained as the final solution.

3.1.5 Boosted Trees

Boosting is a well regarded technique that may be applied to a variety of models including binary tree based models [48]. Boosted regression trees have been introduced in a manner accessible to quantitative scientists by Elith et al. (2008). Akin to bagged trees and random forests, boosted trees incorporate a sequence of binary trees (such as regression trees) fitted to sequentially modified versions

of the original data, the final prediction being drawn from a combination of the predictions of this sequence of trees. In the case of boosted trees, the sequential modifications to the original data take the form of weightings calculated from the results of the tree fitted in the previous step. Boosted regression trees implement a type of functional gradient descent designed to minimise a loss function that quantifies the loss in predictive accuracy resulting from an imperfect model fit [54]. The first tree is fitted to the original data by maximising the reduction in the value of the loss function relative to that from a single node tree. The second tree is fitted to the residuals from the first tree but the predictions from this second step in the fitting process are the result of combining the predictions from the first and the second tree. The third tree is then fitted to the residuals from the combined predictions of the first and second trees and so on. In this manner boosting sequentially focuses the model fitting on the observations that are difficult to explain. Boosting is typically conducted for as many iterations as is computationally feasible and the final prediction from the ensemble of boosted regression trees is a type of weighted average of the predictions from all the constituent trees. Predictions from boosted regression trees have increased stability and accuracy compared to those from a single regression tree model. Furthermore, the introduction of some stochasticity into the boosting algorithm via the inclusion of a bagging step can further improve the accuracy of the predictions and mitigate the effects of overfitting [55]. Boosted trees like bagged trees and random forests involve the sacrifice of the interpretability of a single regression tree for better predictive performance. However, a relatively straight forward metric of covariate importance exists [56, 57] which scores covariates based on the frequency with which each defined a binary division and weights these scores proportionally to the improvement in the model fit that resulted from the inclusion of the associated division. These scores are averaged across all the trees in the boosting sequence and the resulting scores are scaled to sum to 100 for ease of interpretation. Unlike a single regression tree, boosted regression trees can easily model linear relationships, non-linear relationships and relationships that include step-like discontinuities [54]. Unlike bagged trees and random forests, boosted regression trees reduce both the variance and bias associated with predictions [54].

3.1.6 Cubist

Cubist (<https://www.rulequest.com/cubist-info.html>) fits predictive models developed from Quinlan's M5 model tree method [58]. Quinlan's M5 method functions by creating a binary tree structure which is then pruned to reduce tree complexity without greatly reducing the overall fit of the tree to the data. Where this pruning converts former interior nodes of the tree into terminal nodes by collapsing the tree structure below them, MLR models are fitted to each of the subsets of the data thus defined. Each MLR model selects covariates from the set of covariates that were previously used to define the tree structure that was below (child nodes of) this current node. The Cubist method extends the M5 method by incorporating a boosting step.

3.1.7 Bayesian Treed Regression

The Bayesian implementation of CART [50] has been extended to fit MLR models (rather than simple intercept only models) to the subsets of the data defined by the terminal nodes of the associated binary tree in a framework the creators dubbed Bayesian treed regression [59]. One motivation for such a model formulation being the scenario whereby different covariates are most useful for predicting the response in different subsets of the data. The end result is somewhat akin to C5 model trees [58] but the method of attaining such a fit is very different, the model being formulated under the Bayesian paradigm. Much like Bayesian CART, Bayesian treed regression utilises an implicitly specified tree prior and a stochastic search. The Bayesian formulation and stochastic search enable different tree geometries and MLR models in the terminal nodes of these trees to be explored in addition to different error variances to be modelled for each terminal node. The particulars of the model formulation are as follows. For each terminal node of the tree T , the response observations, Y , that are members of this node are assigned a parametric model. In this manner there is a separate parametric model for the response observations contained in each of the unique terminal nodes of the tree. In particular, the distribution of the elements of the response vector \mathbf{Y} that are members of the i^{th} terminal node of the tree T are modeled conditional upon the associated covariate values by the parametric model $Y|x \sim f(y|x, \theta_i)$ where $\Theta = (\theta_1, \dots, \theta_b)$. This contrasts with a Bayesian CART where the distributions of the response observations in each terminal node are not modeled as conditional upon the associated observations of the covariates x there. Where CART and Bayesian CART models utilise a collection of stepwise functions to approximate correlations between the response and the covariates via the binary tree structure Treed regressions may be thought of as a collection of piecewise MLR models. As such Bayesian treed regression models would model linear or non-linear relationships between covariates and the response much more parsimoniously than Bayesian CART models as Bayesian treed models have the facility to invoke MLR in the terminal nodes. This facilitates the transfer of complexity in Bayesian treed models from the tree geometry to the terminal node MLR models. As such, one may reasonably expect shallower more readily intelligible trees to be coupled with the terminal node MLR models in Bayesian treed regression as compared to the more deeper, less readily intelligible trees that are coupled with terminal node single parameter models in Bayesian CART. This may be thought of as the binary tree component of the Bayesian treed models capturing the major subsets of the data and the associated MLR models describing the nuances thereof.

Each fit of a Bayesian treed model is uniquely defined by (Θ, T) and Chipman et al. (2002) outline how the posterior distribution of these (Θ, T) may be explored via stochastic search with the aid of a Metropolis Hastings algorithm. As with their Bayesian CART, Chipman et al. (2002) recommend comparing the fits to which multiple restarts of the stochastic search converge so as to better explore the posterior rather than running a single stochastic search for

a long time. They make this recommendation in light of the noted tendency of their MH explorations to converge on local maxima in the posterior then explore the vicinity of that maxima for many iterations. Chipman et al. (2002) note that their approach is amendable to both Bayesian model selection, should a single model be desired, and Bayesian model averaging (e.g. by posterior or likelihood weighting of the iterations of the stochastic search) should better predictive performance be desired. When model averaging is elected, Chipman et al. suggest model averaging only the better fitting models.

3.2 An Introduction to the Multiple Linear Regression Based Models Considered this Work

Multiple Linear Regression (MLR) predicts observations of the response y_i , $i = 1, \dots, n$, from a linear combination of products of observations of the covariates $x_{i,j}$, $j = 1, \dots, p$, and the associated coefficient estimates $\hat{\beta}_j$ plus an intercept term $\hat{\beta}_0$ as per Equation 7:

$$\hat{y}_i = \hat{\beta}_0 + \sum_{j=1}^p \hat{\beta}_j x_{i,j} \quad (7)$$

In MLR, non-linear effects of covariates upon the response may be modelled via inclusion of additional covariate terms formed as single term polynomials constructed from the original covariates. Interactions between covariates in their effects upon the response may be modelled via including in the regression additional covariate terms constructed by taking products of the covariates constituent to the interaction in question.

3.2.1 MLR Maximum Likelihood

When a MLR model is fitted by ordinary least squares (OLS) likelihood maximisation the vector of coefficient estimates $\hat{\beta}$ is obtained from the design matrix X as per Equation 8.

$$\hat{\beta} = (X^T X)^{-1} X^T \mathbf{y} \quad (8)$$

As such ordinary least squares cannot estimate MLR fits for scenarios where the number of covariates, p , exceeds the number of observations, n , (the $p > n$ or ultrahigh dimensional scenario). Even in situations where the number of covariates is less than the number of observations, consideration of non-linear terms for each covariate and the $\binom{p}{k}$ possible order k interaction terms can lead to the number of considered covariate terms grossly exceeding the number of observations and OLS fitting no longer being possible. Furthermore, fitting MLR by OLS is ill advised when collinearity exists among the covariates [60].

3.2.2 MLR with Information Criterion Based Variable Selection

When the number of covariate terms desired to be considered for use in predicting the response exceeds the number of observations, OLS is often incorporated

into a model comparison framework that fits and compares numerous models that utilise subsets of the available covariates such that $p < n$. This comparison is often effected via an information criterion such as the Akaike Information Criterion (AIC) [61, 62, 63] which assigns models a score that rewards goodness of fit to the data while penalizing model complexity. Where it is computationally feasible to do so all possible models that may be constructed from a particular set of covariates with a particular number of observations may be fitted via an exhaustive search procedure and compared, for instance in terms of information criterion values. Where the computational burden of an exhaustive search is deemed too great, a popular choice is to adopt some greedy algorithm based approach to search for optima in the information criterion values that accompany the set of possible models without having to fit all these models. Stepwise variable selection methods are examples of such techniques. It should be noted that in their base forms both stepwise variable selection and exhaustive searches still utilise OLS fitting procedures and as such are both ill advised in the presence of collinearity among the covariates.

3.2.3 MLR Bayesian

The Bayesian approach may also be used to fit MLR models (see for example Gelman et al. 2004). Furthermore, the Bayesian framework may be used to accomplish variable selection under the ultrahigh dimensional scenario by creating prior distributions that give each regression coefficient a high probability of being zero [64] as performed in the Bayesian variable selection method Spike and Slab priors [65, 66, 67]. A simple Bayesian formulation of an MLR with independent and non-informative priors on all coefficients may be adversely affected by collinearity among covariates. Collinearity among covariates will lead to high posterior variance of the associated coefficient estimates [64] and subsequently slow Markov Chain Monte Carlo convergence. Bayesian analogues to shrinkage techniques for mitigating the undesirable effects of modeling with data that includes collinearity among the covariates are outlined along side their penalized likelihood based counterparts in the following section.

3.2.4 MLR Penalization / Shrinkage

One of the dangers when conducting MLR based modelling with covariates among which collinearity exists is that highly correlated pairs of covariates can be assigned arbitrarily large magnitude positive and negative coefficients that effectively negate each other. One action that may be taken to mitigate or at least control this effect is to impose a penalty upon the combined magnitude of the coefficients as part of the fitting process [48]. A simple choice for such penalty functions is to take the L_γ norm of the regression coefficient vector β , for some value of γ , and search for the $\hat{\beta}$ that minimizes the sum of the residual

sum of squares and this norm of the regression coefficients as per Equation 9.

$$\hat{\beta}^{L_\gamma} = \arg \min_{\beta} \left\{ \sum_{i=1}^n (y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j)^2 + \lambda \sum_{j=1}^p |\beta_j|^\gamma \right\}, \quad \gamma > 0 \quad (9)$$

Solving Equation 9 with $\gamma = 2$ yields a Ridge Regression estimate equivalent to that obtained by solving Equation 10.

$$\hat{\beta}^{\text{ridge}} = \arg \min_{\beta} \left\{ \sum_{i=1}^n (y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j)^2 + \lambda \sum_{j=1}^p \beta_j^2 \right\} \quad (10)$$

Ridge regression shrinks all coefficients towards zero and thus does not perform variable selection and is not useful for regression in the ultrahigh dimensional situation. Interestingly, ridge regression is equivalent to Bayesian MLR with an exchangeable normal prior distribution on the coefficients [64]. Using $\gamma = 1$ in Equation 9 yields an L_1 penalized least squares estimate also known as the Least Absolute Shrinkage and Selection Operator (LASSO) [68]. More complex choices for the penalty function than the L_γ norm in Equation 9 are used in penalized least squares techniques such as adaptive LASSO [69], Smoothly Clipped Absolute Deviation (SCAD) [70] and Minimax Concave Penalty (MCP) [71]. Of these techniques L_γ penalization is perhaps the most easily applied. Solving Equation 9 for any $\gamma < 2$, will shrink the coefficient estimates for some covariates to zero exactly (how many depends on value of tuning parameter λ) thereby performing variable selection in addition to penalized estimation [72]. As such L_γ penalized estimation with $\gamma < 2$ is applicable in scenarios where where the number of potential covariates exceeds the number of observations ($p > n$) and collinearity exists among the covariates. The absolute value in an L_1 penalized estimate requires a computational solution initially provided by quadratic programming [68] and more recently by the more computationally efficient Least Angle Regression (LAR) algorithm [73]. An estimate of the LASSO solution may also be obtained from the posterior mode estimates of a Bayesian MLR with independent and identical, Laplace (double exponential) priors on the regression coefficients [74].

3.3 Comparing Multiple Linear Regression Based Techniques and Binary Tree Based Technique

The comparison of the properties of Multiple Linear Regression (MLR) and Classification And Regression Tree (CART) has most relevance to this work as a foundation to inform the comparison of the various modifications of each technique that would have better suited the coupling of the particular characteristics of the data from our case study and our objective for the analysis of these data. We consider the defining characteristics of our case study data to be: (1) more potential covariate terms than observations (the $p > n$ or ultrahigh dimensional situation for variable selection) (2) a high degree of collinearity among the potential covariate terms and (3) suspected importance of non-linear effects

of covariates and interactions of covariate effects. The primary objective with our case study analysis was the construction of a model for covariate assisted interpolation of the response. Our case study also had the additional context of the modest computational resources provided by one mid-range laptop and our desire for an easily interpretable predictive mechanism.

A plethora of BT based approaches exist today that are variously modifications of and elaborations upon the Classification And Regression Trees [47] (hereafter CART) framework. We confine ourselves here to considering a subset of these that appeared appropriate for the case study data and objectives namely: Bayesian CART [50], bagged regression trees [75], random forests [52], boruta all relevant variable selection [53] (an elaboration upon random forests), boosted regression trees [55], cubist [58] and Bayesian treed regression [59] (an elaboration upon Bayesian CART). Readers unfamiliar with these BT based techniques are referred to Section 3.1. A similar diversity of modifications to and elaborations upon MLR exist though we confine ourselves here to considering the base form and the modification thereof that seem appropriate to the defining features of our case study data. Namely, we consider shrinkage modified MLR for its relevance to the situation where collinearity exists among the covariates with particular attention to LASSO style shrinkage for its additional relevance to the situation where the number of covariates exceeds the number of observations of the response (see Section 3.2 for an introduction to these techniques). We consider both LASSO for MLR fitted under the Bayesian paradigm [74] and LASSO implemented through likelihood penalization as fitted by the LAR algorithm [73]. We consider each of these techniques in light of the key characteristics of our case study data and our objective of building a model to interpolate the response. This allows us to narrow down the choice of methods further and make a final decision. This comparison is also summarised in Table 5.

We commence this consideration with perhaps the most widely know technique introduced above, MLR. A MLR model cannot be fitted by Ordinary Least Squares (OLS) based likelihood maximisation when the number of covariates desired to be included in the model exceeds the number of observations (see Section 3.2). The large number of potential covariate terms and suspected importance of non-linear effects and interactions of covariate effects meant that conducting exhaustive search variable selection on these data was not computationally feasible. Under such scenarios, a common option has been to apply some deterministic variable selection procedure that optimises an information criterion. Stepwise variable selection with the Akaike Information Criterion (AIC) [61, 62, 63] has been a popular choice in this regard having been available in R via the step function [11] through numerous release cycles. However, with linear regression, correlations between potential covariates (as we had in our case study data, see Fig. 2) can be a cause for concern with such techniques that rely on ordinary least squares model fitting (which in the case of linear regression is also maximum likelihood model fitting). The undesirable

effects of correlations among covariates in a MLR model have been explained in the comprehensive book [60]. While Bayesian methods would allow a MLR model to be fitted with more covariates than observations the high degree of collinearity among some pairs of covariates would render the associated pairs of coefficients poorly defined which in turn could complicate convergence of the MCMC iterations upon a fit to these data. In the presence of collinearity among covariates some form of shrinkage is advisable for preventing highly correlated covariates being assigned arbitrarily large magnitude but opposite signed coefficients that all but cancel each other out due to the high correlation of the associated covariates. Coefficient shrinkage is performed by the popular Penalized Least Squares (PLS) family of techniques [72] which forms a subset of the larger family of coefficient shrinkage themed modifications of MLR. We considered shrinkage techniques including Ridge regression [76], LASSO [68], LASSO fitted by LAR [73] and the Bayesian LASSO [74]. Of particular interest were both LASSO for MLR fitted under the Bayesian paradigm [74] and LASSO implemented through likelihood penalization as fitted by the LAR algorithm [73] since LASSO is appropriate to both the situation where the number of covariates exceeds the number of observations and the situation where substantial collinearity exists among the covariates. LASSO conducts shrinkage in a manner that shrinks some coefficients to zero exactly which in effect excludes them from the model and as such is useful for conducting shrinkage and variable selection simultaneously. A LASSO modified MLR fit may be obtained via the deterministic and computationally efficient LAR algorithm [73] which adds or removes the covariate that best optimises its criteria at each iteration. Alternatively, a LASSO modified MLR estimate may be obtained by fitting a Bayesian MLR with independent and identical Laplace priors on all the coefficients [74]. Such an exploration, utilising multiple MCMC chains with a variety of starting values, could well conduct a more thorough exploration of the possible LASSO fits of MLR models that could be constructed from the potential covariates than would be conducted by the deterministic LAR algorithm. Though we note that were non-linear effects and all possible pairwise interactions considered as potential covariate terms, the computational burden involved in obtaining convergence of all chains could be considerable simply due to the number of coefficients to be estimated.

In contrast to MLR, CART may be run in its base form on ultrahigh dimensional data given sufficient computational resources. With more covariates than observations it would be vital to enforce a minimum number of observations per terminal node at which to commence pruning but if this were done there would be no a priori barrier to the implementation of this technique on such data. Were CART models being fitted within the Bayesian paradigm the tree prior would need to be parameterised to perform a similar function guiding the posterior away from deep trees with low numbers of observations per terminal node. Furthermore, the collinearity present among covariates would not prove a barrier to the implementation of CART (here regression trees as we have a continuous response) via the standard algorithms as even small differences among

covariates would be sufficient for one covariate to define a binary partition of the response that reduces some loss metric more than the other covariate. The 63 covariates of our case study analysis, while not prohibitively excessive, would likely lead to some computational burden when fitting Bayesian CART models to these data as the stochastic searches necessary may take many iterations to converge. We postulate this slow convergence due to the collinearity among some covariates and also due to the number of covariates considered. Given the tendency of Bayesian CART to converge upon a local maxima and remain there moving locally for many MCMC iterations thereafter [50] the greater the number of chains with a diversity of starting trees that could be run the greater the conviction one could have that a good enough exploration of the posterior had been conducted to have at least identified some trees that were more than just local maxima of extremely small regions of the posterior. As such a good case exists for devoting significant computational resources to fitting Bayesian CART. Suspected interactions of covariates and non-linear effects of covariates upon the response are a more interesting issue in the context of regression trees. Regression trees by their nature can model complex interactions between covariates and this is a strength they possess relative to MLR. However, as a single regression tree is essentially a collection of stepwise predictions for the response across mutually exclusive partitions of covariate space, deep trees will be necessary to approximate even a simple linear relationship between a covariates and the response well. The limited number of response observations available in our case study and large number of covariates of which several or more may well have non-linear relationships with the response means that trees of sufficient depth to describe multiple non-linear relationships would not be possible with sufficient observations in each terminal node to formulate reliable estimates of the response means there. Subsequently, some form of averaging of multiple trees would be advisable as the average of many slightly different stepwise functions can well approximate a linear or non-linear relationship even when the constituent binary trees are quite shallow. The trees thus averaged could be the trees that scored above some threshold posterior probability or trees fitted to multiple re-samples or re-weightings of the full set of observations.

Fitting regression trees to numerous bootstrap re-samples of the data would be the simplest option for generating multiple regression trees to model average in order construct better approximations to linear and non-linear relationships between covariates and the response than would be possible with a single binary tree. This technique is referred to as bootstrap aggregated trees or bagged trees for short [75]. Furthermore, averaging the predictions from numerous shallow trees each fitted to a different bootstrap re-sample of the data could also allow the various linear and non-linear relationships between different covariates and the response to be incorporated into the model despite the small sample size. If different covariates were most important to different subsets of the response, different trees could result from fitting a regression tree to each bootstrap re-sample and thus this diversity of relationships could be incorporated into the bootstrap aggregation. In this manner, even though the response sample size

limits the depths of trees possible, numerous linear and non-linear relationships between covariates and the response could be incorporated into the predictions through the different shallow trees constituent to the model averaging. As such, bagged trees could incorporate approximations of far more linear and non-linear relationships into predictions of the response than it would be possible to utilise via a single binary tree fitted to the same data. As predictions from bagged trees are constructed from the predictions of numerous individual regression trees the same ability to operate on data with high collinearity among the covariates exists for bagged regression trees as does for single regression trees. Similarly, bagged regression trees retain and even enhance the strength of single regression trees at modelling complex interactions between covariates upon their effect on the response. However, these advantages come at the cost of the easily interpretable structure of a single regression tree. Much the same assertions may be made for random forests.

Random forests [52], like bagged trees, model average predictions from a set of regression trees derived from bootstrap re-samples of the data and thus have similar advantages to bagged trees over single regression trees at approximating linear and non-linear relationships between covariates and the response. Random forests differ from bagged trees in the manner in which regression trees are fitted to the bootstrap re-sample of the data. At each binary partition in such a tree in a random forest, only a randomly selected subset of the potential covariates are made available to the algorithm to define the binary partition. Random forests may be better than bagged trees at approximating different linear and non-linear relationships between the covariates and the response with a limited number of observations due to the combination of bootstrapping and selecting from random subsets of covariates at each partition resulting in a broader range of shallow trees to aggregate than bootstrapping alone. Random forests, being composed of many individual regression trees, should not be adversely affected by collinearity among covariates and should still be able to model complex interactions.

Boruta all relevant variable selection [53] is an elaboration of the random forest method. As such, it shares the advantages of bagged trees and random forests over single regression trees for approximating linear and non-linear relationships between covariates and the response and, like all binary tree based methods, is able to model complex interactions among covariates in their effects upon the response. However, collinearity among the covariates may pose a problem for the application of boruta all relevant variable selection. The problem stems from the boruta algorithm being vulnerable to discarding useful covariates when these covariates are highly correlated with one or more other covariates. If highly correlated covariates were supplied to a random forest algorithm and these covariates contained useful information for predicting the response these covariates could well be used interchangeably throughout particular regression trees for defining binary partitions of the response. Thus the loss in predictive accuracy that would result from permuting the values of just one of these cor-

related covariates and using it in place of the actual covariate would likely be less than would result if this particular covariate had been used in all binary partitions of the tree where the partition was defined based on one of the other covariates with which this covariate was highly correlated. Subsequently, the importance of this particular covariate to the predictive accuracy of the regression tree could be underestimated to the extent that the Boruta all relevant variable selection algorithm would discard this covariate from the set of covariates worth retaining. Furthermore, if the correlated covariates in question were sufficiently correlated that they were being used interchangeably throughout the regression tree then they could all dilute the estimated importance of the others in this manner such that none of these correlated covariates were selected for retention by the boruta all relevant variable selection algorithm despite all containing much the same important information for predicting the response and thus it being well worth retaining one of them.

Boosting iteratively refits a model to data that is re-weighted at each iteration to emphasize the observations that were poorly predicted by the model fitted in the previous iteration. The predictions from this sequence of models are then combined to produce the final prediction. Where these models are regression trees the procedure is referred to as boosted regression trees [55]. Thus the predictions constructed from a sequence of boosted regression trees will have similar advantages to those from bagged trees and random forests over a those from a single regression tree model. Such aggregated predictions will better approximate linear and non-linear relationships than the predictions from a single shallow regression tree given the small number of response observations and numerous potential non-linear and interacting effects of covariates expected to exist in our data. Furthermore, by its very nature boosting will produce a variety of trees each fitted to re-weightings of the data which emphasize a different subset of the response observations. Thus, where the small sample size would hinder a single regression tree capturing the suspected importance of multiple non-linear relationships between covariates and the response, different trees in the sequence of boosted trees will describe different relationships between covariates and the response that are found to be important to observations up weighted at that iteration. Since the ensuing predictions from all such trees will then be aggregated to produce the final predictions all these different identified relationships will be combined into the predictions of the response. Again, similar to the above approaches that aggregate sequences of regression trees to produce predictions, collinearity among the covariates should not greatly hinder the regression tree fitting process constituent to fitting boosted regression trees as differences will exist among even quite collinear covariates sufficient to choose between them when they are useful for defining a binary partition of the response. Furthermore, still being based on binary trees, boosted regression trees will be able to model complex interactions between covariates in their relationships with the response.

The cubist method <https://www.rulequest.com/cubist-info.html> is in essence

an extension of the M5 or model tree approach [58] that incorporates a boosting step. As such the model that is fitted to the iteratively re-weighted data at each iteration of the boosting algorithm starts as a regression tree which is then iteratively pruned back by collapsing parent nodes to the current terminal nodes and fitting MLR models to the subsets of the data defined by these newly created terminal nodes. Each MLR model fitted to the data contained in a newly created terminal node uses as potential covariates all the covariates that defined the binary partitions that were collapsed to create this new terminal node. As such, collinearity among the covariates could lead to poorly defined coefficients in these terminal node MLRs as discussed above (assuming these covariates were also collinear in the subset of observations contained the terminal node in question) unless some shrinkage based fitting was conducted there. This model structure would allow for very flexible description of linear and non-linear relationships between covariates and the response as not only are different stepwise predictions being averaged but constituent in this averaging are also the predictions for various MLR fits. Similarly, being based upon a binary tree structure complex interactions may be modeled implicitly by this method.

Bayesian treed regression [59] is superficially similar to Quinlan’s M5 in that it fits MLR models to subsets of the data contained in the terminal nodes of a binary tree. Bayesian Treed Regression, however, is fitted under the Bayesian paradigm via stochastic search as an elaboration of a Bayesian CART model [50]. The intricate model structures permitted by fitting MLR models in the terminal nodes of binary trees is attractive for data in which complex combinations of linear, non-linear and interaction effects of covariates upon the response are suspected to be important. Some form of shrinkage would be advisable in the terminal nodes to mitigate the concerns of conducting MLR with collinear covariates. As has been discussed above LASSO may be implemented within the Bayesian framework by placing Laplace priors on the regression coefficients [74]. Our greatest concern associated with this method would be the computational burden inherent in conducting a good exploration of the posterior. Collinearity among covariates would slow the convergence of stochastic searches and the sheer breadth of possible models would require numerous chains to be run with a great variety of starting values in order to have any confidence whatsoever that a good exploration of the posterior had been conducted given the noted tendency of this algorithm to rapidly converge on local maxima in the posterior then remain in the neighbourhood of this maxima for many subsequent MCMC iterations [59, 50]. The number of parameters in the stochastic search could be reduced if one were willing to let the binary tree portion of the model be the only form of allowing for potential interactions of covariates (i.e. consideration of the $\binom{63}{2}$ pairwise interaction terms in the MLR models could then be avoided). Further reduction in the number of parameters in the stochastic search could be achieved by only allowing linear terms for each covariates in the terminal node MLRs and relying on model averaging of high posterior probability fits to account for non-linear effects through the averaging of the predictions from

multiple combinations of stepwise functions with MLR fits. Such an approach coupled with a tree prior parameterised to avoid deep binary trees would allow for a rich variety of linear, non-linear and interactions effects to be incorporated into the aggregated predictions.

Of the techniques considered, LASSO penalized MLR fitted via the LAR algorithm, random forests and boosted trees are here proposed as the techniques that were most suitable for our case study analysis (see Table 5). These three techniques: were appropriate for application to data with characteristics like those of our case study, suited our objective of building models to interpolate the response, were straight forward to implement with existing software and seemed unlikely to be adversely computationally burdensome with the computational resources available. Model-averaging of high posterior probability regression trees identified via the Bayesian CART method also appeared quite well suited to data and objectives like ours but would have required more effort to implement (tree priors require some work to parameterise so as to emphasize tree structures that maintain a minimum number of observations in all terminal nodes) and could well have proved quite computationally expensive with the computational resources available. Model-averaging of high posterior probability fits from Bayesian treed regression also appeared very promising for data and objectives like ours provided some form of shrinkage could be implemented in the terminal node MLRs and the 63 covariates were considered only as linear main effects. However, we note that the implementation of shrinkage in the terminal nodes would not be a trivial task and that both these techniques seemed likely to be quite computationally intensive with data like ours. The choice of an easily implemented technique that was appropriate to our data and computationally efficient was thus reduced to a choice between LASSO penalized MLR fitted via the LAR algorithm, random forests and boosted trees.

We have elected to use LASSO modified MLR fitted via the LAR algorithm in our case study analysis. Model-averaging the predictions from the LASSO solutions obtained from LAR executions within a cross validation scheme yielded an aggregate estimate in a manner conceptually similar to the manners in which random forests, bagged trees or boosted trees aggregate predictions from the same model structure fitted to variations on the same data set. Another useful consequence of using a cross validation based approach was that we were able to estimate the shrinkage parameter for the LASSO fits (λ in Equation 9) via cross validation. The motivation for this approach was an attempt to fit models that would perform well at interpolation. For the purposes of building models for interpolation, LASSO modified MLR fitted via the LAR algorithm was a defensible choice of method from a set of good options. Our preference for LASSO modified MLR fitted via LAR was the result of a secondary interest in which non-linear terms for covariates and which particular pairwise interactions were most useful for predicting the response. This was more readily apparent from LASSO modified MLR fits where each covariate term has a coefficient that has been either shrunk to zero, effectively excluding the covariate term

from the model, or assigned a value. When employed within a cross validation scheme this translated into frequencies of selection of these specific covariate terms which were still easy to interpret. In contrast, whether the overall role of a covariate within the aggregated estimate from random forests, bagged or boosted trees was closer to linear or non-linear (and if non-linear what manner of non-linear) would have been harder to judge from the results of such a fit. This ease of interpretability of the LASSO modified MLR came at a cost of having to recenter and rescale (to mean zero and magnitude one) all covariates in each training each set (a requirement of the LAR algorithm [73]) and mirror those transformations on each associated validation set whereas this would not have been necessary for binary tree based techniques.

Method	Data Characteristic	Method Pro	Method Con
MLR (Maximum Likelihood)	ultra-high dimensionality		insufficient degrees of freedom
	collinearity		coefficients poorly defined
	non-linearity	models directly with polynomial terms	consumes degrees of freedom
MLR (Bayesian)	interactions	models directly with product terms	consumes degrees of freedom
	ultra-high dimensionality		computationally expensive
	collinearity		coefficients poorly defined, computationally expensive
	non-linearity	models directly with polynomial terms	computationally expensive
	interactions	models directly with product terms	computationally expensive

Method	Data Characteristic	Method Pro	Method Con
MLR + Ridge Regression [76]	ultra-high dimensionality		shrinks all coefficients, does not perform variable selection
	collinearity	constrains coefficients	
	non-linearity	models directly with polynomial terms	computationally expensive
	interactions	models directly with product terms	computationally expensive
MLR + LASSO via quadratic programming [68]	ultra-high dimensionality	performs variable selection	computationally inefficient
	collinearity	constrains coefficients	
	non-linearity	models directly with polynomial terms	computationally expensive
	interactions	models directly with product terms	computationally expensive

Method	Data Characteristic	Method Pro	Method Con
MLR + LASSO (LAR) [73]	ultrahigh dimensionality	computationally efficient, performs variable selection	
	collinearity	constrains coefficients	
	non-linearity	models directly with polynomial terms	computationally expensive
	interactions	models directly with product terms	computationally expensive
MLR + LASSO (Bayesian) [74]	ultrahigh dimensionality	performs variable selection	computationally expensive
	collinearity	constrains coefficients	
	non-linearity	models directly with polynomial terms	computationally expensive
	interactions	models directly with product terms	computationally expensive

Method	Data Characteristic	Method Pro	Method Con
CART (Frequentist) [47]	ultra-high dimensionality	enforces minimum n per terminal node with stopping rule & pruning	
	collinearity	not adversely affected	
	non-linearity		requires deep tree & large n
	interactions	may describe complex interactions	
CART (Bayesian) [50]	ultra-high dimensionality	tree prior may be tuned to pull posterior towards trees with desired n per terminal node	computationally expensive
	collinearity		extends stochastic search convergence time
	non-linearity	Bayesian model-averaging could approximate well	large n for deep tree required for single tree to approximate well
	interactions	may describe complex interactions	

Method	Data Characteristic	Method Pro	Method Con
Bagged Trees [75]	ultrahigh dimensionality	stopping rule &/ pruning enforces minimum n per terminal node	
	collinearity	not adversely affected	
	non-linearity	approximates well by averaging many shallow trees	
	interactions		may describe complex interactions
Random Forests [52]	ultrahigh dimensionality	stopping rule &/ pruning enforces minimum n per terminal node	
	collinearity	not adversely affected	
	non-linearity	approximates well by averaging many shallow trees	
	interactions		may describe complex interactions

Method	Data Characteristic	Method Pro	Method Con
Boruta All Relevant Variable Selection [53]	ultrahigh dimensionality	enforces minimum n per terminal node by pruning	
	collinearity		may erroneously exclude useful covariates
	non-linearity	approximates well by averaging many shallow trees	
	interactions	may describe complex interactions	
Boosted Trees [55]	ultrahigh dimensionality	stopping rule &/ pruning enforces minimum n per terminal node	
	collinearity	not adversely affected	
	non-linearity	approximates well by averaging many shallow trees	
	interactions	may describe complex interactions	

Method	Data Characteristic	Method Pro	Method Con
Cubist [58]	ultrahigh dimensionality	pruning inherent component of algorithm	
	collinearity		poorly defined coefficients in terminal node MLRs
	non-linearity	may describe complex non-linearity	
	interactions	may describe complex interactions	
Bayesian Treed Regression [59]	ultrahigh dimensionality	tree prior tunable to pull posterior towards trees with sufficient n per terminal node for terminal node MLRs	computationally expensive stochastic search
	collinearity		poorly defined coefficients in terminal node MLRs, slow convergence of stochastic search
	non-linearity	Bayesian model-averaging could approximate well with shallow trees as could polynomial terms in terminal node MLRs	computationally expensive
	interactions	may describe complex interactions	

Method	Data Characteristic	Method Pro	Method Con
Abbreviations: n = number of observations, MLR = Multiple Linear Regression, CART = Classification And Regression Tree LASSO = Least Absolute Shrinkage and Selection Operator, LAR = Least Angle Regression			

4 Appendix D: Design Matrix Filtering

It was apparent from the correlation matrix presented as the raster plot in Fig. 2 that substantial collinearity existed among the 2205 potential covariate terms. We wished to create subsets of the full design matrix in order to explore less collinear sets of potential covariates for two reasons. Firstly, exploring design matrices which included highly collinear pairs of covariates seemed unnecessary. This was because a variable selection algorithm would have selected a covariate from each highly correlated pair of covariates on the basis of minute differences between the covariate values at the locations at which the response was observed and thus this decision could have been adversely influenced by errors in measurement or interpolation. Secondly, the variable selection functions in the version of the ‘leaps’ package [77] we used required that the design matrices explored included no pairs of covariates with correlations coefficient magnitudes greater than 0.4.

The covariates derived from the All Terrain Vehicle (ATV) surveys had the finest spatial resolution of all the covariates considered in our case study. In an effort to build models that would have predicted the response with the greatest spatial accuracy, when faced with highly correlated pairs of covariates we chose to retain the covariates collected by the ATV surveys over any others. Of the ATV survey derived covariates: visible Red reflectance (RED), Near InfraRed reflectance (NIR) and soil apparent electrical conductivity (ECA) only ECA had no other covariate terms derived from it while all the vegetation indices were calculated as functions of the RED and NIR reflectance values. For this reason ECA was retained over NIR, RED and any other highly correlated covariate term. As the vegetation indices were theoretically more indicative of green biomass than raw RED or NIR reflectance, and thus potentially more closely related to SOC levels (see Appendix B), vegetation indices were retained over the raw reflectance values were any such pairs overly correlated. Next in the order of detail of spatial resolution were the Foliar Projective Cover (FPC) Layers. We obtained two data such layers: the projected foliage cover for 2011 (FPCI) and the projected foliage cover for 2012 (FPCII). Since 2011 was less temporally removed from the 2009 soil survey than 2012, FPCI was set to be preferentially retained over FPCII or any other highly correlated covariates. The coarsest spatial resolution data were derived from the Digital Elevation Model (DEM). These data included elevation along with terrain and soil hydrology metrics calculated from the elevation. We considered that these terrain and soil hydrology metrics came closer to describing landscape processes that may have influenced SOC formation, mineralization and or transport and thus the spatial distribution of SOC levels. Subsequently, we elected to retain terrain and hydrology metrics over elevation should elevation have been highly correlated with any of these metrics. Any remaining pairs of highly correlated covariates were then chosen between at random. Once this hierarchy of filtering operations had been applied to the 63 potential covariate terms the remainder was expanded to include all remaining covariate terms to polynomial order four and all pos-

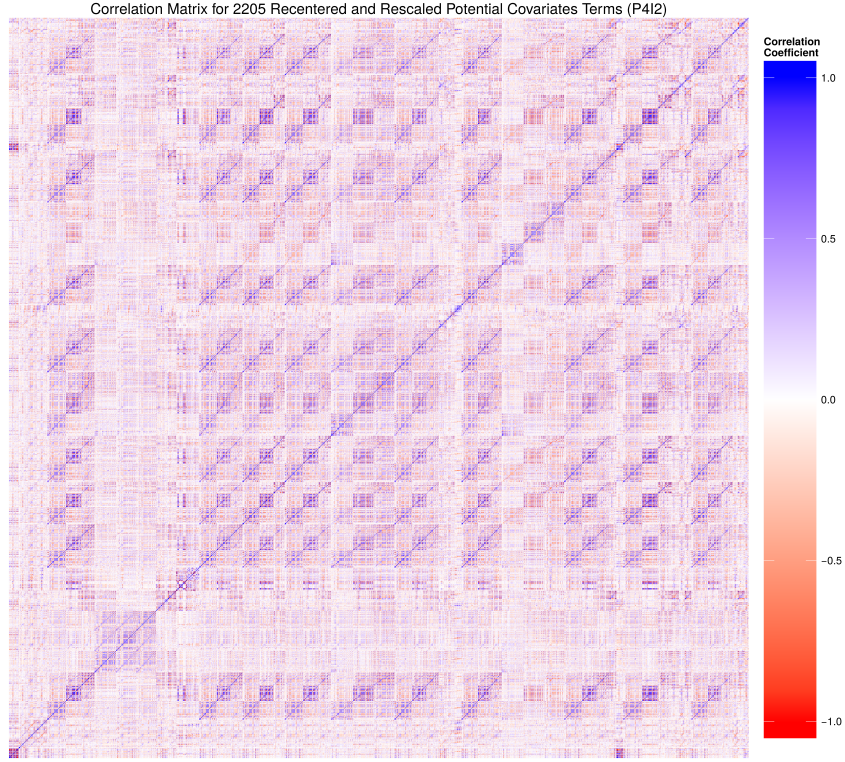


Figure 2: A plot of the matrix depicting the correlation coefficients between all possible pairs of the 2205 potential covariates terms.

sible interactions between pairs of linear terms for these remaining covariates. In the spirit of Occam's razor when searching the expanded design matrix for correlated pairs of covariate terms, single term polynomial terms were set to be retained in preference to any interaction terms with which they were found to be highly correlated. Finally, lower order polynomial terms were set to be retained in preference to any higher order polynomial terms with which they were highly correlated. Once all the above heuristics had been implemented in the order described a selection was made from any remaining pairs of highly correlated covariates at random to complete the enforcement of a maximum permitted correlation coefficient magnitude between covariates in the filtered design matrix.

5 Appendix E: Choice of Training Set Size and Design Matrix Filtering Austerity

We compared the results of using three different pairs of training and validation set sizes combined with each of four different levels of austerity in the design matrix filtering repeating the variable selection and model-averaging routine for each. The training set sizes compared were 35, 45 and 55. Sets of 500 unique training sets of each of these sizes were constructed from design matrices filtered to enforce maximum correlation coefficient magnitudes between remaining covariate pairs of 0.95, 0.8, 0.6 and 0.4. Given our primary objective of interpolating the response between the soil core observations our focus was on out of sample predictive accuracy. The metrics for out of sample prediction accuracy we adopted were the summary statistics for the distributions of the Validation Set Element Prediction Error (VSEPE) absolute values. The distributions of the absolute values of the VSEPE and the coefficients of determination for the model-averaged predictions from models selected from each of these combinations of training set size and design matrix filtering austerity have been summarized in Table 5. For each level of austerity in filtering the design matrix the distribution of the absolute values of the VSEPE appeared to become more compressed towards zero with increasing training set size. However, it should be noted that the number of validation set elements that were predicted under the scenarios where 500 training sets of 35 observations were used was much larger compared to the number predicted under scenarios where 500 training sets of 55 observations were used ($500 * (60 - 35)$ compared to $500 * (60 - 55)$). As such we only recommend comparing the VSEPE distributions obtained from each of the four collections of 500 training sets of the same size that were constructed from design matrices subjected to different the levels of filtering austerity considered.

The final model for each training set was selected from the sequence of models for that training set returned by the LAR algorithm as the model which maximised the predictive accuracy on the associated validation set. As the 35 observation training sets were the smallest of the three sizes of training sets considered, the models selected for these training sets would have had out of sample predictive accuracy most emphasised in this second stage of their selection process where the shrinkage parameter was selected to minimise validation set predictive error. Since areal interpolation of the response from full cover covariate observations via these models selected by LAR is in essence out of sample prediction, we elected to use the results of variable selection on one of the collections of 500 unique 35 observation training sets for this areal interpolation. Of the models selected from the collections of 500 training sets of 35 observations it was those constructed from the design matrix filtered to enforce a maximum permitted correlation coefficient magnitude between covariate pairs of 0.95 that had the first three quarters of the ordered VSEPE absolute values most compressed towards zero. Furthermore, of the scenarios involving 35 obser-

Table 5: Summary statistics for validation set element prediction error distributions and coefficients of determination for model-averaged predictions from each application of Least Angle Regression to a scenario defined by the combination of training set size and design matrix filtering austerity.

$> r $	TSS	Min.	1st Qu.	VSEPE				MAP R^2
				Median	Mean	3rd Qu.	Max.	
0.95	35	1.332e-05	0.1482	0.3184	0.4744	0.5446	4.437	0.5963
0.95	45	5.418e-05	0.1384	0.3067	0.4667	0.5410	4.163	0.6797
0.95	55	3.734e-05	0.1113	0.2656	0.4065	0.4525	3.858	0.8403
0.8	35	5.571e-05	0.1526	0.3288	0.4835	0.5636	4.167	0.4667
0.8	45	5.746e-05	0.1462	0.3204	0.4762	0.5583	4.112	0.5046
0.8	55	0.0002129	0.1163	0.2916	0.4220	0.5026	3.885	0.6284
0.6	35	1.119e-05	0.1495	0.3362	0.4957	0.5769	4.206	0.2796
0.6	45	0.0002723	0.1527	0.3426	0.5065	0.5916	4.184	0.3844
0.6	55	3.376e-05	0.1300	0.2914	0.4193	0.5232	3.875	0.4994
0.4	35	1.097e-05	0.1517	0.3324	0.4776	0.5695	4.063	0.3666
0.4	45	0.0003481	0.1564	0.3333	0.4760	0.5662	3.819	0.4507
0.4	55	7.387e-05	0.1101	0.2388	0.3626	0.4126	3.684	0.5593

Abbreviations:

TSS = Training Set Size,

VSEPE = Validation Set Element Prediction Error,

MAP = Model-Averaged Prediction,

$> |r|$ = the maximum absolute value of the correlation coefficient between covariate pairs permitted to remain in the design matrices supplied to the variable selection algorithms,

R^2 = the coefficient of determination,

Min. = Minimum,

1st Qu. = First Quartile,

3rd Qu. = Third Quartile,

Max. = Maximum.

vation training sets it was the model-averaged prediction from the models fitted to the 500 training sets constructed from the design matrix filtered to enforce a maximum permitted correlation coefficient between covariates of 0.95 that had the best coefficient of determination. For these reasons, we elected to interpolate the response across the study area with the model-averaged predictions from the models selected by applying LAR to the 35 observation training sets constructed from the design matrix filtered to enforce a maximum permitted correlation coefficient magnitude between remaining covariate pairs of 0.95.

References

- [1] National Parks and Wildlife Service. Duval Nature Reserve Plan of Management. Office of Environment and Heritage, Department of Premier and Cabinet, New South Wales, Australia; 2003. November. Available from: <http://www.environment.nsw.gov.au/parkmanagement/DuvalMgmtplan.htm>.
- [2] Isbell R. The Australian Soil Classification. 2nd ed. Collingwood, Australia: CSIRO Publishing; 2002.
- [3] Garraway E, Lamb D. Delineating spatial variations in soil carbon using remotely sensed data: A New England case study. In: Sustaining Our Regions: The Engineering Challenge - Proceedings of the 2011 Regional Convention, Newcastle Division, Engineers Australia; 2011. p. 45–52.
- [4] Chan KY, Cowie A, Kelly G, Singh B, Slavich P. Scoping Paper: Soil Organic Carbon Sequestration Potential for Agriculture in NSW. Orange, NSW: New South Wales Department of Primary Industries; 2008. September.
- [5] Morris ER. Height-above-ground effects on penetration depth and response of electromagnetic induction soil conductivity meters. *Computers and Electronics in Agriculture*. 2009;68:150–156.
- [6] Lamb DW, Mitchell A, Hyde G. Vineyard trellising with steel posts distorts data from EM soil surveys. *Australian Journal of Grape and Wine Research*. 2005;11:24 – 32.
- [7] Holland KH, Lamb DW, Schepers JS. Radiometry of Proximal Active Optical Sensors (AOS) for Agricultural Sensing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*. 2012;5(6):1793–1802.
- [8] Bishop C. Mixture Models and EM. In: *Pattern Recognition and Machine Learning. Information Science and Statistics*. New York: Springer-Verlag; 2006. p. 424–428.
- [9] Miklos M, Short MG, McBratney AB, Minasny B. Mapping and comparing the distribution of soil carbon under cropping and grazing management practices in Narrabri, north-west New South Wales. *Australian Journal of Soil Research*. 2010;48:248–257.
- [10] Conrad O, Bechtel B, Bock M, Dietrich H, Fischer E, Gerlitz L, et al. System for Automated Geoscientific Analyses (SAGA) v. 2.1.4. *Geoscientific Model Development*. 2015;8(7):1991–2007.
- [11] R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria; 2015. Available from: <https://www.r-project.org/>.

- [12] Brenning A. Statistical geocomputing combining R and SAGA: The example of landslide susceptibility analysis with generalized additive models. In: SAGA – Seconds Out (= Hamburger Beiträe zur Physischen Geographie und Landschaftsoekologie, vol. 19). J. Boehner, T. Blaschke, L. Montanarella; 2008. p. 23–32.
- [13] Hijmans RJ, van Etten J, Cheng J, Mattiuzzi M, Sumner M, Greenberg JA, et al.. raster: Geographic Data Analysis and Modeling; 2015. R package version 2.3-40. Available from: <http://cran.r-project.org/web/packages/raster/>.
- [14] Brown RE. Peel South Geoscience Data Package. Sydney: Geological Survey of New South Wales; 2002.
- [15] Brown RE. Peel South Exploration NSW Geophysics - interpretation of new data for exploration and geological investigations in the western New England area of NSW. Quarterly Notes. 2003;(114).
- [16] Barnes EM, Sudduth KA, Hummel JW, Lesch SM, Corwin DL, Yang C, et al. Remote- and Ground-Based Sensor Techniques to Map Soil Properties. Photogrammetric Engineering & Remote Sensing. 2003;69(6):619–630.
- [17] Johnson CK, Doran JW, Duke HR, Wienhold BJ, Eskridge KM, Shanahan JF. Field-Scale Electrical Conductivity Mapping for Delineating Soil Condition. Soil Science Society of America Journal. 2001;65:1829–1837.
- [18] Terra JA, Shaw JN, Reeves DW, Raper RL, van Santen E, Schwab EB, et al. Soil Management and Landscape Variability Affects Field-Scale Cotton Productivity. Soil Science Society of America Journal. 2006;70(1):98–107.
- [19] Simbahan GC, Dobermann A, Goovaerts P, Ping J, Haddix ML. Fine-resolution mapping of soil organic carbon based on multivariate secondary data. Geoderma. 2006 Jun;132(3-4):471–489.
- [20] Vitharana UWA, Van Meirvenne M, Simpson D, Cockx L, De Baerdemaeker J. Key soil and topographic properties to delineate potential management classes for precision agriculture in the European loess area. Geoderma. 2008 Jan;143(1-2):206–215.
- [21] Siri-Prieto G, Reeves DW, Shaw JN, Mitchell CC. World’s Oldest Cotton Experiment: Relationships between Soil Chemical and Physical Properties and Apparent Electrical Conductivity. Communications in Soil Science and Plant Analysis. 2006 Jun;37(5-6):767–786.
- [22] Lal R, Follett RF, editors. Soil Carbon Sequestration and the Greenhouse Effect. 2nd ed. Madison: Soil Science Society of America; 2009.

- [23] Pinter PJ, Hatfield JL, Schepers JS, Barnes EM, Moran MS, Daughtry CST, et al. Remote Sensing for Crop Management. *Photogrammetric Engineering & Remote Sensing*. 2003;69(6):647–664.
- [24] Franzluebbers AJ. Water infiltration and soil structure related to organic matter and its stratification with depth. *Soil & Tillage Research*. 2002;66:197–205.
- [25] Rouse JWJ, Haas RH, Schell JA, Deering DW. Monitoring vegetation systems in the great plains with ERTS. In: 3rd ERTS symposium NASA SP-351. Washington DC: US Government Printing Office; 1973. p. 309–317.
- [26] Huete AR. A soil-adjusted vegetation index (SAVI). *Remote Sensing of Environment*. 1988;25:295–309.
- [27] Stamatiadis S, Christofides C, Tsadilas C, Samaras V, Schepers JS, Francis D. Ground-Sensor Soil Reflectance as Related to Soil Properties and Crop Response in a Cotton Field. *Precision Agriculture*. 2005;6:399–411.
- [28] Trotter MG, Lamb DW, Donald GE, Schneider DA. Evaluating an active optical sensor for quantifying and mapping green herbage mass and growth in a perennial grass pasture. *Crop and Pasture Science*. 2010;61(5):389–398.
- [29] Mulder VL, de Bruin S, Schaepman ME, Mayr TR. The use of remote sensing in soil and terrain mapping A review. *Geoderma*. 2011 Apr;162(1-2):1–19.
- [30] Muñoz Robles C, Frazier P, Tighe M, Reid N, Briggs SV, Wilson B. Assessing ground cover at patch and hillslope scale in semi-arid woody vegetation and pasture using fused Quickbird data. *International Journal of Applied Earth Observation and Geoinformation*. 2012 Feb;14(1):94–102.
- [31] Payero JO, Neale CMU, Wright JL. Comparison of Eleven Vegetation Indices For Estimating Plant Height Of Alfalfa And Grass. *Applied Engineering in Agriculture*. 2004;20(3):385–393.
- [32] Hibbard KA, Archer S, Schimel DS, Valentine DW. Biogeochemical Changes Accompanying Woody Plant Encroachment in a Subtropical Savanna. *Ecology*. 2011;82(7):1999–2011.
- [33] Graham S, Wilson BR, Reid N, Jones H. Scattered paddock trees, litter chemistry, and surface soil properties in pastures of the New England Tablelands, New South Wales. *Australian Journal of Soil Research*. 2004;42:905–912.
- [34] Moore ID, Gessler PE, Nielsen GA, Peterson GA. Soil Attribute Prediction Using Terrain Analysis. *Soil Science Society of America Journal*. 1993;57(2):443–452.

- [35] Lal R. Soil carbon dynamics in cropland and rangeland. *Environmental pollution* (Barking, Essex : 1987). 2002 Jan;116(3):353–62.
- [36] Terra JA, Reeves DW, Shaw JN, Raper RL. Impacts of landscape attributes on carbon sequestration during the transition from conventional to conservation management practices on a Coastal Plain field. *Journal of Soil and Water Conservation*. 2005;60(6):438–446.
- [37] Mueller TG, Pierce FJ. Soil Carbon Maps: Enhancing Spatial Estimates with Simple Terrain Attributes at Multiple Scales. *Soil Science Society of America Journal*. 2003;67(1):258–267.
- [38] Terra JA, Shaw JN, Reeves DW, Raper RL, van Santen E, Mask PL. Soil Carbon Relationships with Terrain Attributes, Electrical Conductivity, and a Soil Survey in a Coastal Plain Landscape. *Soil Science*. 2004;169(12):819–831.
- [39] Quinn P, Beven K, Chevallier P, Planchon O. The Prediction Of Hillslope Flow Paths For Distributed Hydrological Modelling Using Digital Terrain Models. *Hydrological Processes*. 1991;5(1991):59–79.
- [40] Van Oost K, Govers G, Desmet P. Evaluating the effects of changes in landscape structure on soil erosion by water and tillage. *Landscape Ecology*. 2000;15:577–589.
- [41] Van Rompaey AJJ, Verstraeten G, Van Oost K, Govers G, Poesen J. Modelling mean annual sediment yield using a distributed approach. *Earth Surface Processes and Landforms*. 2001 Oct;26(11):1221–1236.
- [42] DluogoßV, Fiener P, Schneider K. Layer-Specific Analysis and Spatial Prediction of Soil Organic Carbon Using Terrain Attributes and Erosion Modeling. *Soil Science Society of America Journal*. 2010;74(3):922–935.
- [43] Cook SE, Corner RJ, Groves PR, Grealish GJ. Use of airborne gamma radiometric data for soil mapping. *Aust J Soil Res*. 1996;34:183–194.
- [44] Rawlins BG, Marchant BP, Smyth D, Scheib C, Lark RM, Jordan C. Airborne radiometric survey data and a DTM as covariates for regional scale mapping of soil organic carbon across Northern Ireland. *European Journal of Soil Science*. 2009;60(February):44–54.
- [45] McKenzie NJ, Ryan PJ. Spatial prediction of soil properties using environmental correlation. *Geoderma*. 1999;89:67–94.
- [46] Singh K, Murphy BW, Marchant BP. Towards cost-effective estimation of soil carbon stocks at the field scale. *Soil Research*. 2012;50:672–684.
- [47] Breiman L, Friedman J, Stone CJ, Olshen RA. *Classification And Regression Trees*. Belmont, Calif: Chapman and Hall/CRC; 1984.

- [48] Hastie T, Tibshirani R, Friedman J. The Elements of Statistical Learning Data Mining, Inference, and Prediction. 2nd ed. Springer; 2009.
- [49] Ripley BD. Pattern Recognition and Neural Networks. Cambridge: Cambridge University Press; 1996.
- [50] Chipman HA, George EI, McCulloch RE. Bayesian CART Model Search. Journal of the American Statistical Association. 1998;93(443):935–948.
- [51] Tibshirani R, Knight K. Model Search by Bootstrap "Bumping". Journal of Computational and Graphical Statistics. 1999;8(4):671–686.
- [52] Breiman L. Random Forests. Machine Learning. 2001;45(1):5–32.
- [53] Kursa MB, Rudnicki WR. Feature Selection with the Boruta Package. Journal Of Statistical Software. 2010;36(11):1–13.
- [54] Elith J, Leathwick JR, Hastie T. A working guide to boosted regression trees. Journal of Animal Ecology. 2008;77(4):802–813.
- [55] Friedman JH. Stochastic gradient boosting. Computational Statistics & Data Analysis. 2002;38:367–378.
- [56] Friedman JH. Greedy Function Approximation: A Gradient Boosting Machine. The Annals of Statistics. 2001;29(5):1189–1232.
- [57] Friedman JH, Meulman JJ. Multiple additive regression trees with application in epidemiology. Statistics in Medicine. 2003;22(9):1365–1381.
- [58] Quinlan JR. Learning with continuous classes. In: Adams A, Sterling L, editors. AI' 92: proceedings of the 5th Australian Joint Conference on Artificial Intelligence. Hobart, Tasmania: Singapore: World Scientific; 1992. p. 343–348.
- [59] Chipman HA, George EI, McCulloch RE. Bayesian Treed Models. Machine Learning. 2002;48(1-3):299–320.
- [60] Belsley DA, Kuh E, Welsch RE. Regression diagnostics: Identifying influential data and sources of collinearity. New York: Wiley; 1980.
- [61] Akaike H. A New Look at the Statistical Model Identification. IEEE Transactions on Automatic Control. 1974;19(6):716–723.
- [62] Venables WN, Ripley BD. Modern Applied Statistics with S. 4th ed. New York: Springer; 2002.
- [63] Konishi S, Kitagawa G. Information Criteria and Statistical Modeling. New York: Springer; 2008.
- [64] Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian Data Analysis. 2nd ed. Chatfield C, Tanner M, Zidek J, editors. London: Chapman and Hall/CRC; 2004.

- [65] Mitchell TJ, Beauchamp JJ. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*. 1988;83:1023–1036.
- [66] Geweke J. Variable selection and model comparison in regression. *Bayesian Statistics*. 1996;5:609–620.
- [67] George EI, McCulloch RE. Approaches for Bayesian variable selection. *Statistica Sinica*. 1997;7(2):339–373.
- [68] Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society Series B (Methodological)*. 1996;58(1):267–288.
- [69] Zou H. The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*. 2006;101(476):1418–1429.
- [70] Fan J, Li R. Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties. *Journal of the American Statistical Association*. 2001;96(456):1348–1360.
- [71] Zhang CH. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*. 2010;38(2):894–942.
- [72] Ahmed SE. *Penalty, Shrinkage and Pretest Strategies Variable Selection and Estimation*. London: Springer; 2014.
- [73] Efron B, Hastie T, Johnstone I, Tibshirani R. Least Angle Regression. *The Annals of Statistics*. 2004;32(2):407–451.
- [74] Park T, Casella G. The Bayesian Lasso. *Journal of the American Statistical Association*. 2008;103(482):681–686.
- [75] Breiman L. Bagging predictors. *Machine Learning*. 1996;24(2):123–140.
- [76] Hoerl AE, Kennard RW. Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*. 1970;12(1):55–67.
- [77] Lumley T, Miller A. leaps: regression subset selection; 2009. R package version 2.9. Available from: <http://cran.r-project.org/package=leaps>.