THE UNIVERSITY of ADELAIDE

Land Suitability Assessments for Agriculture using the Data By-products of Mining Exploration

A thesis submitted in fulfilment of the requirements

for the degree of *Doctor of Philosophy*

Ingrid Ahmer

B.A., Grad Dip Comp Sci, B.Sc(Hons), M.Eng

School of Biological Sciences

The University of Adelaide

October 2022

Dedication

In memory of my nephew Tom, sister Kirstan and mother Gwendith,

Thomas Djemmy Ahmer

27 August 1999 – 23 May 2016

Kirstan Margrethe Ahmer

31 October 1956 – 19 August 2016

Gwendith Ahmer

24 August 1925 – 16 July 2019

Abstract

This thesis was proposed in the context of the social licence to operate for a new gold mine in West Africa that necessitated the relocation of subsistence farmers and asked the question:

> *Can land suitability assessments for agriculture be carried out effectively using the data by-products of mining exploration, without the necessity for and cost of additional data acquisition?*

The project objectives were to develop a land suitability assessment process that, firstly, was locally relevant with outputs usable by farmers in the vicinity of a new gold mine in the tropical south-west of Burkina Faso and, secondly, is transferable to other sites with different terrain, climate and styles of agriculture.

The compensation maps prepared as part of the application for the mining licence provided known occurrences for locally grown crops, so facilitating the use of data-driven species distribution methods to produce crop suitability maps. However, the clustered occurrences presented the likelihood of spatial sampling bias affecting models and the evaluation of results was complicated by the lack of test data outside of these areas.

The maximum entropy algorithm (Maxent) was used to produce crop suitability models using a methodology that took advantage of the geographical separation of the presence data sites to develop cross-validation models trained on different sets of presence points. The accuracy of model predictions (using the area under the curve (AUC) of the receiver operator curve for test data) and similarity of resulting suitability maps (from correlations) were compared to assess model accuracy, robustness and sensitivity to sampling bias; however, region wide validation of results was not possible with the available data. The methodology was then applied to two local sites in South Australia for which region-wide verification data were available in order to validate the methods used and to demonstrate their transferability to other sites with different terrain, climate and styles of agriculture.

Neither the categorical soil map supplied by the exploration company nor the publicly available global maps of commonly used soil properties were useful as model predictors. Algorithmic methods were devised to process both sources of soil data into a new set of hybrid soil layers (combining the fine spatial detail of the supplied map with the multi-dimensionality of the soil property maps) that proved effective in modelling. The thesis contributes to the application of species distribution modelling by presenting this new method for converting categorically valued maps into continuously valued raster layers for effective use by modelling algorithms.

The thesis also demonstrates effective cost-free and language-independent mapping solutions that overcome the local challenges of illiteracy and poor access to technology. Paper maps were designed for map users without access to other technology, and interactive maps were produced for map users with access to electronic devices, with and without internet access.

The method of predicting local agricultural land suitability presented in the thesis has been shown to be transferrable to other sites. It is particularly well suited to mining applications in developing countries where detailed data on local agriculture are collected as part of the environmental and social impact assessments. As such, it could become a model for future mining projects and contribute to more successful collaborations between the mining sector and local communities in developing countries.

Thesis Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Acknowledgements

This research journey has been an adventure. I am deeply grateful to my principal supervisor Associate Professor Bertram Ostendorf for inspiring me to undertake the spatial science research degree that has become this thesis, also to Professor Megan Lewis for taking on the associate supervisor role and to the donors of the C J Everard Scholarship and the Norman and Patricia Polglase Supplementary Scholarship for the financial support that made it possible. Thank you to Peter Williams and Simon Bolster of Gryphon Minerals Ltd for formulating the idea for the project and then offering it to the University of Adelaide with supporting data - so giving me the opportunity for such interesting research. I am sincerely grateful to the Burkinabe people who have helped me understand the context of this research. I am especially grateful for my friendship with Sylvie Nombre and for her invaluable assistance throughout this project in understanding the agriculture, customs and languages of the local Burkina Faso communities. Sincerest thanks to Dr Moumouni Konate who provided very valuable guidance at the start of the project (while still living in Adelaide in 2016) and to his family for hosting me in Bobo-Dioulasso in 2018. I especially wish to thank Salimata Konate for organising the day trip to Banfora that allowed me to observe first-hand the local agriculture in the vicinity of my project area. I sincerely thank Bertram for his ongoing guidance, encouragement and friendship and acknowledge his role in initiating the project and providing the resources to facilitate it. Thank you to the Biodiversity and Climate Change Virtual Laboratory for the excellent platform to experiment with species distribution modelling. Thank you also to Dr Neville Crossman for his interest early in the project and his suggestion to use the SoilGrids global soil maps. Special thanks to both Professor Megan Lewis and Patricia Hanlon for their very high-quality feedback on drafts of the thesis. I am also indebted to those people whose interest and encouragement has motivated me at difficult times. I particularly thank Sharon Rowett for her excellent feedback and advice when I was grappling with soil data issues and Professor Andrew Metcalfe from the School of Mathematical Sciences for his ongoing interest in the project and his excellent feedback and advice when I was grappling with how to write up my results. Sincere thanks to all the members of the Spatial Science Group (initially at Waite Campus and now on 3rd floor of the Oliphant Building) for their friendship during this project, with particular mention of Dr Iffat Ara who welcomed me so warmly when I joined the group in 2016 and Professor William Breed with whom I have been sharing an office and who has so generously shared his love of biology with me. Thank you to Steve Howard for the generous gift of his library of GIS and remote sensing books as I was starting my degree. Finally, I wish to thank my friends and family for their patience and support throughout this period of study. Sadly, my mother, sister Kirstan and nephew Tom were not able to see this project concluded.

Table of Contents

Table of Figures

List of Tables

Table of Abbreviations

ABS	Australian Bureau of Statistics
AHP	Analytical hierarchy process
ALOS	Advanced Land Observing Satellite
ASTER	Advanced Spaceborne Thermal Emission and Reflectance Radiometer
BCCVL	Biodiversity and Climate Change Virtual Laboratory
cmolc	Centimoles of charge
DEM	Digital elevation model
DFAT	Department of Foreign Affairs and Trade
ESRI	Environmental Systems Research Institute
FAO	Food and Agriculture Organisation (of the United Nations)
GIS	Geographical information system
IFC	International Finance Corporation
ISRIC	International Soil Reference Information Centre
J	Joule
MAE	Mean absolute error
MCDM	Multi-criteria decision making
MRRTF	Multi-resolution ridge top flatness
MRVBF	Multi-resolution valley bottom flatness
nGy/hr	nanoGray per hour
PC	Principal component
PCA	Principal component analysis
PIRSA	Department of Primary Industries and Regions SA
ppm	parts per million
RMSE	Root mean squared error
SA	South Australia
SAGA	System for Automated Geoscientific Analyses
SDM	Species Distribution Modelling
SLGA	Soil and Landscape Grid of Australia
SRTM	Shuttle Radar Topography Mission
SWI	SAGA Wetness Index
TWI	Topographic Wetness Index
USDA	United States Department of Agriculture
USGS	United States Geological Survey
UTM	Universal Transverse Mercator
WGS	World Geodetic System

Chapter 1 Introduction

1.1 Overview

This research project was proposed in the context of the establishment of a large new industrial gold mine in south-west Burkina Faso and the ongoing disruption to the livelihoods and way of life of the existing communities in that region. Large scale surface gold mining projects generate extreme environmental and social impacts. Badly managed, they can result in damaging long term effects for affected regions and the resident populations. However, well managed projects that are able to engage positively with host communities have the potential to stimulate regional development and result in sustainable improvements in standards of living and quality of life for resident communities.

Source: https://www.iamgold.com/English/operations/operating-mines/essakane-gold-mine-burkina-faso

Figure 1-1 Map of West Africa showing the site of the new gold mine in Burkina Faso

Burkina Faso is a very poor country and the dry Sahel region in the north is very vulnerable to impacts of climate change. The south-west, where this project is located, is very fertile, and has the potential to become the bread-basket for a country struggling to ensure national food security for its people into the future. However, farming in the south-west is still primarily rain-fed subsistence agriculture, and farmers are constrained in their attempts to improve production by lack of access to capital, credit and agricultural technology. The injection of capital by way of wages and compensation payments from mining will stimulate regional development in the short term. But lasting benefits will only be derived by ensuring a productive and sustainable agricultural sector that will last beyond the life of the mine.

In this project, mining exploration data is used to develop an objective evidence-based agricultural land suitability assessment tool to support the sustainable development of the region. The background information provided below addresses the three major aspects of this project: (1) mining context, to understand the geopolitical factors behind the project; (2) land evaluation, to

explain the particular nature of the research; and (3) research site and local communities, shaping the character of the land evaluation task. The research questions explored in this project are discussed in Section 1.7 and the thesis outline is presented in Section 1.8.

1.2 Mining in the developing world

Abundant mineral resources (and especially gold) can offer enormous potential wealth to a country. And yet, the term "resource curse" has been coined to describe mining in the developing world, where the countries *without* mineral resources typically out-perform the countries rich in mineral resources in terms of economic growth Kumah (2006).

The impacts of mining are wide ranging and include environmental damage, land conflict, and societal impacts for affected communities, with the most extreme impacts sometimes including environmental catastrophes and human rights abuses (Kumah 2006, Schueler, Kuemmerle et al. 2011). Following an investigation of twenty years of land use change due to mining in Brazil, Sonter, Moran et al. (2014) concluded that there is a need to have sustainable development goals both within the mining community and beyond to manage environmental and societal impacts.

Marco Chown Oved/Toronto Star

Hugh Brown/Pulitzer Center

Hugh Brown/Quartz Africa

Figure 1-1 Essakane open cut gold mine in Burkina Faso (left), artisanal gold mining scenes (right)

In relation to the impacts of surface gold mining in West Africa, a 2006 assessment of the impact of gold mining in Ghana (a southern neighbour of Burkina Faso) reported multiple negative impacts, as follows: (1) only 10% of the earnings from gold exports were actually retained in the country; (2) there were low labour absorption rates in mining communities (the mines were not creating jobs for Ghanaians); (3) the displacement of farmers and resident small scale miners had caused unrest; (4) there had been multiple cyanide spills; (5) land degradation had occurred as a result of large scale mining and agricultural intensification; (6) dust pollution was causing disease; (7) relocation schemes had led to loss of land and resources, chronic impoverishment and social disruption; and (8) there had been human rights abuses (Kumah 2006). The damaging environmental impacts of surface gold mining can be observed in Figure 1-1 which shows both an industrial gold mine and scenes of artisanal mining in Burkina Faso.

Mining projects are capital intensive, and so poor countries rely heavily on foreign investment for financing their mining projects. The recognition that African countries have generally drawn little benefit from foreign owned mining projects is now leading to efforts to obtain better outcomes from mining, such as negotiating fairer deals, greater transparency, better channelling of revenues back to local populations and, particularly related to this project, positive engagement by the mining companies to benefit local communities (Kimani 2009).

1.2.1 Social license for mining

Local communities have emerged as important governance actors in the mining sector and there is now widespread recognition that developers need to gain a *social license to operate* (SLO) from local communities to avoid potentially costly conflict and exposure to social risks. An SLO is considered to exist when a mining project is seen as having the ongoing approval and broad acceptance from society and host communities to conduct its activities (Prno and Slocombe 2012).

The term 'social license to operate' was introduced in 1997 by Jim Cooney when talking about political risk at a World Bank conference on the future of mining (Canada Science and Technology Museum 2017). In 1999 Ian Thompson explained that a social licence to operate had become essential for success in international exploration as a consequence of globalization. He cited the main drivers as: (1) the internationalisation of environmental and social issues; (2) the rise of international civil society to monitor the activities of companies and hold them accountable; and (3) the explosive growth in communications which has empowered communities to challenge mining companies as they now have access to support and information networks to assist them in challenges (Thompson 1999).

An SLO is not a single licence granted by a particular community but rather a continuum of multiple licences that must be achieved across various levels of society. Dare, Schirmer et al. (2014) observe that achieving an SLO is important for organisations with long time horizons and high exposure to global markets, and that community engagement is critical in achieving it. The path to securing and holding an SLO is through building trust with local communities. The results of Moffat and Zhang (2014) highlight the importance of fair treatment and high quality engagement when dealing with communities to build this trust, alongside mitigation of operational impacts. Meaningful dialogue is central to the process of engagement. Mercer-Mapstone, Rifkin et al. (2017) examined dialogue outcomes from the perspective of community engagement practitioners and identified sixteen outcomes from meaningful dialogue that included trust, relationships, social acceptance, shared decision-making and legitimacy.

in 1998, the International Finance Corporation (IFC) adopted a set of environmental and social safeguard policies to assist private sector businesses address development challenges. These policies were revised to form the IFC Performance Standards on Environmental and Social Sustainability, introduced in 2006 and further revised in 2012 (IFC 2016). They prescribe a set of eight performance standards that an IFC client is to meet throughout the life of an investment by IFC. Performance Standard 1 (Assessment and Management of Environmental and Social Risks and Impacts) deals with the social licence to operate (IFC 2012).

Although early applications of SLO focused primarily on positive engagement with local communities, Komnitsas (2020) looks to the future of mining and observes that "the SLO may prove an important tool in future mining in order to safeguard the supply of raw materials, minimize the environmental footprint and improve the quality of life in the affected regions."

1.3 Gold mining in Burkina Faso

Burkina Faso is one of the poorest nations in the world with 80% of its population engaged in agriculture, mainly supporting subsistence livelihoods (USAID 2021). In 2020, Burkina Faso had a per capita gross annual income of US$ 858 (World Bank Group 2022) and was ranked 173 out of 192 countries in terms of GDP(PPP)[1] by Global Finance Magazine (Ventura 2022). However, Burkina Faso is also very rich in gold reserves. In 2009 gold surpassed cotton as the country's leading commodity export (USAID 2017), and by 2012, Burkina Faso had become the fourth largest gold producer in Africa (Côte 2016). So the issue of obtaining benefits for its population and mitigating the damaging impacts from mining is now of critical importance in Burkina Faso.

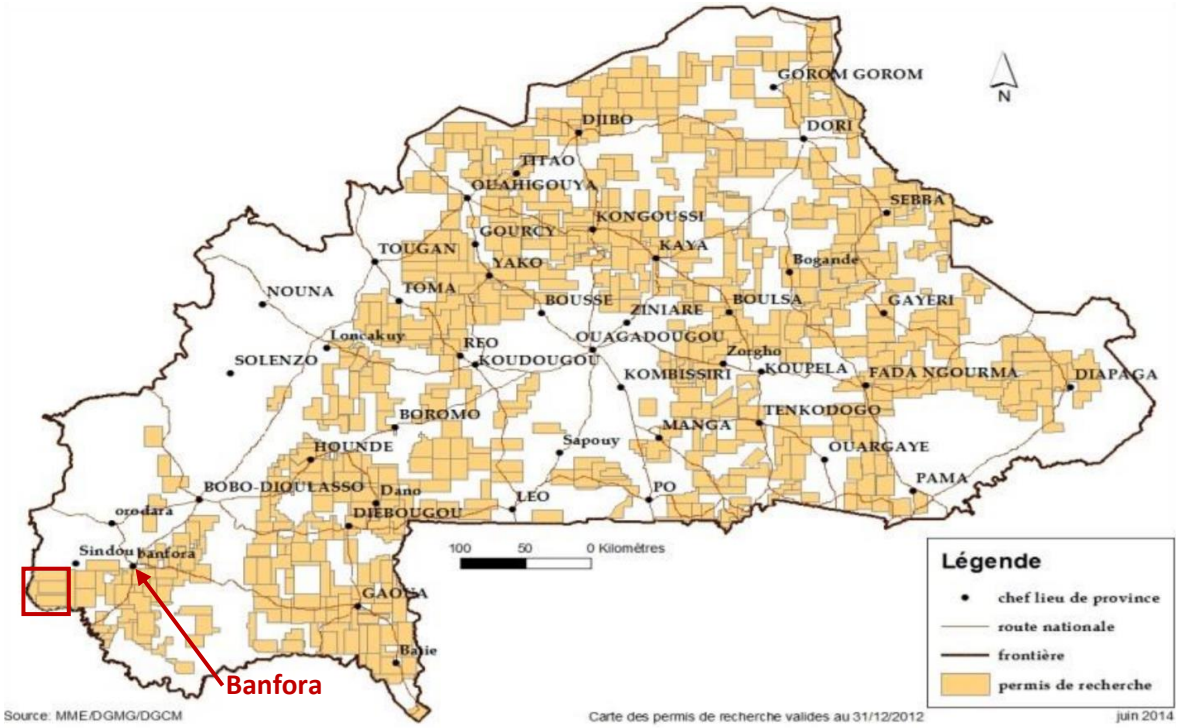

Figure 1-2 Map of 2012 exploration permits in Burkina Faso (shaded in gold) - reproduced from 2012 Burkina Faso EITI Report (IETI 2014). The Banfora Gold Project region is outlined in red.

Gold mining has a long history in Burkina Faso with gold being extracted since ancient times from the Poura mines near the Black Volta River (Rupley, Bangali et al. 2013). However, the recent increase in the importance of gold mining has only occurred since the 1980s when several severe droughts resulting in famine caused many farmers to turn to gold mining (Luning 2008) – see Figure 1-1. Artisanal gold mining is dangerous work and often involves the use of toxic chemicals (mercury and cyanide) for extracting gold (Al Jazeera English 2012, Al Jazeera English 2014). Artisanal gold mining continues in Burkina Faso as a means of livelihood for an estimated 400,000 miners, but it is the industrial gold mines set up by international mining companies that are now having the biggest impact, and this change has happened very rapidly (Gongo and Bax 2016).

[1] Gross domestic product (GDP) based on purchasing-power-parity (PPP) per capita - a measure commonly used for comparing the relative wealth of countries that takes into account the impact of exchange rates.

The liberalization of the Burkina Faso mining code in 2003 led to increased investment in prospecting by over 30 international companies (Rupley, Bangali et al. 2013). So, whilst only two exploration permits were granted in Burkina Faso in 2001, by 2012, 660 were held covering most of the country (Côte 2016) – see Figure 1-2.

For large scale mining projects, companies referred to as Juniors typically specialise in exploration, and companies referred to as Majors specialise in extraction (Luning 2008). Several Australian companies have been involved (as Juniors) in gold exploration in Burkina Faso. One of these, Gryphon Minerals Limited, formerly based in Perth, supplied the data for this research project from its Banfora Gold Project in the south-west corner of the country.[2] In June 2016, Gryphon Minerals was acquired by a Major (Teranga Gold Corporation) in order to bring the Banfora Gold Project (now renamed as the Wahgnion Gold Project) into production.

As a nation, Burkina Faso has to manage its extractive industry to avoid the damaging impacts of mining and ensure revenues are properly channelled to promote development and enhance the wellbeing of its people. Burkina Faso was one of the first countries to join the Extractive Industries Transparency Initiative, which is a global standard to promote open and accountable management of natural resources (Kimani 2009). And the mining code compels exploration companies to negotiate with communities and acknowledge local practices (Luning 2012).

The state owns all land in Burkina Faso and issues the licences for exploration and mining. Research permits are issued for plots of 250 square kilometres. In order to transform a research permit into a mining permit, a company must submit an application to the Ministry of Mines that includes a feasibility study containing geological data and an inventory of houses to be moved and agricultural fields to be compensated (Luning 2008). Once the mining permit is granted, local residents can proceed to negotiate a compensation deal. The potential benefits to a community for hosting an industrial gold mine include development of the village, paid jobs, water if there is construction of dams, and improved transport connections to broader markets. So there can be competition between communities to host the mine. Luning (2012) has observed that initial contacts with the mining companies are often marked by promises and expectations but, later, disappointments can start affecting alliances and increasing tensions in neighbourly relations.

1.3.1 Wahgnion Gold Project

The Wahgnion gold project (formerly known as the “Banfora Gold Project”) is located in the south-west corner of Burkina Faso in one of the wettest and most fertile parts of the country. The project comprises a mine licence area of 89 km^2 and a regional exploration land package of 1,093 km^2 in the Niankorodougou Commune of Léraba Provence in the Cascades Region. The mining licence was granted in February 2014 and the social impact assessment was completed in November 2014 (Intersocial Consulting 2014). Early construction works commenced in 2017 and the expected life of the mine is nine years following the first gold pour in 2019 (Mining technology 2018).

The area impacted by the Wahgnion gold project mining activity is presented in Figure 1-3. The establishment of the new mine will involve compensation for, and relocation of, four agricultural communities (Nogbele, Fourkoura, Stinger and Samavogo). For some of the displaced people from these communities there will be paid jobs in the mine, for others there will be new business

[2] The new mine is situated in Léraba province. Banfora, in Comoé province, with a population around 100,000 people is the largest town in the vicinity.

opportunities from the access to capital and new markets, but for most there will be the need to re-establish farming livelihoods in new locations.

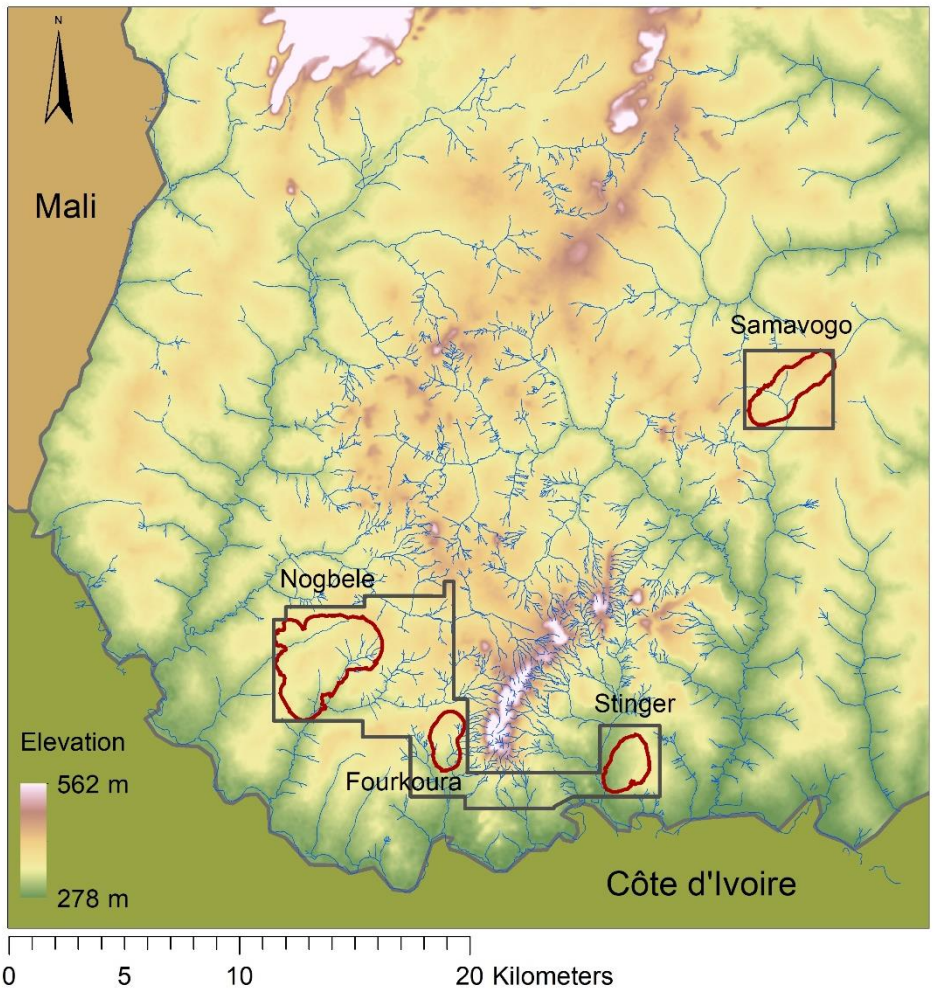

Figure 1-3 Map of the Burkina Faso project region showing: (1) elevation and drainage channels; (2) mine boundary (in black); and (3) areas where compulsory land acquisition will occur (in red)

Concerned about its social licence for mining, Gryphon Minerals offered The University of Adelaide the detailed data acquired during exploration to aid agricultural land use mapping with the goal of identifying suitable new farming land close to the planned mining complex. This is a land evaluation process that involves assessing the suitability of available fertile land for particular land uses and this process is discussed in the next section.

1.4 Land evaluation process

1.4.1 FAO framework

Decisions on land use are part of the evolution of human society (FAO 1976). The development of a standardised international framework for land evaluation commenced in the early 1970s when the Food and Agriculture Organisation (FAO) of the United Nations standardised existing country specific land assessment systems in a draft framework that was widely circulated for comment. In 1976 the FAO published "A framework for land evaluation" (FAO 1976) setting out the principles and procedures for land suitability assessment in what continues as the international standard. This was followed by the publication during 1983 to 1993 of specific FAO guidelines for rain-fed agriculture, forestry, irrigated agriculture, extensive grazing and land use planning.

Using the FAO Framework, relevant land uses are identified and the requirements for these land utilisation types determined. These requirements are then matched against measured land qualities (such as climate, soil type, landform, water availability, vegetation, etc.) to assess the degree of suitability of particular land mapping units for particular land uses. Other factors such as availability of technology, labour intensity and socio-economic factors are also relevant in the land suitability assessments. A hierarchical classification scheme with five main classes is used in the FAO framework for rating of the fitness of a given type of land for a defined land use.[3]

The FAO Framework is being revised in recognition of the fact that the scope and purpose of land evaluation has shifted from land use planning and land development projects to the sustainable management of land resources (FAO 2007). The value of local knowledge is stressed more strongly in the proposed revision, and the use of environmental models is recommended for quantifying changes and identifying future risks. Importantly, improved data availability and knowledge of its spatial variability (with the aid of Earth observation techniques) has allowed the creation of many global datasets relating to climate and land resources for use in these models.

The growing availability of computer-based information systems led to the development of automated land use planning tools and databases based on the FAO Framework (see reviews in Rossiter 1996, George 2005). Early land suitability evaluators used expert systems designed for individual problems. Later, Rossiter's microcomputer-based Automated Land Evaluation System (ALES) allowed users to build their own knowledge-based systems (Rossiter 1990). Integration with geographical information systems (GIS) (from the mid-1990s onwards) has greatly increased the usefulness and user-friendliness of such systems (for example, Elsheikh, Shariff et al. 2013).

1.4.2 GIS-based approaches

GIS-based land suitability assessments typically combine multiple GIS layers representing land qualities according to subjectively defined suitability criteria. Malczewski (2004) provides a critical overview of these methods. *Overlay mapping* is the simplest and typically uses Boolean operations and weighted linear combinations of layers – both easy to implement in a GIS environment using map algebra. *Multi-criteria decision making (MCDM)* methods manipulate

[3] The five main classes in the FAO Framework classification are: Class S1 = suitable; Class S2 = moderately suitable; Class S3 = marginally suitable; Class N1 = unsuitable for economic reasons but otherwise marginally suitable; and Class N2 = unsuitable for physical reasons. Subclasses are used to reflect kinds of limitations (e.g. moisture deficiency, erosion hazard), and further subdivision into land suitability units allows detailed interpretation at the farm planning level if this is required.

geographical data according to specified rules defined by decision maker preferences, with the *Analytical Hierarchy Process (AHP)* technique often employed in GIS-based MCDM applications to aid decision making and derivation of suitable weights. *Fuzzy logic techniques* (allowing degrees of set membership between 0 and 1) are also often applied to spatial data to develop fuzzy suitability ratings for use in MCDM, with experts typically involved to specify the membership functions.

Attua and Fisher (2010) illustrate the MCDM process with fuzzy logic in their land suitability assessment case study for pineapple production in Ghana:

1. Seven criteria maps were standardised to a continuous suitability scale using seven fuzzy membership functions, and three Boolean constraint maps were produced;
2. Weights for each criterion were derived by a panel of experts using the AHP;
3. Weighted linear combination of GIS layers was performed and the resultant image multiplied with three Boolean constraint maps to produce a final fuzzy suitability map;
4. The fuzzy suitability map was then reclassified as marginally, moderately and highly suitable based on threshold values for the pixels.

Other case studies using GIS-based MCDM methods to assess land suitability for crops in tropical regions include: an assessment of sites in Jamaica suitable for coffee growing, also using AHP (Mighty 2015); assessment of mango suitability in Malaysia, as a prototype for a decision support tool for tropical and subtropical crops (Elsheikh, Shariff et al. 2013); and the development of an agroforestry suitability map for a region in eastern India (Ahmad, Goparaju et al. 2017).

Criticisms of GIS-based approaches include: inappropriate methods of standardising suitability maps; unverified assumptions of independence among selection criteria; and oversimplification by focusing only on what can be represented in a GIS (Malczewski 2004).

Other sources of uncertainty in GIS-based land suitability assessments are described by Liu, Zhan et al. (2017) and include:

- Spatial data uncertainty deriving from the degree of accuracy of the source data;
- Raster data uncertainty necessarily occurring with the discretization of continuous features and with unavoidable loss of accuracy during the conversion of vector data to raster data;
- Index weight determination uncertainty;
- Classification uncertainty when suitability results are converted to a choropleth map.

1.4.3 Data-driven approaches

The methods described above are knowledge-based approaches to land suitability assessment that synthesise the opinions of experts and other relevant knowledge to formulate defined algorithms for deriving suitability results. Often the detailed knowledge needed by these methods is not easily available or is unreliable.

Correlative approaches, on the other hand, do not require detailed knowledge of the necessary environmental conditions for a species to thrive. Instead, they are empirical methods that are based on the assumption that the current distribution of a species is a good indicator of its ecological requirements, and so link the locations of known occurrences of a species (referred to as “presence” data) with geospatial environmental data to derive the species-environment relationships. Such methods are now commonly used in species distribution and environmental niche modelling.

The simplest correlative approach is profile modelling, which defines the environmental niche for a species using the ranges of values of the environmental variables at presence locations. More complex correlative approaches include statistical models that estimate the coefficients for the models from the input data using statistical regression, and machine learning where models are typically "trained" using one part of the input data, with the remaining data used to test the accuracy of the trained model.

The Biodiversity and Climate Change Virtual Laboratory (BCCVL)[4] online platform developed in Australia provides an excellent introduction to species distribution modelling and the commonly used algorithms (Hallgren, Beaumont et al. 2016). The BCCVL interactive interface offers two profile models, six machine learning models and four statistical regression models (as listed below). All of these models have also been implemented as R software packages if a programming environment is preferred.[5]

Profile Models

- Bioclim
- Surface Range Envelope (SRE)

Machine Learning Models

- Artificial Neural Network (ANN)
- Boosted Regression Tree (BRT)
- Classification Tree Analysis (CTA)
- Generalised Boosting Model (GBM)
- Maximum entropy modelling (Maxent)
- Random Forest (RF)

Statistical Models

- Flexible Discriminant Analysis (FDA)
- Generalised Additive Model (GAM)
- Generalised Linear Model (GLM)
- Multivariate Adaptive Regression Splines (MARS)

The statistical and machine learning models are able to fit complex relationships and can also account for interactions between variables. All can be rerun multiple times using different sets of input data in order to test robustness. Most of the algorithms (with the exception of the profile models and Maxent) use both *presence* and *absence* occurrence data, or can generate *pseudo-absence* data if true absence data is not available.

Data-driven correlative methods such as these are increasingly being used in the context of agricultural land suitability assessment. An explicit comparison of a mechanistic model (DSSAT[6]) and empirical models (GAM and Maxent) was performed by Estes, Bradley et al. (2013) in the context of crop suitability and productivity of dryland maize in South Africa. The empirical models achieved the same or better results than the mechanistic model, with the authors noting

[4] See http://bccvl.org.au/

[5] Bioclim and BRT are implemented in the 'dismo' package, GBM is implemented in the 'gbm' package, and all others are implemented in the 'biomod2' package.

[6] Decision Support System for Agrotechnology Transfer (DSSAT) is a software application program that comprises dynamic crop growth simulation models for over 40 crops. See https://dssat.net/

"Empirical modelling thus appears to be a better choice for mapping suitability/unsuitability in this study area, since it had comparable accuracy while requiring less research effort."

1.4.4 Presence-only models

Presence-only correlative models (such as Maxent) would seem very appropriate for the task of agricultural land suitability assessment as actual cropping locations are a very reliable source of presence data, and there is no need to speculate with regard to absence data. The usefulness of geographic presence-only modelling for land suitability mapping was assessed by Heumann, Walsh et al. (2011) who used Maxent to map agricultural crop suitability of lowland paddy rice and upland field crops in rural Thailand. They performed 1,000 runs of each model (using a random seed to partition crop locations into different training and test sets) in order to generate confidence intervals of model output and concluded that presence-only modelling is a very promising technique. In the Hawai`i Island Crop Probability Map project, Maxent was also used to assess the environmental conditions at the current locations of agricultural crops on Hawai`i Island in order to predict the probability of suitable conditions existing for the same crop at other locations on the island (Kemp 2012). Kemp noted that "the Maxent technique can produce good results even when given a large number of correlated variables" and that "unlike traditional statistical techniques, such correlations do not invalidate the modelling process."

1.5 The research site and local communities

The project area corresponds to the regional exploration land package for the Wahgnion gold project and covers almost 1,100 square kilometres of tropical woodland savannah in the south-west corner of Burkina Faso at altitudes of 300-400 metres (refer Figure 1-3 above). It is located in Léraba province of the Cascades region and is one of the wettest and most fertile parts of the country. The Léraba River defines the boundary with Mali and Côte d'Ivoire in this area. The river is shaded by riparian forests along most of its length and provides an important habitat for wildlife.

Inland valleys are common landscapes in the upper reaches of African river systems and many occur in this region. They are defined as seasonally flooded wetlands comprising valley bottoms and hydromorphic fringes, and generally have high agricultural production potential (Rodenburg, Zwart et al. 2014). Inland valleys represent 36% of the total area covered by wetlands in sub-Saharan Africa and perform important ecological functions such as water purification, carbon sequestration, protection against flooding and erosion, and providing habitat for many wildlife species (Kiepe 2006). From this perspective they may be viewed as fragile ecosystems. However, from an agricultural perspective they are assumed to form the basis of robust production systems. Their economic functions include crop and vegetable production, fishing, as a source of materials for thatching, fencing and basket weaving, and as a water source and grazing area for livestock during the dry season. They are less sensitive than adjacent uplands to degradation due to the frequent inflows of water bringing nutrients and organic debris. (Kiepe 2006).

Rodenburg, Zwart et al. (2014) addressed the conflicting agricultural and ecological perspectives on wetland management in their examination of African inland valleys and have proposed a methodology for fulfilling the agricultural potential of inland valleys to benefit rural livelihoods while safeguarding other ecosystem services. In characterising typical agricultural use of these valleys they observe that the valley bottoms are usually planted with rice as it is the only major food crop that can be grown in temporarily flooded conditions; dry land rice and cash crops like cotton are grown on the hydromorphic slopes; high value fruit trees (e.g. mango and cashew) and

fodder crops are grown on the upper slopes; and maize and sorghum are grown on the crests. In peri-urban areas inland valleys are used mainly for vegetable production (Rodenburg, Zwart et al. 2014). Examples of inland valley agriculture within the project region are shown in Figure 1-4. The valley bottoms (*bas-fonds*) are used almost exclusively for growing rice, while the hydromorphic slopes are suitable for other field crops.

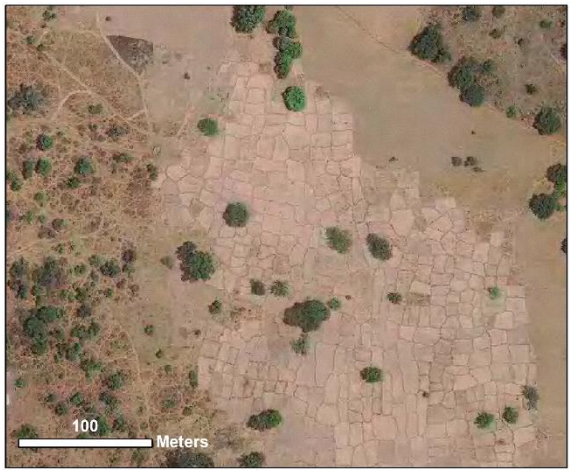

Worldview2 image April 2010

Worldview2 image April 2010

Figure 1-4 Dry season satellite images of inland valleys showing rice fields (left) and irrigated vegetable plots along a permanent watercourse within close proximity to a village (right)

Within the project region agriculture is mostly subsistence, with food crops including maize, sesame, peanut, millet, sorghum and beans, and cotton is grown as a cash crop. Rice is grown in the water courses and small market gardens are also cultivated. There is little mechanical agriculture and most fields are ploughed by oxen (see Figure 1-5). Additional income is derived from mangos, cashews and kerite nuts (used to produce shea butter and widely used in cosmetics).


Source: Gryphon Minerals Ltd

Figure 1-5 Oxen used for ploughing fields in the project region, June 2011

Agroforestry, whereby trees are integrated into agricultural landscapes to bolster food security and environmental resilience, is widely practised in much of sub-Saharan Africa (Garrity, Akinnifesi et al. 2010). Agroforestry is very important in the project region with trees typically grown around crops and on pastureland, as well as being native vegetation. The trees provide important sources of income and of wild food, and perform environmental functions such as

regulating climate, mitigating erosion and providing habitat for wildlife. Figure 1-6 shows examples of agroforestry practices near Banfora.

Source: I. Ahmer

Source: I. Ahmer

Figure 1-6 Cropping with agroforestry near Banfora, July 2018

The local climate has distinct wet and dry seasons with rainfall of around 1100mm per year falling primarily during May to October and peaking in August. Drainage is seasonal but there are some permanent water courses. Agriculture is primarily rain-fed with preparation and planting occurring at the beginning of the wet season although dry-season farming has been increasing since the introduction of diesel powered water pumps to the region in 2004. Dry season crops had been mainly rice and vegetables for market; however, a national food security campaign (Operation Bondofa) was launched in the region in 2011 to promote dry season maize (Dowd-Uribe, Roncoli et al. 2012). Export agriculture has enabled greater access by farmers to technology and inputs (fertilisers, insecticides, improved varieties, etc.) compared to the rest of the country, nevertheless local farmers are still constrained in their attempts to improve agricultural productivity by lack of access to credit, capital and agricultural technologies (Ingram, Roncoli et al. 2002).

Source: Gryphon Minerals Ltd

Figure 1-7 Village community in the project region, June 2011

Diversification of production and livelihood systems is the main strategy for risk mitigation used by Burkinabe households (Ingram, Roncoli et al. 2002, Tincani 2012). Within households, diversification is achieved by engaging in non-farm income generating activities (e.g. crafts, trade, brewing) and by supplementing the diet with wild foods (Tincani 2012). Agricultural diversification involves planting combinations of crops and crop varieties that have different growing times and water requirements in a constellation of fields in different locations so as to minimise the impact

of adverse climatic events. These plantings can be adjusted to respond to different climate predictions. For example, if higher than average rain is predicted then farmers can plant more upland fields, plant more cash crops, and plant rice instead of maize or cotton in lower lying fields. Conversely, if lower than average rainfall is predicted they may plant maize and sorghum in low lying fields (Ingram, Roncoli et al. 2002).

Source: Gryphon Minerals Ltd

Source: Gryphon Minerals Ltd

Figure 1-8 Farm scenes from the project region, June 2011: granaries (left) and farmstead (right)

Land is not privately owned in Burkina Faso. In the south-west customary institutions play the main role in allocating land parcels, either as a donation (*don*) to authochthons[7] or as a loan (*prêt*) with restricted usage rights to migrants (Engels 2014). Most people in the project region live in village communities, maintaining a permanent residence in the village and only travelling to their farms when there is work to do there (see Figure 1-7 and Figure 1-8).

Education levels amongst the adult population of the project region are very low. The Banfora Gold Project Social Impact Assessment reports that 55% of males and 78% of females aged 20+ have had no formal education (Intersocial Consulting 2014). Although French is the official language for Burkina Faso many adults are not literate in French, but may be literate in local languages. Adult literacy programs have been run across the country in a number of local languages with the three most commonly used languages being Dioula (spoken in the south-west), Moore and Fulfuldé (Konate 2016). The social impact assessment reported that only 3% of the local population could read French. However, 8% of males and 7% of females in the survey population indicated that they had recently attended literacy schools, primarily in the Dioula and Senoufo languages (Intersocial Consulting 2014).

Major changes to the region will occur over the lifetime of the mine and many of these will affect agricultural production. For example, the mine itself will have an impact on local water resources: the water pumped out of the mine will cause drawdown of the water table in the vicinity of the mine but will be far in excess of the needs of the mine. The excess water can be stored in dams and so be available for use in irrigation to improve local agricultural production, and can be filtered to provide clean water for drinking and domestic use to improve health outcomes (Williams 2016). The mine will bring other potential opportunities for agricultural development, including the increased access to capital to stimulate the use of new farming technologies, and access to broader markets which may lead to the cultivation of new crops.

[7] Original or indigenous inhabitants of a place.

The land suitability assessment tool proposed in this research project is being developed to aid sound land use planning by local farmers and land managers in this changing environment. The food sovereignty movement (launched in 2007 in Mali) is based on environmental, social and economic sustainability and aims to empower local farmers in the production, distribution and consumption of food (Via Campesina 2007). In keeping with these principles the assessment tool will offer information relevant and useful to its users that is independent of the technologies and resources that may be available to them or that they may choose to use. In this way it is hoped that the tool will be able to offer ongoing value as the region develops and so support decisions affecting sustainable agriculture.

1.6 Aims and objectives

This project aims to demonstrate that agricultural land suitability assessment can be done effectively using the data by-products of mining exploration. The objective of this project is to develop a land suitability assessment process that, firstly, is locally relevant with outputs usable by farmers in the vicinity of a new gold mine in the tropical south-west of Burkina Faso (West Africa) and, secondly, is transferable to other sites with different terrain, climate and styles of agriculture.

1.7 Research questions

The primary purpose of this research project was to answer the question:

1. *Can land suitability assessments for agriculture be done effectively using the data by-products of mining exploration, without the necessity for and cost of additional data acquisition?*

In undertaking this research other secondary research questions were answered, including:

2. *How can existing cropping patterns inform us with regard to potential expansion areas for particular crops?*
3. *Are the soil categories used for mining exploration useful as soil categories for agricultural land suitability models?*
4. *How can land suitability assessment results be delivered to subsistence farmers in remote locations?*

1.8 Thesis outline

This chapter has outlined the background to the research project, noting that the high-resolution geospatial and Earth observation data collected in modern mining exploration potentially has immense value for regional planning by communities and governments. Chapter 2 explores the data supplied by Gryphon Minerals Ltd and its potential for use in agricultural land suitability modelling for the project region.

Chapter 3 explains maximum entropy modelling and its use in identifying the fundamental environmental niches for species. Chapter 4 then applies this modelling technique for the purpose of agricultural land suitability modelling for the West African research site. The final crop suitability maps result from crop specific models produced by the Maxent algorithm. The robustness of these models is tested by comparing the results of crop models trained on different

sets of presence data; however, region wide validation of results is not possible with the available data.

The purpose of Chapter 5 is twofold: firstly, to validate the methodology used in Chapter 4 by applying it to sites for which region-wide validation data are available; and secondly, to demonstrate the transferability of the method to other sites with different terrain, climate and styles of agriculture. The methods used in Chapter 4 are duplicated for two local sites in South Australia for which region wide validation data are available and that can be visited to assess results and evaluate prototype presentation methods.

Chapter 6 addresses the special challenges that illiteracy and poor access to internet resources bring to the task of disseminating informative maps. The crop suitability maps produced in Chapter 4 are presented in a printable format suitable for illiterate users without access to electronic devices. An electronic version suitable for display on a computer, tablet or mobile phone is also developed, but the author is mindful that commonly used base-maps served from the internet are unlikely to be available to most target users.

Finally, Chapter 7 concludes that the method of predicting local agricultural land suitability presented in the thesis is transferrable to other sites with different physical characteristics and styles of agriculture, and is particularly well suited to mining applications in developing countries such as Burkina Faso and Ghana where detailed data on local agriculture is collected as part of the environmental and social impact assessments. As such, it could become a model for future mining projects and contribute to more successful collaborations between the mining sector and local communities in developing countries.

Chapter 2 Data for land suitability assessment

2.1 Introduction

The FAO Framework for land evaluation (FAO 1976) defines land evaluation as:

> "… the process of assessment of land performance when used for specified purposes, involving the execution and interpretation of surveys and studies of landforms, soils, vegetation, climate and other aspects of land in order to identify and make a comparison of promising kinds of land use in terms applicable to the objectives of the evaluation."

The objectives of this evaluation are to assist farmers from four farming communities subject to compulsory land acquisition in south-west Burkina Faso to identify suitable alternative locations for cultivation of their crops. The reason for the compulsory acquisition of this land and resettlement of its farmers is the establishment of a new surface gold mine at these sites.

Most of the data for this land evaluation were supplied by the exploration company, Gryphon Minerals Ltd, who successfully applied for the mining licence for the new mine. These data include both the geospatial and earth observation data for the region collected for the purpose of exploration, as well as the detailed community land use maps of the areas subject to compensation compiled for the mining licence application (refer Appendix A). The mining company, Teranga Gold Corporation, supplied the social impact assessment report containing socio-economic data for the region and, later, soil sampling results. Additional publically available data for the region has also been sourced.

This chapter explores the data available for the evaluation and considers its potential for use in agricultural land suitability modelling across the project region. The community land use files provide detailed information on the existing local agriculture in the four affected communities and are examined first. Following this, data on the climate, terrain and soils of the region are investigated with the goal of producing useful maps of land qualities relevant to the evaluation.

2.2 Local agriculture

Section 1.5 described the agriculture of the general area where the mine leases are located. It is mostly subsistence with food crops predominating. Cotton is grown as a cash crop and additional income is derived from mangos, cashews and kerite nuts. The agriculture is primarily rain-fed and there is little use of technology.

2.2.1 Current agricultural land use

The community land use maps produced for Gryphon Minerals as part of the 2014 social impact assessment provide a very detailed picture of agricultural activity being undertaken in the four affected communities at that time. Figure 2-1 shows the community land use map for the largest compulsory acquisition site (Nogbele). Maize and cotton are the dominant crops at this site, and many small rice plots are cultivated along the river courses.

Figure 2-2 shows the compensation maps for the three smaller sites (Fourkoura, Samavogo and Stinger) for comparison. It can be observed that a larger proportion of the land is used for

growing cotton at Nogbele and Stinger compared to the other two sites, and that cashews are not being cultivated at Fourkoura.

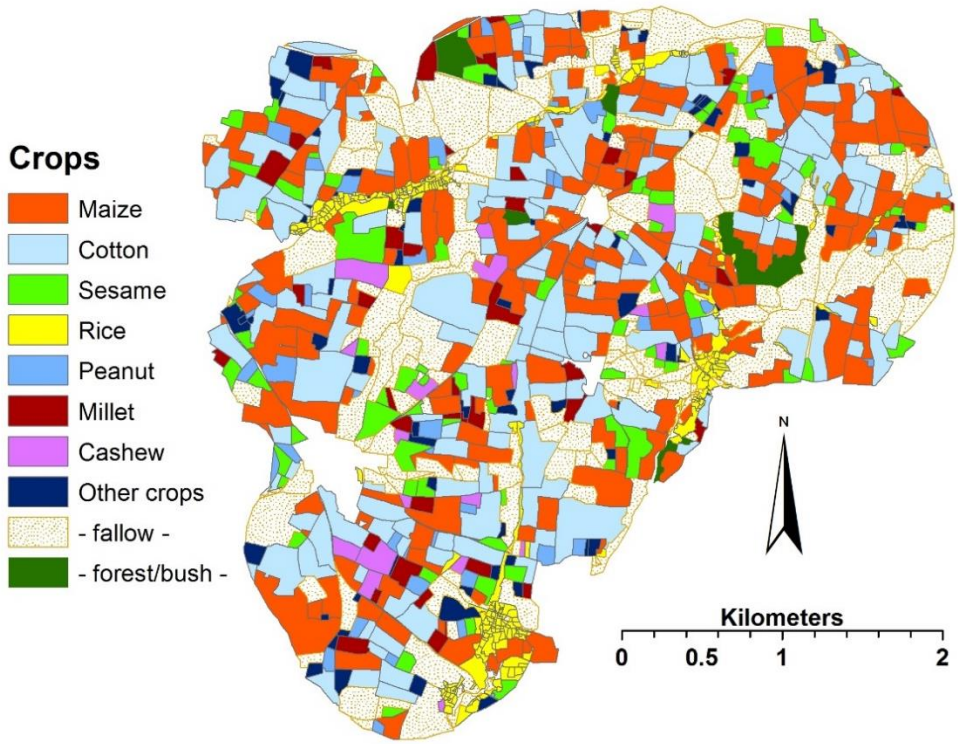

Figure 2-1 Crop compensation map for largest compulsory acquisition site (Nogbele)

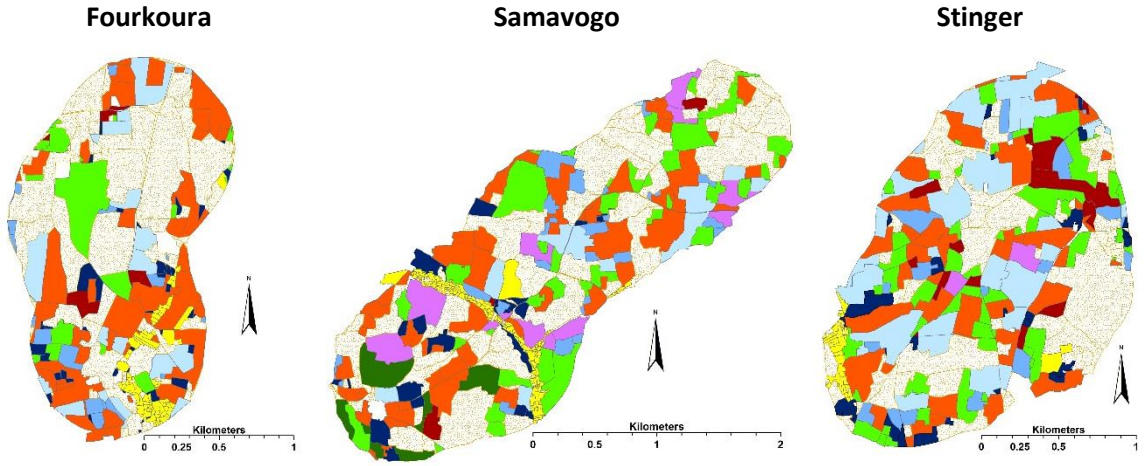

Figure 2-2 Crop compensation maps for the three smaller compulsory acquisition sites (using the same legend as Figure 2-1)

Analysis of the compensation maps yields detailed information regarding the land use and crops cultivated in these areas. The four mining lease communities contained 2,230 agricultural plots in an area of 2,557 hectares (refer Table 2-1 and Appendix B). Nogbele is the largest community, accounting for more than half of the total area, and has the largest number of plots. The average plot size across the four sites is 1.15 hectares. Samavogo has the largest average plot size at 1.26 hectares and Fourkoura has the smallest at 0.93 hectares.

Table 2-1 Comparison of mining lease communities: cultivable area and number of plots

	FOURKOURA	NOGBELE	SAMAVOGO	STINGER	Total
Area (hectares)	286.0	1,382.9	547.6	340.8	2,557.3
% total area	*11.2%*	*54.1%*	*21.4%*	*13.3%*	*100.0%*
Plots (number)	308	1,186	434	302	2,230
% total plots	*13.8%*	*53.2%*	*19.5%*	*13.5%*	*100.0%*
Average plot size (hectares)	0.93	1.17	1.26	1.13	1.15

Other aspects of land use that can be determined from these community land use maps include:

- **Land ownership**: The majority (86% by area) of the land is cultivated by the customary land owners ("Owner and user"). The proportion of "Users only" cultivation is less than 10% overall – it is almost negligible in Nogbele (at 2%), but accounts for almost a third of cultivation at Samavogo.
- **Irrigation**: The mining lease agriculture is almost entirely rain-fed with only 0.3% of the combined area having amenities for irrigation. Stinger has the highest proportion of irrigated land although still less than one percent. Almost all irrigation was used for rice.
- **Land management**: Land that has been fallow for over two years makes up 21.4% of the mining lease area, varying from 30.9% at Fourkoura down to 18.6% at Nogbele, with a further 1.7% fallow for less than two years. Only 0.2% of the land is described "uncultivable", although 4.7% is recorded as "never cultivated". A local erosion mitigation method referred to as "stony cords" occurs on plots making up 3.2% of the Fourkoura site.

2.2.2 Crops grown

Table 2-2 Main crops grown in the mining lease communities: area cultivated and number of plots[*]

Crops grown	Cultivated area (hectares)	%	Number of plots	%	Average plot size (hectares)
Beans	27.3	*1.1%*	64	*2.9%*	0.43
Cashew	65.5	*2.6%*	32	*1.4%*	2.05
Cotton	510.9	*20.0%*	224	*10.0%*	2.28
Cowpea	36.5	*1.4%*	41	*1.8%*	0.89
Earth pea	34.1	*1.3%*	39	*1.7%*	0.87
Maize	615.6	*24.1%*	346	*15.5%*	1.78
Millet	86.0	*3.4%*	71	*3.2%*	1.21
Peanut	122.8	*4.8%*	132	*5.9%*	0.93
Potato/Yam	7.9	*0.3%*	40	*1.8%*	0.20
Rice	103.2	*4.0%*	674	*30.2%*	0.15
Sesame	269.1	*10.5%*	207	*9.3%*	1.30
Sorghum	60.7	*2.4%*	42	*1.9%*	1.44
Fallow	*683.5*	*24.1%*	*314*	*14.1%*	*2.18*
Mine leases	**2557.3**	***100%***	**2230**	***100%***	***1.15***

[*] Note: Cultivated hectares and numbers of plots per crop do not sum to the totals for the mine leases as not all crops are included and some plots support more than one crop.

The twelve main crops grown in the mining lease communities are listed in Table 2-2, and a breakdown of these crops by community is provided in Appendix B. Maize and cotton are the most important, accounting for 24% and 20% of the area respectively. The plot sizes vary significantly with the type of crop grown: an average size rice plot (0.15 hectares) is less than a tenth the size of an average maize plot (1.78 hectares); and potatoes, yams and beans also have smaller plots.

Other local crops identified from the community land use maps, but accounting for very few plots, include banana, sorrel, okra, aubergine, cassava and sterculia (used in traditional medicine). The main crop cultivated was listed as *orchard* on five plots and as *mixed* on eleven plots without additional information to identify the particular crops grown.

The soil type of each plot was characterised as one of four types: clayey, sandy, gravelly or bas-fonds. The distribution of the crop planting locations across these soil types is shown in Table 2-3. Bas-fonds are the valley bottoms of seasonally flooded inland valleys and are used almost exclusively for growing rice. Gravelly soils dominate for most other crops (typically making up more than 80% of the area cultivated), with the exception of potato/yam, which are often cultivated in clayey soils and sometimes in bas-fonds.

Table 2-3 Soil type for the main crops grown in the mining lease communities (% area cultivated)

Crops grown	Bas-fonds	Gravelly	Sandy	Clayey		Total
Beans	*1.8%*	*80.4%*	*14.8%*	*3.0%*		*100%*
Cashew	*0.4%*	*86.0%*	*13.6%*			*100%*
Cotton		*94.6%*	*3.5%*	*1.9%*		*100%*
Cowpea	*0.4%*	*78.7%*	*18.9%*	*1.5%*	*0.4%*	*100%*
Earth pea		*97.6%*	*0.8%*	*1.6%*		*100%*
Maize	*1.8%*	*82.1%*	*10.9%*	*4.9%*	*0.3%*	*100%*
Millet		*91.0%*	*8.9%*		*0.1%*	*100%*
Peanut	*1.1%*	*88.3%*	*10.0%*	*0.1%*	*0.5%*	*100%*
Potato/Yam	*23.9%*	*27.4%*	*15.2%*	*33.5%*		*100%*
Rice	*75.7%*	*14.0%*	*0.3%*	*8.2%*	*1.8%*	*100%*
Sesame		*85.7%*	*8.5%*	*5.8%*	*0.0%*	*100%*
Sorghum	*4.0%*	*94.4%*	*1.6%*			*100%*
Fallow	*3.0%*	*87.4%*	*5.0%*	*4.5%*		*100%*
Mine leases	***4.7%***	***77.1%***	***10.4%***	***4.1%***	***3.8%***	***100%***

2.2.3 Implications for land suitability modelling

Implicit in the current agricultural land use pattern is the (heuristically acquired) local knowledge of the land qualities necessary for the successful agriculture of these crops in this locale. That the farmers have chosen to cultivate these crops implies their suitability to the local climate and farming methods and that markets for this produce exist. The planting locations embody existing local knowledge about land suitability and so should allow individual crops to be matched to the land qualities best supporting their cultivation.

2.3 Climate

Burkina Faso crosses three major climatic zones: (1) across the north is a hot desert climate referred to as *Sahelian*; (2) across the centre is a hot semi-arid climate referred to as *Sudano-Sahelian*; and (3) in the south is a tropical savanna climate referred to as *Sudanian*. The total annual rainfall in the Sahelian zone is less than 600mm, it is between 600 and 900mm in the Sudano-Sahelian zone, and greater than 900mm in the Sudanian zone (De Longueville, Hountondji et al. 2016).

The project region is situated in the Sudanian zone, which is alternatively classified as having a *tropical savanna climate (Aw)* according to the Köppen-Geiger climate classification system or a *wet and dry tropical climate (V3)* according to the Troll-Paffen system (Müller 1982). Historical climate records are not available at fine spatial resolution in Burkina Faso as there are only seven synoptic stations and thirteen rain gauges across the country. The closest synoptic weather station to the project region is at Bobo Dioulasso, approximately 80km north-east, and the closest rain gauge is at Banfora to the east (De Longueville, Hountondji et al. 2016).

The average monthly rainfall and temperatures at Bobo Dioulasso (sourced from Climate-Data.org[8]) are plotted in Figure 2-3 and listed in Table 2-4.

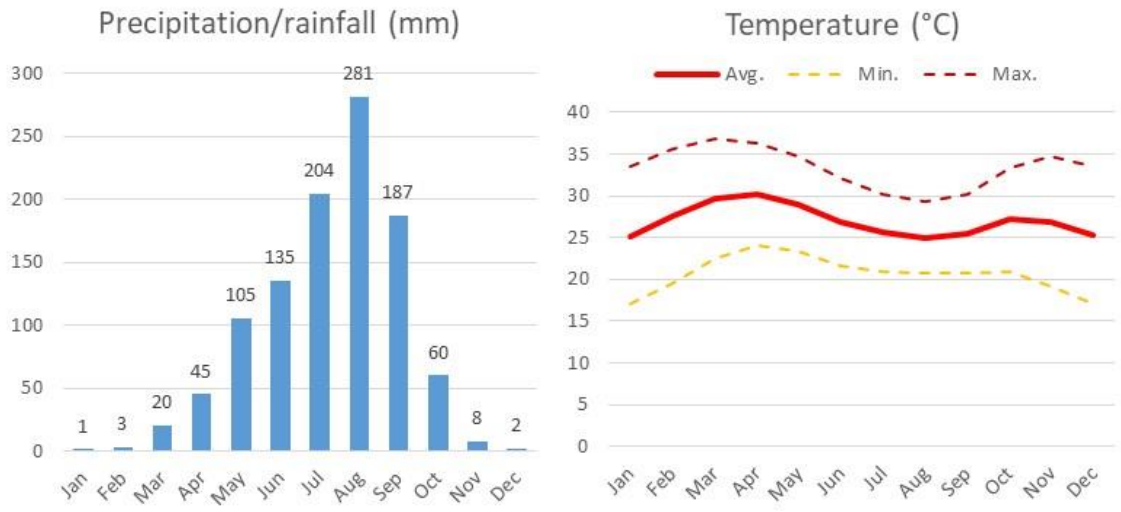

Figure 2-3 Average monthly rainfall and temperatures for Bobo Dioulasso

Table 2-4 Average monthly climate data for Bobo Dioulasso (Source: https://en.climate-data.org/)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Avg. Temperature (°C)	25.2	27.6	29.7	**30.2**	29	26.9	25.6	25	25.4	27.2	26.9	25.3
Min. Temperature (°C)	17	19.5	22.5	**24**	23.3	21.7	21	20.7	20.7	21	19.1	17.1
Max. Temperature (°C)	33.5	35.7	**36.9**	36.4	34.7	32.1	30.2	29.4	30.2	33.4	34.8	33.5
Precipitation (mm)	1	3	20	45	105	135	204	**281**	187	60	8	2

The climate has distinct wet and dry seasons with rain falling primarily during May to October and peaking in August. Temperatures are lower during the wet season and in mid-winter (December

[8] Climate-Data.org publishes global climate data at 30 arc second resolution from a climate model based on twenty years of observation data from thousands of weather stations around the world.

and January), and highest in March and April at the end of the dry season. The wet season corresponds to the growing season for field crops and harvest is usually completed by the end of December. The extensive cloud cover during the wet season hampers the use of satellite imagery to view this region during the growing season.

The average annual rainfall is estimated at 1,051mm at Bobo Dioulasso and 1,086mm at Banfora. Rainfall increases further towards the south, with estimates for two locations within the project area being 1,122mm at Baguera and 1,139mm at Fourkoura.

Figure 2-4 shows a time series of annual recorded rainfall at Bobo Dioulasso climate station from 1981 to 2017. Complete records exist for only eleven years out of the 37. Nevertheless, the variability of the rainfall is apparent with 1,346mm rainfall recorded during 1982, compared to just 644mm during 2011, and only four of the eleven years exceeding the average predicted rainfall.

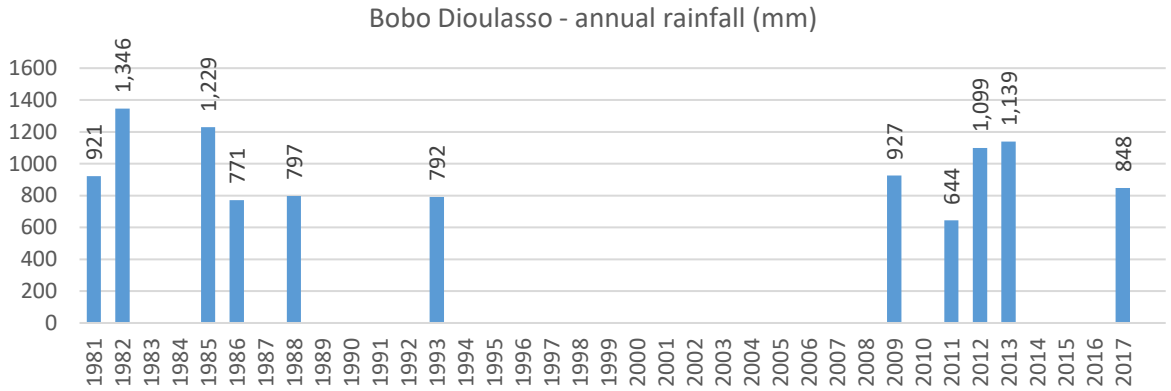

Figure 2-4 Annual rainfall recorded at Bobo Dioulasso climate station 1981 to 2017 for years where complete records exist (Source: https://en.tutiempo.net/climate/ws-655100.html)

2.3.1 Implications for land suitability modelling

De Longueville, Hountondji et al. (2016) observe that West Africa is characterized by a very sparse network of weather stations and the spatial distribution of the existing rain gauges throughout Africa is not sufficient to allow relevant description of local, regional or sub-regional climatic phenomena. Interpolated rainfall maps are typically used as inputs for agricultural land suitability models, reflecting the critical importance of available water for agriculture. However, spatially variable rainfall data are not available for this small region so cannot be included in the models developed here. Instead, rainfall will be assumed to be constant across the project area (and this is not an unreasonable assumption given the small area), with the patterns of dispersal of the rainfall runoff across the terrain modelled instead to produce an alternative and more useful measure of available water for rain-fed agriculture.

2.4 Terrain

Geomorphometry is the science of quantitative land surface analysis (Pike, Evans et al. 2009). It involves the extraction of spatial measures and features from digital topography to produce digital models of terrain.[9] Raster maps showing topographic elevation, referred to as digital elevation models (DEMs), provide the basis for digital terrain modelling in a GIS environment. From these maps many land-surface parameters may be derived using GIS analysis, which is typically implemented using neighbourhood operations on DEM cells (Jordan 2007). Note, however, that because these parameters can be generated by different algorithms or sampling strategies, and vary with spatial scale, no DEM-derived map is definitive (Pike, Evans et al. 2009).

The topography of a landscape strongly influences its land use and suitability for agriculture. Measurable terrain characteristics include elevation, slope and orientation. Many other relevant land qualities may also be derived from the terrain, such as the surface water flow and distribution, and the amount of solar radiation that may be received at a location. Several DEMs showing the surface structure of the landscape at different scales were available for further terrain analysis in the project region. These DEMs and the derivation of terrain parameters relevant for agricultural land suitability assessment are described in the following sections.

Most GIS software platforms provide functions for terrain analysis from DEMs. The System for Automated Geoscientific Analyses (SAGA) is a comprehensive and globally established open source GIS platform for scientific analysis and modelling. It was originally developed as a specialized tool for digital terrain analysis and was first released in 2004 (Conrad, Bechtel et al. 2015). SAGA can be run interactively through its graphical user interface (GUI) or using a command line interface. The R software package RSAGA provides access to geo-computing and terrain analysis functions of the SAGA GIS from within R by running the command line version of SAGA (Brenning, Bangs et al. 2018). All of the terrain layers developed for this project were created using SAGA.

2.4.1 Digital elevation model

Gryphon Minerals DEMs

The data supplied by Gryphon Minerals included several files representing digital elevation models for parts of the project area (refer Figure 2-5):

1. Single raster image taken by the **Shuttle Radar Topography Mission (SRTM)** in 2000 and covering 100% of the project region. (Cell size of 30m by 30m, signed integer values for altitude.)
2. Single raster image taken by the **Advanced Land Observing Satellite (ALOS)** in 2011 and covering approximately 85% of the project region, but not including the southern-most area and three of the mining leases. (Cell size of 5m by 5m, floating point values for altitude.)

[9] Digital terrain models have been used in geoscience applications since the 1950's and are now a major constituent of geographical information processing. Chapter 19 of Maguire, Goodchild et al. (1991) provides a review of fundamental digital geomorphometry techniques.

3. Raster image of the combination of 8 tile images taken by the **WorldView-2** satellite in 2010 covering approximately 40% of the project area in the south and including the regions missed by the ALOS image. (Cell size of 1m by 1m, floating point values for altitude.)
4. Raster image taken by the **WorldView-2** satellite in 2012 covering approximately 10% of the project area towards the north-east, including the Samavogo mining lease and adjoining the larger WorldView-2 image. (Cell size of 1m by 1m, floating point values for altitude.)

All the DEMs used the WGS 1984 UTM Zone 30N projection and had been resampled using the nearest neighbour method.

The SRTM DEM supplied by Gryphon Minerals covered the entire project region and so appeared potentially suitable as a basis for further terrain modelling. However, it had retained integer values for altitude meaning that all altitude differences between cells could only be measured as integer multiples of a metre. This coarse discretization of a naturally continuous variable suggested that this DEM (left as is) would be undesirable as a basis for further terrain modelling. Alternatively, a finer resolution DEM with 5 metre cell size and floating-point altitude could be created by combining the ALOS DEM with the southern WorldView-2 DEM, however the compatibility of the two DEMs would first need to be assessed.

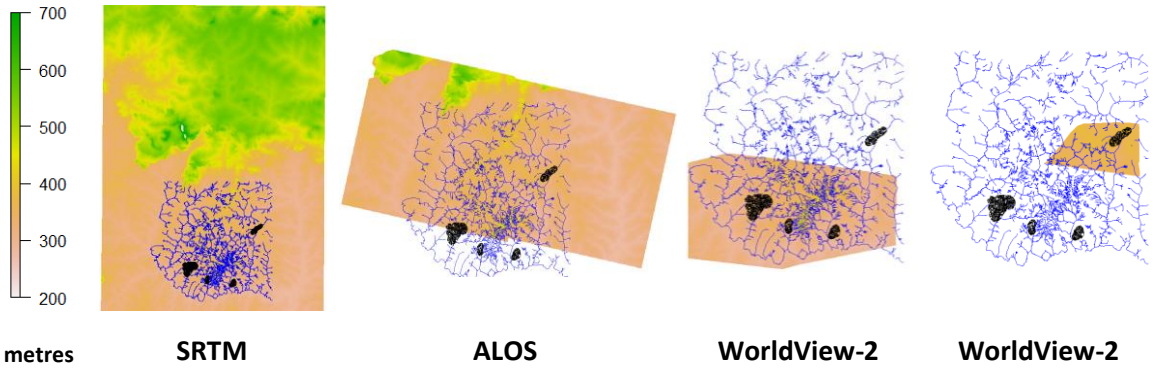

Figure 2-5 Digital elevation models supplied by Gryphon Minerals, overlaid with drainage channels (in blue) and mining leases (in black)

SRTM DEM

The high-resolution topographic data resulting from NASA's Shuttle Radar Topography Mission in 2000 has been publicly available since 2014 and can now be downloaded as 1° by 1° grids with spatial resolution of 1 arc-second (approximately 30 metres) from the USGS Earth Explorer website.[10] The SRTM 1 Arc-Second Global tile containing the project region (SRTM1N10W006V3) was thus acquired to allow comparison with the supplied DEMs. The coordinate reference system of the downloaded SRTM tile was GCS_WGS_1984, with altitude recorded using signed integer values.

The new SRTM tile was reprojected to the WGS 1984 UTM Zone 30N coordinate system and resampled using the bilinear method to 30m cell size with floating point values for altitude. Each of the supplied DEMs was then subtracted from the new SRTM DEM to observe any differences. Table 2-5 summarises the differences between the DEMs.

[10] https://earthexplorer.usgs.gov/

Table 2-5 Summary of differences between the new SRTM DEM and the four DEMs supplied by Gryphon Minerals Ltd

DEM	Minimum (m)	Maximum (m)	Mean (m)	Standard deviation (m)	Root mean squared error (m)	Mean absolute error (m)
Gryphon SRTM	-30.25	55.49	0.00	1.37	1.37	0.97
ALOS	-116.11	81.62	3.18	4.92	5.85	4.63
WorldView-2 northern	-37.22	-11.07	-24.27	1.96	24.35	24.27
WorldView-2 southern	-31.76	26.43	4.93	2.30	5.45	4.99

The most extreme differences in altitude between the DEMs occurred in the north-west corner of the project region where an extensive and dramatic series of rock formations called the Sindou Peaks are situated (see Figure 2-6). The large differences in altitude in this location and other areas with rugged terrain are readily understandable as data by-products of the different grid sizes of the ALOS and WorldView2 DEMs and the different interpolation method used for the Gryphon SRTM DEM.

The various mean differences between the supplied DEMs and the new SRTM DEM demonstrated that the ALOS and WorldView-2 DEMS were inconsistent with respect to true metres above sea level. Given such variations in base altitude detected between the supplied DEMs, the new SRTM DEM was selected as the most reliable basis for further terrain modelling.

Source: www.traveladventures.org/continents/africa/sindou-peaks11.html

Source: www.traveladventures.org/continents/africa/sindou-peaks08.html

Figure 2-6 Images of Sindou Peaks rock formations in the north-west of the project region

Terrain modelling DEM

The final DEM that was used in all subsequent terrain modelling was derived from the downloaded SRTM tile and post-processed by: (1) projection to the WGS 1984 UTM Zone 30N coordinate system with resampling using the bilinear method to 30m cell size and floating point values for altitude; and (2) applying a simple smoothing filter of radius 3.

2.4.2 Terrain layers

Five new terrain layers were created from the DEM using the SAGA's digital terrain analysis functions. The RSAGA package was used to create terrain layers for slope, solar radiation and wetness index. Layers showing the index values for multi-resolution valley bottom flatness (MRVBF) and multi-resolution ridge top flatness (MRRTF) were created interactively using the SAGA GUI as these functions are not available in RSAGA. The derivation of these five layers is described below and Figure 2-7 shows the maps and histograms of their values.

Slope

The slope layer was generated using the SAGA morphometry function for calculating slope, aspect and curvature, implemented in RSAGA as the function *rsaga.slope.asp.curv*.

Flat and gentle slopes dominate across the project region, as is shown by the histogram values in Figure 2-7. The mean slope is 1.68 degrees with standard deviation of 1.94, implying that 95% of the terrain has slope of less than 5.6 degrees and 68% has slope less than 3.6 degrees.

Solar radiation

The solar radiation layer was generated using the SAGA lighting function for calculating incoming solar radiation, implemented in RSAGA as the function *rsaga.pisr2*. This function will calculate the solar radiation at a given latitude using dates and time steps specified by the user. Latitude was specified as 10 degrees and the time variables were left to the default of half hour time steps on 31 October 2015 which is around the end of the growing season and the start of harvest.

Wetness index

The patterns of dispersal of rainfall runoff across a landscape can be digitally modelled from its DEM. The Topographic Wetness Index (TWI) is a commonly used measure of soil moisture availability that is derived from the catchment area size and local slope at sites in a landscape (Mattivi, Franci et al. 2019). However, the TWI was designed for hillslope catenas and so the physical concepts behind it are not valid for channels and extremely flat areas. The SAGA Wetness Index (SWI) is similar to the TWI, but it is based on a modified catchment area calculation that tends to assign more realistic and higher potential soil wetness to those grid cells situated in valley floors and having a small vertical distance to a channel (Brenning, Bangs et al. 2018). The wetness index values reveal the relative variation in soil moisture availability across the landscape but are themselves unit-less.

An important pre-processing step in automated DEM-based modelling of surface rainfall runoff is the creation of a hydrologically sound elevation model. This is done by identification and removal of surface depressions (referred to as 'filling sinks') and preservation of downward slopes along flow paths. One of the methods implemented in SAGA for filling surface depressions is based on the concept of spill elevation and uses the least-cost search technique to progressively build the optimal flow paths (Wang and Liu 2006).

A hydrologically sound DEM was generated using the SAGA terrain analysis pre-processing function that fills sinks using the Wang Liu method, implemented in RSAGA as the function *rsaga.fill.sinks* with the parameter *method="wang.liu.2006"*. The wetness index layer was then generated from this new DEM using the SAGA hydrology function for calculating the SAGA Wetness Index, implemented in RSAGA as the function *rsaga.wetness.index*.

Multi-resolution valley bottom flatness

The distinction between hillslopes and valley bottoms can be very important in agriculture. It was observed in Table 2-3 that the bas-fonds (valley bottoms of seasonally flooded inland valleys) had a very distinct cropping profile compared to other soil types and were used almost exclusively for growing rice.

Distinguishing valley bottoms from hillslopes is important in hydrologic and geomorphic analysis to separate erosional and depositional areas and in characterising sediment deposits (Gallant and Dowling 2003). The multi-resolution valley bottom flatness (MRVBF) measure was developed to identify valley bottoms based on their topographic signature as flat and low-lying relative to their surroundings. The algorithm to compute MRVBF is described in Gallant and Dowling (2003) and uses a combination of local lowness and flatness indices computed at a range of scales from a DEM. The MRVBF is a continuous measure that divides into classes corresponding to the different resolutions and slope thresholds as follows:

- Values less than 0.5 are not valley bottom areas.
- Values from 0.5 to 1.5 are considered to be the steepest and smallest resolvable valley bottoms for 25m DEMs.
- Flatter and larger valley bottoms are represented by values from 1.5 to 2.5, 2.5 to 3.5, and so on.

Both the MRVBF and MRRTF layers for the project region were created interactively in SAGA using the morphometry function for calculating Multi-resolution Index of Valley Bottom Flatness. It can be observed in Figure 2-7 that there is a strong correspondence between high MRVBF values and high wetness index values in areas at lower elevations.

Multi-resolution ridge top flatness

MRRTF is a topographic index designed to identify high flat areas at a range of scales that complements the MRVBF index. Zero values for MRRTF indicate areas that are steep or low, with values 1 and larger indicating progressively larger areas of high flat land (Gallant and Dowling 2003). There is a strong correspondence between high MRRTF values and high wetness index values in areas at higher elevations (see Figure 2-7).

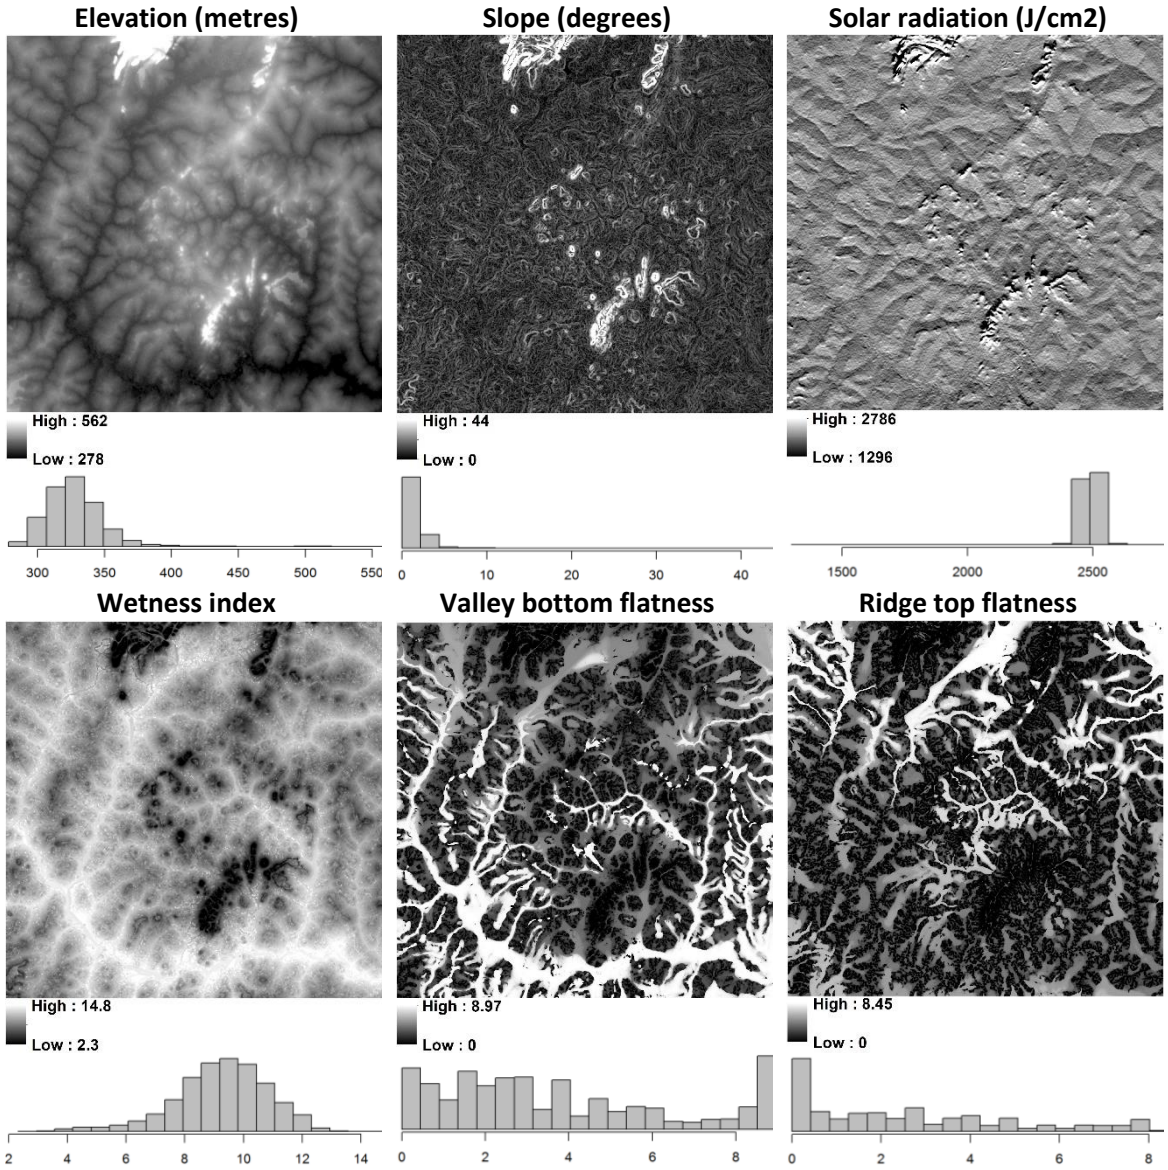

Figure 2-7 Maps and histograms of values for terrain layers

2.5 Soils

The project region is situated in a geological domain dominated by volcanic and volcano-sedimentary belts and is regarded as the result of juvenile crust formation and the reworking of older crust that has been disaggregated (Belousova 2016 ref). In West Africa large areas are characterised by soils with surface layers that can become hardened by iron and clay compounds. Figure 2-8 has been reproduced from the Soil Atlas of Africa (Jones, Breuning-Madsen et al. 2013) and shows the major soil types occurring in south-west Burkina Faso. In the project region Plinthosols (marginal soils with iron accumulation that harden irreversibly when exposed to air and sunlight) are dominant and there is a localised area of Acrisols along the border with Ivory Coast (these are strongly acidic marginal soils that also harden irreversibly). However, the central third of the area is characterised as Cambisols (young soils that are moderately developed) that are classed as excellent soils for agriculture.

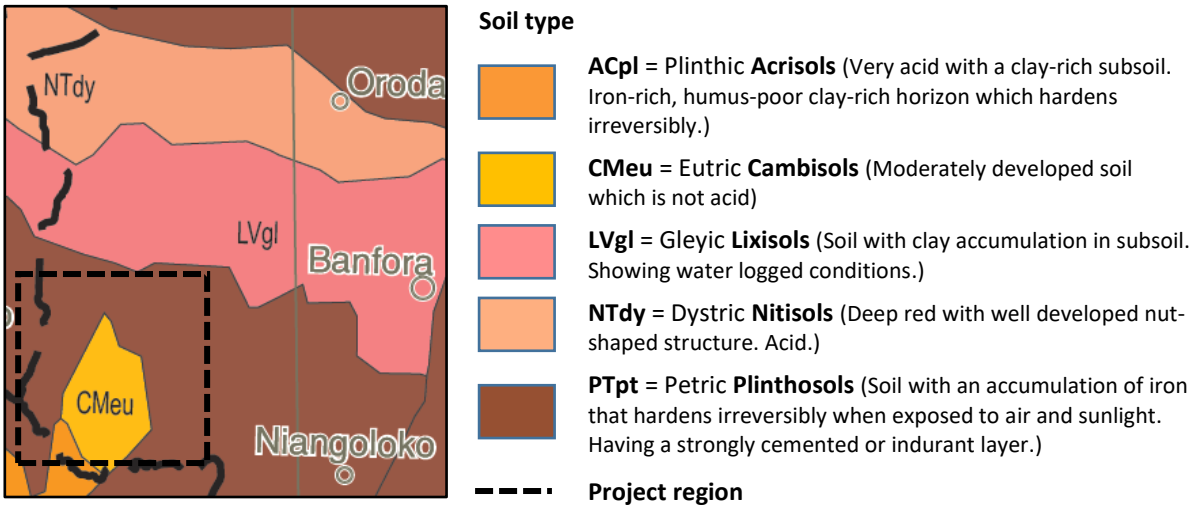

Figure 2-8 Soils of south-west Burkina Faso, from Soil Atlas of Africa (plate 9, pages 96-97)

The maps available in the Soil Atlas of Africa are informative at a regional level, but lack the fine detail required for land suitability analysis at community level.

2.5.1 Gryphon Minerals soil map

In contrast, the soil map of the project region supplied by Gryphon Minerals Ltd (shown in Figure 2-9) contains very fine spatial detail. It is derived from a regolith map and identifies eight major soil types grouped into four soil and terrain forming regimes (Bolster 1999). These are:

- **curaisse** (soil codes Fp and Fh) – often referred to as “desert armour” and having a hard crust formed from the weathering of iron-rich soils;
- **erosional** (soil codes Eo and Es);
- **depositional** (soil codes Da, Dc and Dd); and
- **residual** (soil code Rs).

The detailed descriptions for the soil types are shown in the map legend. The distribution of these soil types across the project area and in the mining lease communities (both individually and combined) is presented in Table 2-6, following the map.

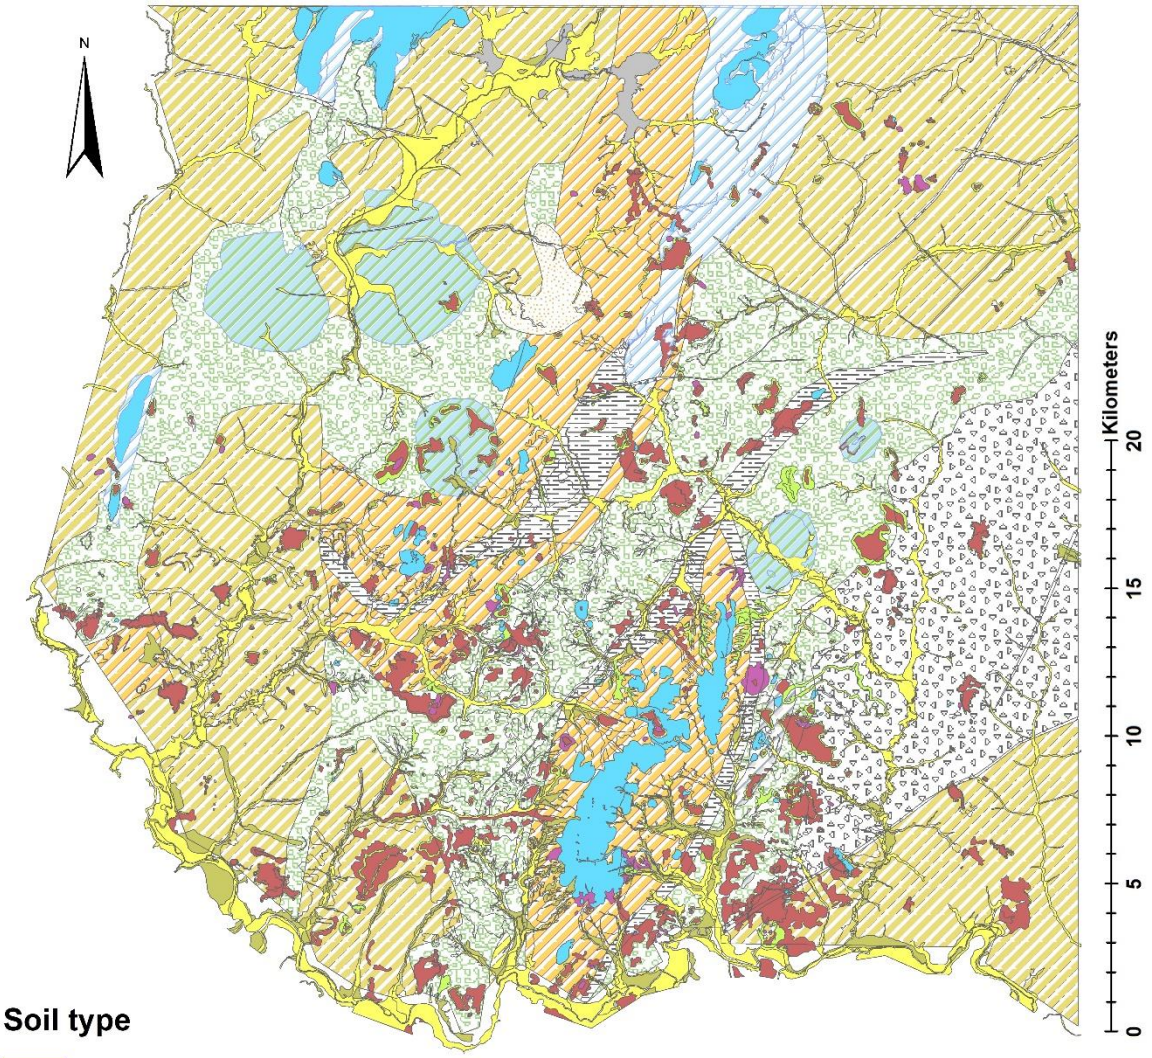

Soil type

- Da. Alluvium
- Dc. Colluvial clay and sandy soils
- Dd. Swamp / black soil plains
- Eo. Skeletal soils and outcrop
- Es. Latosol / pedolith and saprolith on escarpments
- Fh. Cuirasse dominated hills and slopes over pedolith / saprolith
- Fp. Cuirasse plateaus with abundant hard duricrust and pisolitic gravels
- Rs-Bas-pyr. Residual lateritic soils after basaltic and andesitic volcanics including proximal volanoclastic facies
- Rs. ME1. Residual lateritic soils after variably foliated granitoids of TTG affinity
- Rs-ME2. Residual lateritic soils after syn-tectonic granodiorite-granite
- Rs-ME3. Residual lateritic soils after post tectonic granite-monzonite stocks
- Rs-Taoudeni. Residual soils and outcrop after Phanerozoic siliclastic sandstone
- Rs-dol dyke. Residual soils after dolerite dykes
- Rs-dol-gab-dio. Residual soils after gabbroic to dioritic intrusions including dolerites
- Rs-gab-prox. Residual soils after coarse grained pyroxene-gabbro
- Rs-sed. Residual soils after phyllite and arkose
- Rs-vs. Residual soils after Fine grained mafic to intermediate composition tuffaceous sediments

Figure 2-9 Soil map of the project region supplied by Gryphon Minerals Ltd[11]

[11] **Note**: Villages and artisanal mining sites were also identified in the supplied map, but accounted for less than one percent of the total area (refer Table 2-6). The ArcGIS geoprocessing tool *Dissolve* was used to simplify the map by merging these tiny polygons with neighbouring soil polygons.

As can be observed in Table 2-6, residual soils make up 80% of the landscape. These soils have weathered in place and so their chemical makeup is determined by their underlying geology. There are ten distinct classes of residual soils identified in the Gryphon Minerals soil map, with mafic, felsic and intermediate residual soils prevailing. Depositional soils account for almost 10% and occur along the river valleys. Curaisse plateaus and curaisse dominated hills account for almost 5% of the landscape. The remaining 5% of soils are in the erosional regime either as skeletal soils and outcrop, or erodible soil on escarpments. The Stinger mining lease, with 50% of its area being curaisse plateau, has a very different soil profile from the other mining leases and from the project area as a whole. The Fourkoura mining lease also has a different soil profile, having a much higher proportion of depositional soils (18%).

Table 2-6 Distribution of soil types in the Burkina Faso project area and the four mining lease communities (%)*

Soil code	Soil description	FOURKOURA	NOGBELE	SAMAVOGO	STINGER	Mining leases	Project area
Da	Alluvium	8.8	1.9	3.0	2.7	3.0	7.0
Dc	Colluvial clay and sandy soils	9.4	3.5	0.0	0.0	2.9	2.0
Dd	Swamp / black soil plains	0.0	0.0	0.0	0.0	0.0	0.3
Eo	Skeletal soils and outcrop	0.4	0.2	0.0	0.0	0.2	2.9
Es	Latosol / pedolith and saprolith on escarpments	0.3	1.7	2.4	0.6	1.5	1.7
Fh, Fp	Cuirasse (primarily plateaus)	4.9	2.5	2.6	49.8	9.3	4.5
Rs	Residual soils	74.8	89.3	91.7	46.3	82.4	80.8
Vil, Orp	Not classified (village/town, artisanal gold mining)	1.4	0.9	0.3	0.6	0.8	0.8
Total		*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*

* Percentages based on numbers of pixels in a 30m x 30m cell raster of the supplied soil map

Table 2-7 Distribution of soil types in the planting locations of crops in the four mining lease communities (%)*

Soil code	Beans	Cashew	Cotton	Cowpea	Earth pea	Maize	Millet	Peanut	Potato/yam	Rice	Sesame	Sorghum	Total
Da	1.3	1.6	0.5	2.7	0.8	2.3	0.3	0.3	7.8	38.5	0.3	1.9	3.0
Dc	3.5	0.0	0.5	0.2	0.5	2.7	0.4	0.8	3.3	22.0	0.8	0.1	2.9
Dd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eo	0.0	0.0	0.1	0.5	0.0	0.1	0.0	0.0	0.0	0.0	0.3	0.0	0.2
Es	0.3	0.1	0.9	0.0	0.3	0.7	0.0	0.9	0.0	0.4	0.5	0.3	1.5
Fh, Fp	15.2	1.1	7.5	0.7	9.1	6.0	11.8	11.3	16.7	3.7	9.7	5.2	9.3
Rs	78.4	96.1	90.0	95.7	87.8	87.6	86.7	86.0	72.2	35.2	88.2	91.2	82.4
Orp, Vil	1.3	1.1	0.5	0.2	1.6	0.7	0.8	0.6	0.0	0.2	0.3	1.2	0.8
Total	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*	*100.0*

* Percentage based on numbers of pixels in a 30m x 30m cell raster of the supplied soil map

In Table 2-7 the soils used for the planting locations of the twelve main crops are compared. Residual soils are predominant in the planting locations for most of the crops. However, Rice, has a very different soil profile from the other crops, with the majority of rice fields (60.5% by area) being on depositional soils (alluvium and colluvium).

The Gryphon Minerals soil map was produced using a combination of field work and the remote interpretation of high-resolution satellite imagery and topographic data, supported by radiometric and geophysical image interpretation. It was supplied as an ESRI shape file with the attributes identifying the soil and terrain forming regime and the underlying geology (as can be seen in the map legend), but lacking any of the attributes that would typically be used in agricultural suitability modelling (such as soil depth, pH, salinity, and the proportions of clay, sand, silt, carbon etc).

2.5.2 SoilGrids

Global soil maps that predict many of these soil properties are now freely available. In 2014, ISRIC (International Soil Reference Information Centre) released a global soil information system called SoilGrids that provided predictions of soil properties and soil classes at 1km spatial resolution. These predictions were refined to produce SoilGrids250m released in 2016, the development of which is documented in Hengl, Mendes de Jesus et al. (2017).

SoilGrids250m is a global 3D soil information system providing predictions at 250m spatial resolution for many standard numeric soil properties at different depths, plus depth to bedrock and soil classes in both World Reference Base (WRB) and USDA classification systems. The predictions of soil qualities (in raster format) were made at seven standard depths (0, 5, 15, 30, 60, 100 and 200cm). [12]

The SoilGrids250m layers likely to be relevant to this project were the soil properties and depth to bedrock layers (listed in Appendix D). In order for these files to be usable in the project, initial processing was required to download the global raster layers (119 tiff files, each approximately 1 Gb in size), clip each raster to the project area, and then resample and reproject at 30m pixel resolution (resulting in 119 rasters varying in size from 1.4 Mb to 17 Mb).

Suitability modelling using seven depth predictions for each soil property was not practical. Instead, an estimate for the top 30cm of soil was produced (following the recommendation by Hengl, Mendes de Jesus et al. (2017)) by numerically integrating the layers using the trapezoidal rule:

$$\frac{1}{b-a}\int_a^b f(x)dx \;\approx\; \frac{1}{b-a}.\frac{1}{2}\sum_{k=1}^{N-1}(x_{k+1}-x_k)(f(x_k)+f(x_{k+1}))$$

where N is the number of depths, x_k is the k-th depth and $f(x_k)$ is the value of the soil property at depth x_k. This resolves to a simple weighted average of the soil property values in the top four layers, easily implemented as the raster calculation:

$$\frac{1}{60}(5(sl_1+sl_2)+10(sl_2+sl_3)+15(sl_3+sl_4))$$

where sl_i is the cell value in the i–th soil layer for the target soil property.

[12] SoilGrids250m map layers are available for download via www.SoilGrids.org under the Open Database License (ODbL). GeoTiffs can also be obtained from ftp://ftp.soilgrids.org/data/.

Two examples of the resulting soil layers are shown in Figure 2-10. Other layers are displayed as maps in Appendix D. Figure 2-11 presents boxplots of the values in these two soil layers showing the median, interquartile range and outliers for: (a) the project area as a whole and each mining lease community; (b) the planting locations for the twelve main crops; and (c) each of the soil types in the Gryphon Minerals soil map. Analysis of such boxplots for all the new soil layers (not reproduced in this thesis) revealed considerable variation in raster values between the different soil types and crop planting locations.

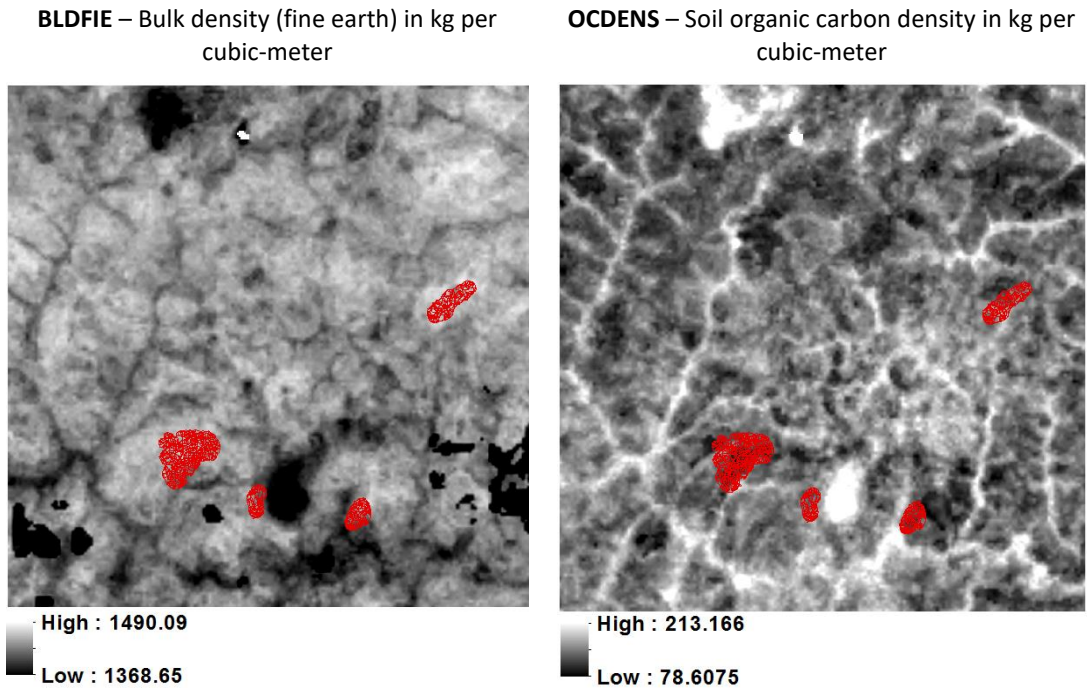

Figure 2-10 SoilGrids250m layers resampled to 30m pixels and consolidated for the top 30cm of soil

Maps for the new multi-dimensional set of soil layers derived from the SoilGrids250m data are provided in Appendix D. If the SoilGrids250m source maps are reliable for this region then these new SoilGrids layers may provide useful inputs for the agricultural suitability modelling task and overcome the limitations of the categorical map of soil type supplied by Gryphon Minerals.

However, caution must be exercised when using these new soil layers. The SoilsGrids250m predictions have not been validated globally and, in particular, have not been validated for this local area, so their accuracy cannot be relied upon. In addition, the spatial resolution of the source data (250m cell size) is much coarser than processing resolution (30m) and so may not capture important changes in soil characteristics occurring at the finer scale.

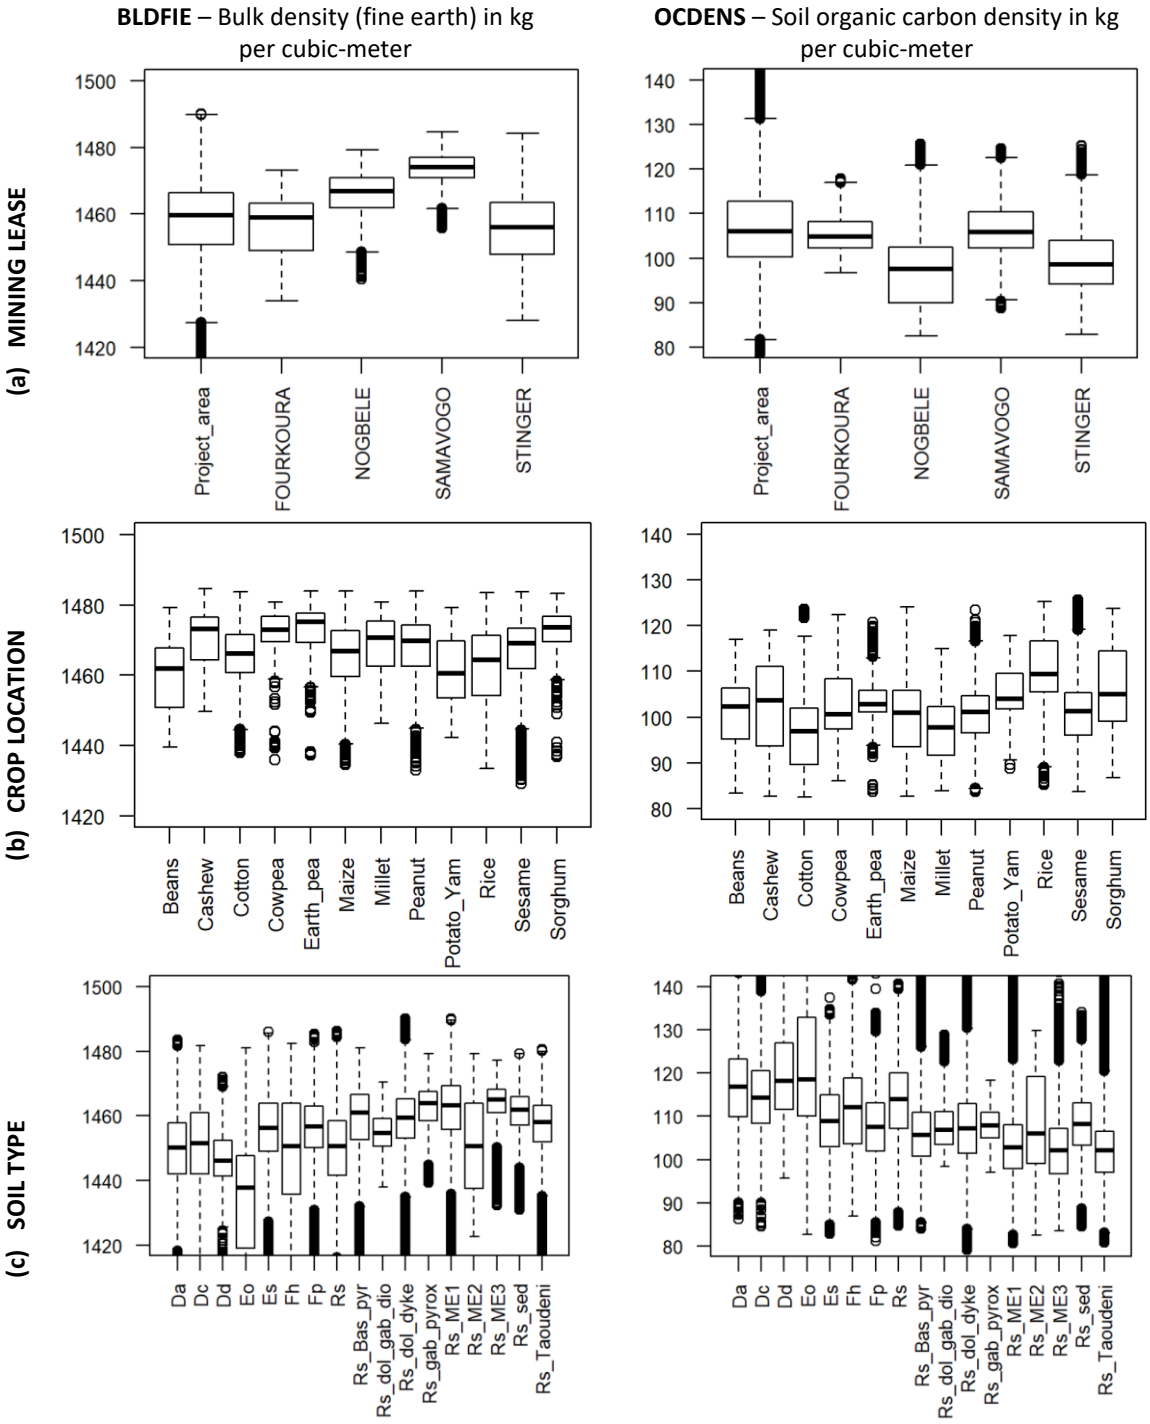

Figure 2-11 Analysis of SoilGrids values for the region and by mining lease, crop location and soil type

2.5.3 Hybrid soil layers

Attempts to use the Gryphon Minerals map of soil types and the SoilGrids layers in the agricultural suitability modelling (described later in Chapter 3) each produced very unsatisfactory results. The two sources of soil data are compared in detail in Figure 2-12 and Figure 2-13. Each figure shows high-resolution satellite imagery for a small cultivated area alongside the corresponding soil polygons from the Gryphon Minerals soil map and the raster values from four SoilGrids layers. The influence of the soil type on agriculture is clearly visible in both figures.

In Figure 2-12, the boundary between residual and depositional soils is clearly observable in the satellite image and is marked by the change in style of agriculture. Intensive rice cultivation occurs on the alluvial soils whereas the residual soils support dryland cropping and agroforestry. The impediment to agriculture posed by the erosional escarpments is also observable in the image as uncleared bush.

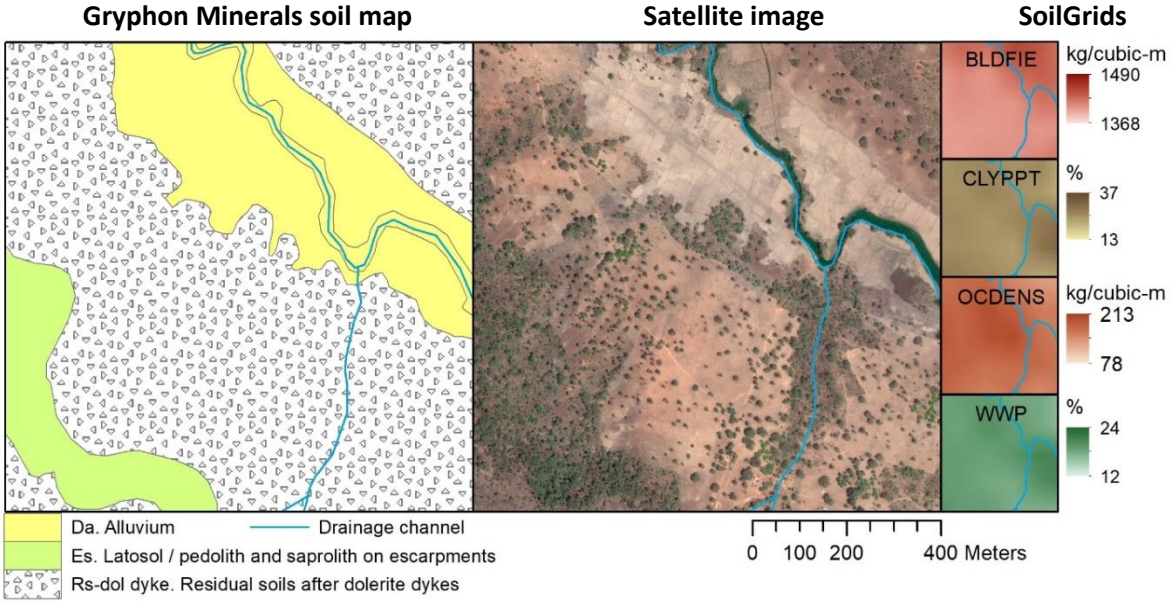

Figure 2-12 Comparison of Gryphon Minerals soil map and four SoilGrids layers with Worldview2 satellite image (true colour) April 2010

Again, in Figure 2-13, the influence of soil type can be observed in the cultivation patterns. Here, the rice fields visible in the satellite image match closely the large polygon of colluvial clay and sandy soils from the Gryphon Minerals map. Agroforestry of other dryland crops occurs on the residual soils and the erosional escarpments have not been cleared.

In comparison, the corresponding SoilGrids maps reveal no readily observable relationships between the SoilGrids values and (1) the soil types from the Gryphon Minerals map, or (2) the agricultural patterns and topographic features visible in the satellite images. Figure 2-11 demonstrated that the distribution of SoilGrids values does vary between soil types and also between the planting locations for different crops. However, the changes in soil type observable in the figures above, that are very relevant to agricultural land suitability modelling, could not be detected using just the SoilGrids data. To be able to include useful soil data in the modelling it became desirable to devise a way to combine the detailed and accurate soil map supplied by Gryphon Minerals with the coarse resolution but multi-dimensional soil data from SoilGrids.

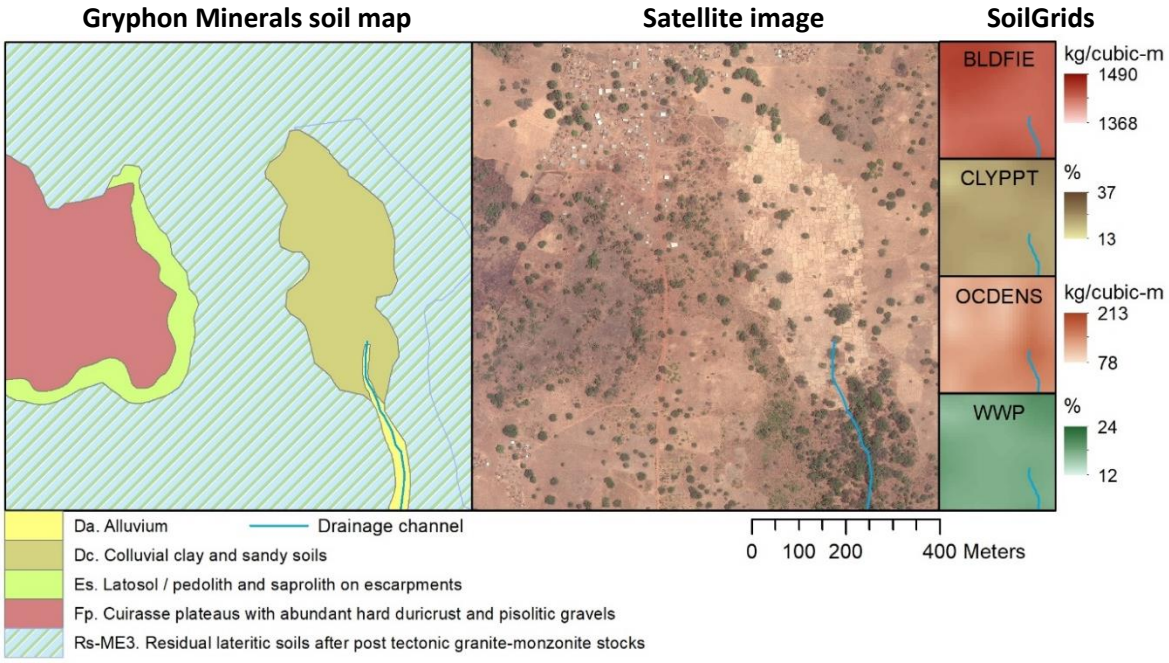

Figure 2-13 Comparison of Gryphon Minerals soil map and four SoilGrids layers with Worldview2 satellite image (true colour) April 2010

New sets of hybrid soil layers were created using the soil polygon boundaries from the Gryphon Minerals map and the SoilGrids values for the respective soil types. Figure 2-14 illustrates this process. The top left map shows the SoilGrids values for bulk density (fine earth) in kg per cubic-meter (BLDFIE) in the top 30cm of soil. The top right map is a raster derived from the Gryphon Minerals map that plots the mean BLDFIE value for each soil type. The general pattern of low BLDFIE values in areas of skeletal and depositional soils and higher values for residual soils is sharpened in the new map. However, with only 17 distinct soil types, this is still essentially a categorical map. To address this limitation variability was introduced to the layers using three different methods:

Adding Gaussian noise - a noise layer was generated using random values from a normal distribution with mean=0 and standard deviation=0.5. To prevent anomalous outliers, any noise values with magnitude greater than 2 were reset to zero. A raster of the standard deviation of BLDFIE values for each soil type was also generated (centre left map). The final hybrid soil layer (centre right map) was generated from the mean, noise and standard deviation (std) rasters using raster arithmetic:

*hybrid_1 = mean + noise*std*

Adding SoilGrids values in the ratio 50:50 - the bottom left map shows the mean by soil type raster (top right map) added to the corresponding SoilGrids raster (top left map) in equal proportion:

*hybrid_2 = mean*50% + SoilGrids**50%

Adding SoilGrids values in the ratio 75:25 - the bottom right map shows the mean by soil type raster (top right map) added to the corresponding SoilGrids raster (top left map) in the proportion 75% to 25%:

*hybrid_3 = mean*75% + SoilGrids*25%*

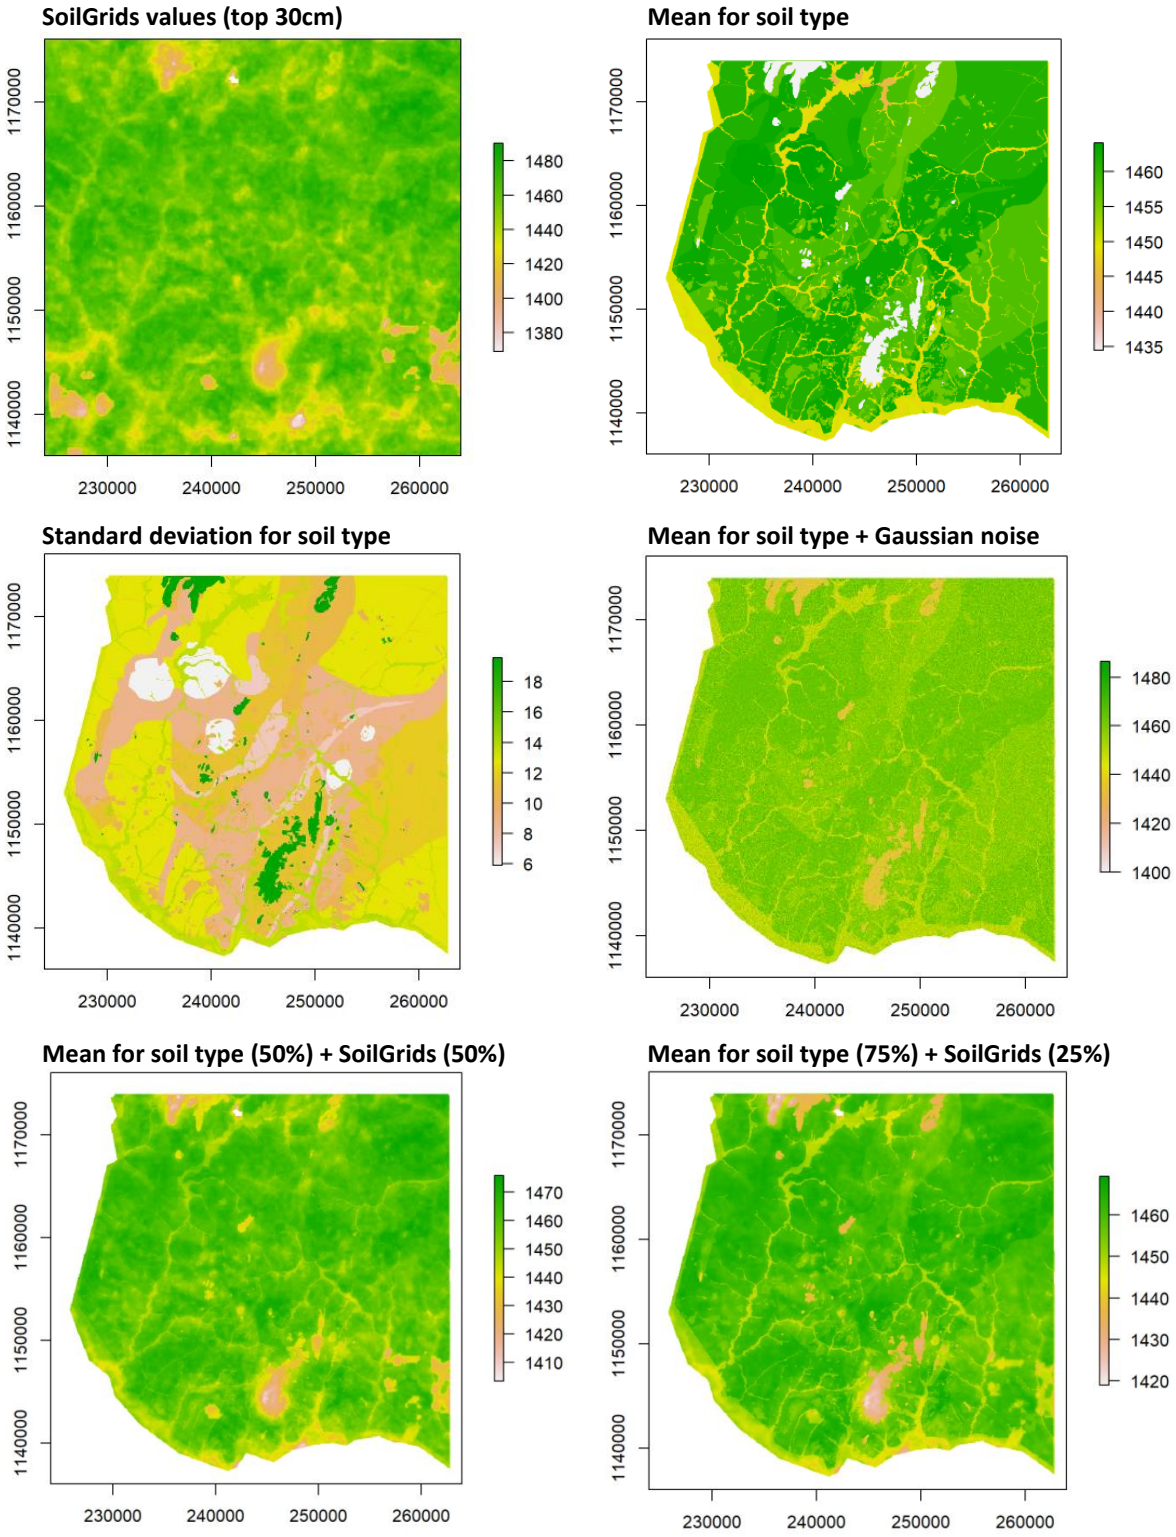

Figure 2-14 Derivation of hybrid soil layers from Gryphon Minerals soil map and the SoilGrids (top 30cm) BLDFIE layer

Sets of twenty hybrid soil layers were derived in this manner for potential use in the agricultural suitability modelling (see Table 2-8). Each raster layer reflects the precise landscape detail from the Gryphon Minerals soil map and is informed by the SoilGrids250m soil quality values averaged by soil type.

Table 2-8 List of hybrid soil layers derived from SoilGrids250m

Layer name	Soil quality	Description
ACDWRB	Acid grade	Grade of a sub-soil being acid e.g. having a pH < 5 and low BS
BDRICM	Bedrock depth	Depth to bedrock (R horizon) up to 200cm
BDRLOG	Bedrock depth	Probability of occurrence (0-100%) of R horizon
BDTICM	Bedrock depth	Absolute depth to bedrock (in cm)
BLDFIE	Bulk density	Bulk density (fine earth) in kg per cubic-m
OCDENS	Carbon - density	Soil organic carbon density in kg per cubic-m
ORCDRC	Carbon - organic	Soil organic carbon content (fine earth fraction) in g per kg
OCSTHA	Carbon - stock	Soil organic carbon stock in tons per ha
CECSOL	Cation capacity	Cation exchange capacity of soil in cmolc/kg
PHIHOX	pH acidity	Soil pH x 10 in water
PHIKCL	pH KCl	Soil pH x 10 in Potassium chloride
CLYPPT	Texture - clay	Clay content (0-2 micrometre) mass fraction in %
CRFVOL	Texture - coarse	Coarse fragments volumetric in %
SNDPPT	Texture - sand	Sand content (50-2000 micrometre) mass fraction in %
SLTPPT	Texture - silt	Silt content (2-50 micrometre) mass fraction in %
AWCh1	Water available	Available soil water capacity (volumetric fraction) for h1
AWCh2	Water available	Available soil water capacity (volumetric fraction) for h2
AWCh3	Water available	Available soil water capacity (volumetric fraction) for h3
AWCtS	Water available	Saturated water content (volumetric fraction) for tS
WWP	Water available	Available soil water capacity (volumetric fraction) until wilting point

2.6 Radiometrics

Radiometrics (gamma ray spectrometry) measures the natural radiation at the Earth's surface and a radiometric survey typically measures the spatial distribution of three radioactive elements (potassium - K, thorium - Th and uranium - U) in the top 30-45cm of the Earth's crust. The abundances of K, Th and U are measured by detecting the gamma rays produced during the natural radioactive decay of these elements. Since weathering modifies the concentration and distribution of the radioactive elements relative to fresh bedrock, radiometric measurements provide information on geomorphic processes and properties of the regolith and soils (IAEA 2003).

A low level airborne geophysical survey of the project area was commissioned by Gryphon Minerals Ltd, with magnetic, radiometric and elevation data acquired during the period 10-19 April 2010. The radiometric component of this survey (only) was supplied to the research project. It comprised eight raster files in ER-Mapper data format showing gamma ray counts and estimated ground concentrations as follows:

- Total count (counts/sec)
- Dose rate (nGy/hr)
- Potassium (counts/sec)
- Potassium ground concentration (%)
- Thorium (counts/sec)
- Thorium ground concentration (ppm)

- Uranium (counts/sec)
- Uranium ground concentration (ppm)

Each raster had cell size of 25m x 25m and had been projected to the WGS 1984 UTM Zone 30N coordinate system. To maintain consistency with the terrain files, all rasters were reprojected to 30m x 30m cell size.

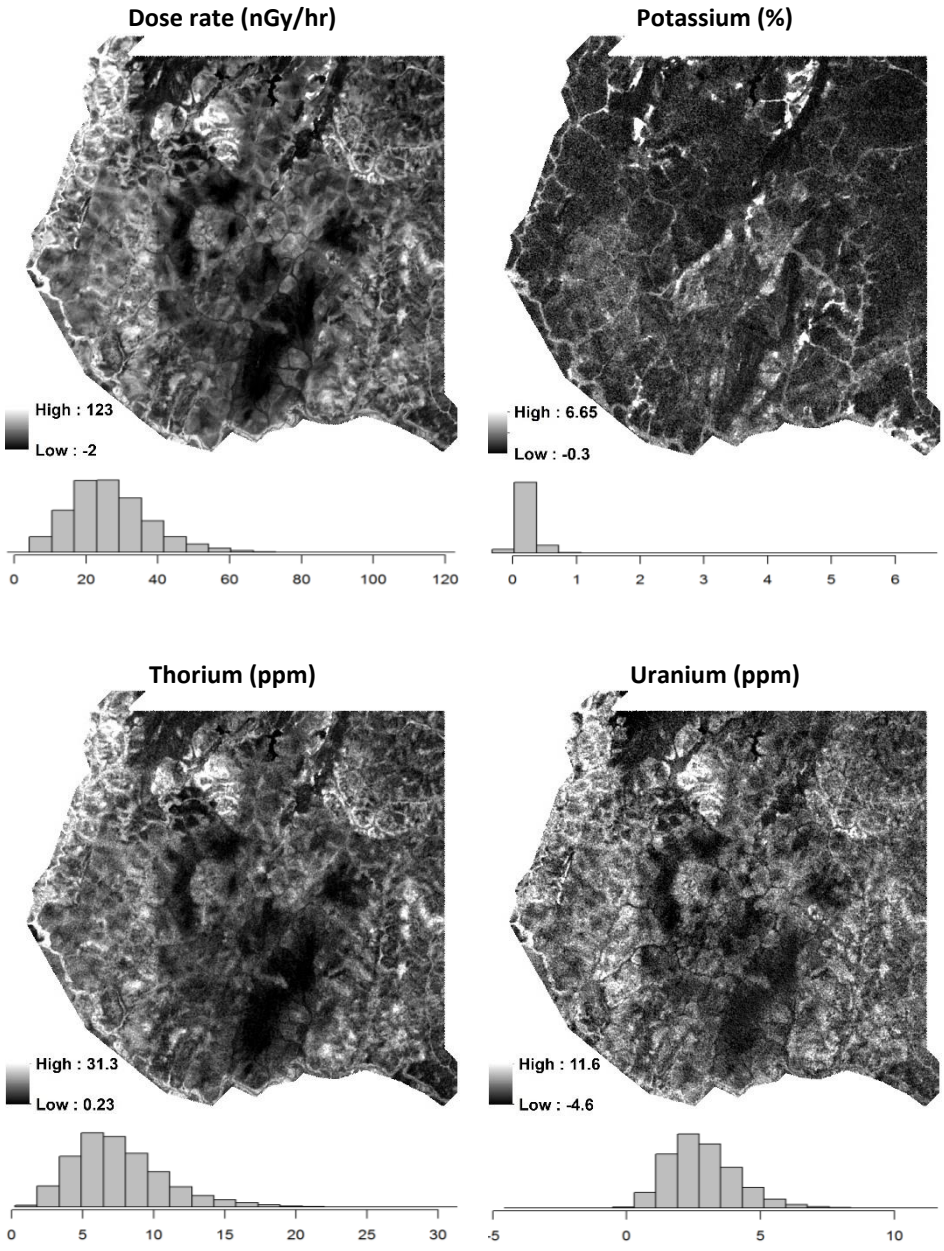

Figure 2-15 Maps and histograms of values for radiometric layers

The maps and histograms in Figure 2-15 show the numeric ranges and spatial distribution for the intensity of radioactive dose rate, the percent ground concentration of potassium, and the ground concentrations for thorium and uranium estimated in parts per million. The maps for dose rate, thorium and uranium exhibit a similar spatial pattern of low and high values, with the lowest values occurring in areas with high altitude and rugged terrain. Higher concentrations of potassium are observable along some of the water courses and drainage channels.

The radiometric response differs for fresh bedrock, weathered rocks and transported material and so these data may be used to identify areas of active weathering and distinguish erosional and depositional terrain forming processes. However, the gamma ray response to weathering is typically specific to the local geology so radiometric data are usually interpreted in conjunction with other sources of information such as topography, aerial photography and satellite imagery (IAEA 2003). The radiometric data described in this section were used by Gryphon Minerals Ltd in refining the soil map discussed in Section 2.5.1, but may also be useful for the agricultural land suitability assessment task.

2.7 Other data

Appendix A lists the mining exploration spatial data that were supplied for use in this research project. The community land use maps, the regional soil map, the DEMs and the radiometric data have been described in this chapter. The other datasets, not used directly to generate environmental layers for suitability modelling, are described below.

2.7.1 Satellite imagery

The very fine resolution WorldView2 imagery is extremely useful when trying to gain insight into the terrain and land use across the region. Detail from the 2010 dry season image that covers the entire project region was seen in Figure 1-4, Figure 2-12 and Figure 2-13. Similar satellite imagery for parts of the project region in 2012, 2013 and 2014 was also supplied by Gryphon Minerals Ltd but has not been used in this project.

Note that the extensive cloud cover during the wet season hampers the use of satellite imagery to view this region during the growing season. Table 2-9 shows the average cloud cover by month for images from two high-resolution Earth observation satellites over the period 2010-2018. Images for wet season months had high levels of cloud cover peaking at an average of 78% for August. For this reason, no further satellite imagery was sourced for this project.

Table 2-9 Average cloud cover for images of the project region taken by the AIRBUS satellites SPOT and Pléiades during 2010-2018 (192 images), by month

Month	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Cloud cover	4.3%	6.9%	7.6%	26.5%	29.3%	49.2%	66.6%	78.1%	58.6%	12.5%	16.3%	8.2%

2.7.2 Drainage

The location of drainage channels is powerfully informative with regard to the shape of a landscape and this vector layer has been frequently used as an overlay to improve the clarity of terrain maps. In particular, it has been used in the map outputs from this project (described later in 0) to make these maps more intuitively understandable and so enhance their readability.

Although descriptive in their shapes, the drainage channels had no measurements to indicate the size or duration of flows. Referring back to Figure 1-3, a sparse network of drainage channels can be observed in the low flat areas, compared with a much denser network in rugged terrain. This suggests that distance to a drainage channel would not be a valid indicator for soil moisture in this terrain. The wetness index layer described in 2.4.2 provides much richer data on water availability.

2.7.3 Contours

The elevation contour lines have not been used analytically as all raster terrain layers have been derived from the DEM. However, they have been used as descriptive map overlays.

2.7.4 Regolith and soil samples

The regolith files initially looked very promising. However, they proved to be a raster tile set showing an image only (with no attributes) of a soil/regolith map with the locations for soil samples marked. An attributed version of this map was reported to exist, as was a detailed spreadsheet of soil sample properties. An effort was made to acquire these additional datasets from ex-Gryphon Minerals members and from Teranga Gold Corporation. Eventually, Teranga Gold Corporation supplied two Access databases containing soil sample results; however, they had a great deal of missing data and used undocumented codes rather than numeric values for the soil attributes. The code descriptions were able to be obtained from an ex-Gryphon Minerals member but proved unsuitable for interpolation to generate any environmental layers of interest.

2.7.5 Vegetation

An Excel spreadsheet containing the locations of 66 named trees was supplied. Unfortunately these locations did not match tree crowns in the satellite image so could not be used for spectral analysis or object detection, and there were insufficient numbers of trees of any particular species to use as presence data for modelling.

Gryphon Minerals Ltd also supplied shape files showing the distribution of Baobab, Kerite and Nere trees in the area around Nogbele. However, these were found to be the results of in-house trials of directed classification using WorldView-2 imagery. As the accuracy of these shape files could not be established they were not used in the project.

2.7.6 ASTER imagery

Thirty-one raster images derived from Advanced Spaceborne Thermal Emission and Reflectance Radiometer (ASTER) were supplied showing mineral distributions. The data files are enhanced compressed wavelet (.ecw) images with three raster bands. Only ten images display in ArcGIS (for alumite, b1_grey, calcite1, calcite2, carbonate, dictate_b5, iron, jarosite1, jarosite2 and kaolinite) with the other images triggering an “invalid raster dataset” message. No documentation was provided on how these images were derived from the 14 ASTER bands or what the values in each of the three bands represented and so they have not been used in the project.

2.8 Summary

This thesis asks whether land suitability assessments for agriculture can be done effectively using the data by-products of mining exploration, without the necessity for and cost of additional data acquisition. This chapter has examined in detail the spatial data supplied for the research task by the exploration company (Gryphon Minerals Ltd) and has also investigated other sources of freely available spatial data that could potentially be used to supplement it.

The annotated shape files of the crop compensation maps proved a rich source of data regarding the local agriculture in the affected communities, showing both the crops grown and their chosen planting locations. These data show the existing cropping patterns in four small areas of the project region and the locations used to grow individual crops can be spatially linked with other data to reveal the environments selected by farmers for their cultivation.

The topography of a landscape strongly influences its land use and suitability for agriculture. Many relevant land qualities are derivable from digital elevation models and six terrain rasters were created for the project region (elevation, slope, received solar radiation, and indices of wetness, valley bottom flatness and ridge top flatness). The supplied digital elevation data had been purchased by the exploration company and represented an expensive data acquisition at that time. However, since 2014, high-resolution topographic data resulting from NASA's Shuttle Radar Topography Mission in 2000 has been publicly available, and so acquiring a DEM no longer represents a cost to projects.

A recently updated and finely detailed categorical soil map that classified soil according to regolith forming regime and original rock type was supplied by the exploration company; however, these soil categories were not easily interpretable in an agricultural context. Global maps predicting multiple commonly used soil properties are publicly available, although mostly at very coarse spatial resolutions that would be unsuitable for this project. In 2016 SoilGrids released its 3-dimensional soil information system at 250m resolution and these data were downloaded to supplement the supplied soil map. However, the resolution of these SoilGrids rasters was still too coarse for effective modelling, and so algorithmic methods were devised to process both sources of soil data into new sets of hybrid soil layers (that combined the fine spatial detail of the supplied map with the multi-dimensionality of the SoilGrids soil property predictions) that might prove useful for modelling.

The radiometric survey commissioned by the exploration company provided further detailed and accurate source data potentially useful for modelling. However, other supplied data that had been purchased or explicitly created by the exploration company were found to be useful for their descriptive value only (high resolution satellite imagery and the map of drainage channels) or insufficiently documented to be able to be used in modelling (regolith image tiles, vegetation data, ASTER imagery). Spatially relevant climate data were not supplied and were not elsewhere available (due to the very sparse network of weather stations in Burkina Faso).

Chapter 3 Environmental niche modelling and Maxent

3.1 Introduction

This research is performed in the context of a new mine in West Africa that necessitates the relocation of subsistence farmers from four sites. It asks the question:

> *Can land suitability assessments for agriculture be carried out effectively using the data by-products of mining exploration, without the necessity for and cost of additional data acquisition?*

If the answer to this question is yes, then meaningful assistance in determining suitable alternative farming locations could be offered to the affected farmers by the mining company. The major challenge with this research is in determining *how* such a land suitability assessment can be performed given the available data.

Chapter 2 has documented the main data resources available for the research. These were supplied by Gryphon Minerals Ltd (the exploration company for the new mine) at the start of the project or were derived from publicly available data sources. Attempts to obtain further data of use to the project are constrained by several significant issues:

- In 2016, Gryphon Minerals Ltd was taken over by Teranga Gold Corporation and it is now very difficult to obtain further exploration data not included in the original data supply.
- It will not be possible to visit the project region to inspect environmental conditions, to talk with agricultural experts and affected farmers, or to ground truth results, due to the project region bordering Mali and falling within an Australian Department of Foreign Affairs and Trade defined "do not travel" zone[13].
- Global datasets, e.g. for land cover, land use or other environmental factors typically have coarse spatial resolutions of inappropriate scales for this analysis and their accuracy for the project region cannot be relied upon as they have not been locally validated. In addition, most of the freely available satellite data commonly used in identifying growth patterns of crops proves ineffective for this region due to the extensive cloud cover during the growing season.

Reviewing the available data resources, some secondary research questions emerge:

> *How can existing cropping patterns inform us with regard to other potential planting locations for particular crops?*
>
> *Are the soil categories used for mining exploration useful as soil categories for agricultural land suitability models?*

The agricultural compensation maps provide a rich data source for identifying significant crops and existing cropping patterns. Using the FAO land evaluation approach (refer Section 1.4.1), the physical requirements of these crops would be matched against measured land qualities such as climate, soil type, landform, water availability, etc. to assess the degree of suitability of particular land for particular crops. However, closer inspection of these crop data reveals sufficient

[13] Parts of Burkina Faso are currently rated "do not travel" due to the threat of kidnapping and terrorist attack.

ambiguity in the crop descriptions to undermine rigorous application of this approach. For example, “rice” could refer to high yield Asian varieties or native African varieties grown using irrigated, dryland or floating cultivation (National Research Council 1996). Does “potato” refer to white potatoes or sweet potatoes? Which species of legume are referred to by the term “bean”? Are multiple crop varieties planted by farmers and do these varieties differ in their responses to the environment? Without the opportunity to clarify such questions it is not possible to reliably specify the physical requirements of the crops.

Attempts to craft a GIS-based solution (refer Section 1.4.2), e.g. using multi-criteria decision making (MCDM), are also hampered by the prohibition on travel to the region. Effective MCDM solutions rely heavily on expert opinion to ensure that input layers and overlay weights are realistic and to assist with the interpretation of final results. The travel ban makes finding such experts almost impossible, and also prevents in situ validation of results if such an approach was tried.

The limitations to further data collection and access to relevant experts preclude a knowledge-based solution in this project. In contrast, data-driven approaches that use empirical methods to derive the species-environment relationships are not similarly constrained by a lack of expert knowledge. These methods correlate known occurrences of a species with the environmental characteristics at those locations and then extrapolate these relationships to identify other areas that have similar environmental conditions. They simply require the locations for known occurrences of a species and geospatial environmental data for the target region to match it against. And very detailed crop location data are available to the project in the form of the crop compensation maps.

3.2 Exploring data-driven environmental niche modelling using the BCCVL

Environmental niche models predict suitability for a species as a function of given environmental variables. The *fundamental niche* for a species is the set of conditions that allow for its long term survival, whereas the *realised niche* is the subset of the fundamental niche that it actually occupies. The environmental conditions at the presence locations, by definition, constitute samples from the realised niche, and the realised niche in relation to agricultural crops are those places chosen for their cultivation.

The BCCVL (Biodiversity and Climate Change Virtual Laboratory) offers a useful platform for experimenting with species distribution and environmental niche modelling algorithms (refer Section 1.4.3) and was used to explore their performance in relation to the research task.

Initial tests of the ten BCCVL statistical regression and machine learning algorithms used presence data for the twelve main crops (derived as the centroids of the fields in the crop compensation maps) and three terrain derived environmental layers (slope, solar radiation, wetness index). The suitability maps generated by seven of these algorithms showed clearly the influence of terrain on the areas identified as more suitable for individual crops.[14] The results were consistent with common sense expectations regarding suitable planting locations for particular crops and so were very encouraging as evidence that this approach was appropriate for the task.

[14] The ANN model failed in all runs so produced no output, and the suitability maps generated by the Boosted Regression Tree (BRT) and Random Forest (RT) models were not meaningful.

The use of absence data was explored. Most of the SDM algorithms allow absence data to be supplied by the user but will simulate absence data if the user does not supply it. An absence dataset was explicitly created by selecting points in the landscape unlikely to support cropping (steep slope, skeletal soils, swampy). The resulting suitability maps using these absence data were very different from the initial tests, with these poor modelling results most likely due to the exaggerated differences between the presence and explicitly (but naively) defined absence data. Following these results, user defined absence data were not used in further tests.

The number of pseudo-absence/background points generated by the models is determined by the number of presence points and the absence ratio. Increasing the absence ratio from 1 (the default) to 20 improved the appearance of the suitability maps giving smoother results with less speckling. However, further increasing the absence ratio to 200 resulted in no further discernible improvement.

The creation of effective soil layers was challenging and ultimately unsuccessful during this exploratory phase. The soil map provided by Gryphon Minerals has categorical data (17 soil types). The BCCVL statistical and machine learning models can use environmental layers with categorical values, so the soil map was converted to a raster and uploaded to the BCCVL as a categorical layer. The models were run using the three terrain layers plus the new soil layer. The resulting suitability maps were very unsatisfying. The shapes of the soil polygons dominated the maps producing intuitively unrealistic results. As some soil types did not occur at the training data sites the algorithms would have no data to assess the influence of these soil types on cropping locations. The suitability maps generated using the rasterised soil map may well have been accurate for the four training data sites (this was not investigated), but they were a very poor solution for the region.

The unsuitability of the supplied soil map led to the search for an alternative source of soil data. The SoilGrids global layers for soil properties had been released at 250m resolution in 2017 – these were downloaded and processed to create suitable environmental layers for the modelling (refer Section 2.5.2). Five of these new layers were uploaded to the BCCVL and tests were run, first using the five layers alone, and again with the addition of the three terrain layers from the initial tests. These results were also very unsatisfying. There was no evidence of the influence of terrain factors in the blotchy suitability maps. As many soil forming processes are related to terrain, the results were also clearly unrealistic.

Experimentation with different environmental layers and combinations of layers exposed serious problems with interpretation and comparison of suitability maps. There was very limited knowledge about agricultural conditions outside of the four training data sites and intuition was an unreliable guide due to the unfamiliar landscape and styles of agriculture. Another site would be needed to properly test the methodology – one with validation data for the entire region; that was familiar to allow intuitive assessments; and that could be visited if necessary to verify results or collect more data. A parallel test region located near Adelaide in South Australia was defined (refer Chapter 5 for more detail). To duplicate the conditions of the Burkina Faso task, the same environmental layers were created and presence data for six locally grown crops was simulated in four small training sites. Any crop suitability maps generated by the BCCVL algorithms could be readily compared to published data on land use potentials for these crops allowing region-wide validation of results.

At the time of testing the BCCVL had not fully implemented a bulk download function for the outputs of its experiments. With twelve crop models generated for each experiment for the Burkina Faso site and six crop models for the Adelaide site, managing the manual download of

outputs for post-processing was an unwieldy task. So, although the BCCVL had proven to be a valuable tool for exploring the multiple data-driven algorithms and experimenting with combinations of environmental layers, it was not a suitable platform for the main research.

The major conclusions drawn from the BCCVL testing were as follows:

1. Presence data reflecting cropping patterns are derivable from the crop compensation maps;
2. Absence data are not required;
3. Most of the SDM models (excluding ANN, BRT and RT) are appropriate for the agricultural land suitability task;
4. The Generalised Boosting Model (GBM) and Maxent produced the best results;[15]
5. Neither the soil map supplied by the exploration company nor the generated SoilGrids layers were useful in their current forms;
6. A parallel site that allowed region-wide testing of results was needed to validate the methodology;
7. A more easily manageable software platform was needed.

3.2.1 Selected modelling algorithm (Maxent)

Of the ecological niche modelling algorithms tested, the Maxent algorithm seemed the most naturally suited to the agricultural land suitability task. Although often referred to as a presence-only model, it is technically a presence/background method (Peterson, Soberón et al. 2011) with the predictions interpretable as indices of habitat suitability (Merow, Smith et al. 2013). The open-source release of the Maxent Java software package (Phillips, Anderson et al. 2017) was tested and proved suitable for use in the main research.

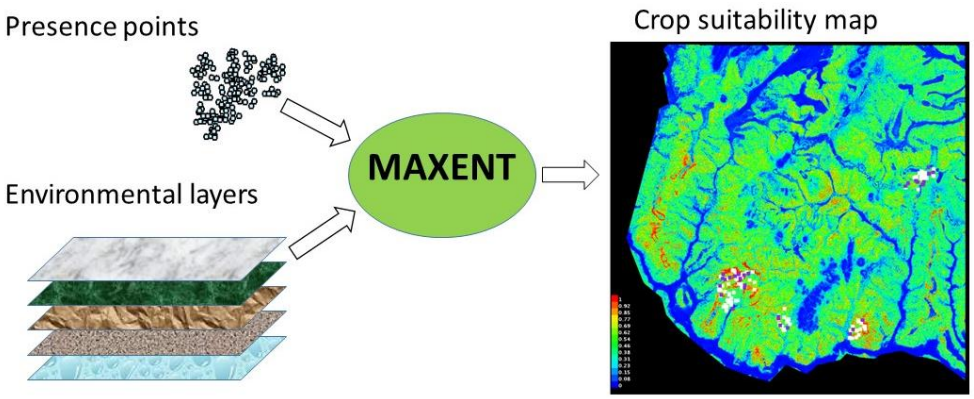

Figure 3-1 Maxent inputs and outputs. Warmer colours on the map indicate areas of higher suitability, and the presence points used for training and testing are shown in white and purple, respectively.

Figure 3-1 depicts the inputs and outputs for the Maxent software package and shows a typical crop suitability map for the West African study region. The presence points are the surveyed planting locations for that crop at the four mine sites (refer Section 2.2.1) and are visible on the

[15] As measured by the area under the curve (AUC) of the receiver operating characteristics (ROC) curve, refer section 3.3.6 below.

map. The environmental conditions occurring at the presence locations form the basis for the suitability predictions across the entire study region (shown in the map).

The experimental development of agricultural land suitability models for the crops grown locally in the West African project region (using the Maxent software package) is documented in Chapter 4. Chapter 5 then duplicates the methodology for parallel sites in South Australia for which validation data are available. The remainder of this chapter explains how the Maxent algorithm works to generate ecological niche models, describes the Maxent software package and identifies evaluation techniques for assessing models.

3.3 Maximum entropy modelling (Maxent)

Maximum entropy modelling is a method for making predictions from incomplete data that comes from Information Theory, where 'entropy' is used as a measure of how much choice is involved in the selection of an event (Shannon and Weaver 1949). The probability function of maximum entropy is the one that is closest to uniform that agrees with everything that is known without assuming anything that is not known. The maximum entropy algorithm used for modelling species' geographic distributions from presence-only data was developed by Steven J. Phillips at AT&T Labs with Miroslav Dudik and Robert E. Schapiro from the Department of Computer Science at Princeton University. It was presented first at machine learning conferences in 2004 (Dudík, Phillips et al. 2004, Phillips, Dudík et al. 2004) before being published for the ecological modelling community in 2006 (Phillips, Anderson et al. 2006). Its goal is to predict, from observations of a species' occurrence, those areas within a region that satisfy the requirements of the species' fundamental ecological niche and so form part of the species' potential distribution. It is a statistical machine learning approach that allows all reasonable predictors to be input to the model and lets the algorithm decide which ones are important in relation to the presence data (Phillips, Anderson et al. 2006).

3.3.1 Overview

The three components needed for this type of species distribution modelling are a *data model* of species presence, an *ecological model* of the study area and a *statistical model* to model these inputs and derive relationships (Phillips, Anderson et al. 2006). Inputs for the data and ecological models are formulated in geographical space – the study area is defined as a set of map pixels for which measurements of environmental variables are available, with species presences being sample points from this set, referred to as presence points. These data are projected into environmental space to model and compare: (1) the environment where species are found; and (2) all available environments across the study area. Available environments in the study area are characterised by randomly sampling a large number of pixels for which presence is unknown (default is 10,000) and adding the known presence locations to create a set of *background* points that are used in the statistical modelling. An environmental niche model for the species is then optimised for these data and projected back into geographical space to produce a map rating the suitability of the environment for the species at each pixel in the study area.

The goal of Maxent's algorithm is to estimate the species' potential distribution by finding the most uniform (spread out) probability distribution across the study area subject to constraints that are derived from the presence data.

3.3.2 Covariates / features

The covariates for the model in their original geographic form are raster layers for environmental variables. These are transformed by Maxent into an expanded set of simple real-valued functions referred to as *features* to model the distribution of the species in environmental space. Maxent has six feature classes as shown in Table 3-1. The *constraints* of the model are derived from the presence data and require that the expected value of each feature over all background points be close (within some error bounds governed by a regularisation parameter) to the empirical mean over the presence points. All features are rescaled to the interval [0,1] to make coefficients comparable.

Table 3-1 Maxent feature classes (Phillips, Anderson et al. 2006, Phillips and Dudík 2008)

Feature class	Description	Samples
Linear	A continuous variable	2+
Quadratic	The square of a continuous variable	10+
Product	The product of two continuous variables	80+
Threshold	Step function based on a continuous variable, that has value 1 when the variable value is above a given threshold and 0 otherwise	80+
Hinge	Similar to threshold but allowing a linear response below or above the threshold: • A *forward hinge* has value 0 below the threshold and increases linearly to 1 above the threshold • A *reverse hinge* has value 1 at the minimum value and decreases linearly to 0 at the threshold	15+
Categorical	A binary feature is created for each categorical value, with feature value of 1 for that categorical value and 0 otherwise	15+

Maxent can build complex response curves by using multiple features for a variable. The feature classes used for a model by default depend on the sample size for presences (refer Table 3-1) but may also be explicitly chosen. One linear and one quadratic feature are constructed for each predictor, product features are constructed for each pair of predictors, and binary features are created for each categorical value. The number of possible piecewise features (threshold and hinge) depends on the number of presences, with Maxent permitting these features between each pair of successive values of a predictor. Only the most useful features from this potentially enormous collection are retained by Maxent for the final model (Merow, Smith et al. 2013).

Note that hinge features were not part of the original release of Maxent and were added later to model arbitrary piecewise linear responses to environmental variables (Phillips and Dudík 2008). With their introduction the default use of threshold features has been omitted, resulting in improved model performance generally and models that are smoother and simpler and hence likely to be more realistic (Phillips, Anderson et al. 2017). These authors also found that the use of product features barely improved average performance and could usually be omitted in order to make simpler and more easily interpretable models. Quadratic features, on the other hand, are recommended by Merow, Smith et al. (2013) who observe that ecological theory suggests response curves for fundamental niches are often unimodal.

3.3.3 How it works

To produce the final probability distribution the Maxent algorithm starts with a fully uniform distribution over all background points and conducts an iterative optimisation routine to maximise the probability of the presences. Without regularisation the error bounds for all feature constraints are 0. In this situation the distribution with maximum entropy[16] (among all possible distributions satisfying the constraints) has been proved to be the same as the unique Gibbs distribution maximising the likelihood of samples (Dudík, Phillips et al. 2004). This distribution is exponential in a linear combination of features and takes the form:

$$P(x) = \frac{e^{c_1 f_1(x) + c_2 f_2(x) + c_3 f_3(x) + \cdots}}{Z} \qquad (3\text{-}1)$$

where the c_i are constants, the f_i are features and Z is a scaling constant that ensures P sums to 1 over all background points.[17]

The likelihood of the samples indicates how closely the model is concentrated around the presences. It is expressed in terms of *gain* which is defined as the average log probability of presence samples minus a constant that makes the uniform distribution have zero gain. Gain starts at zero (the gain of the uniform distribution) at the beginning of a Maxent run and increases with each iteration (as the weights on the features are adjusted) until the change in gain between consecutive iterations falls below a convergence threshold or until a specified maximum number of iterations have been performed. This final distribution becomes the basis for the predictor variable coefficients that are used to estimate the probability of presence (Phillips 2017).

The Maxent algorithm adjusts the weights on the features sequentially (one at a time) rather than in parallel, so is robust for large number of features. It starts with the uniform distribution (c_i=0 for all f_i in equation 3-1 above). At the start of each iteration Maxent calculates (for all features) the change in gain resulting from a small change in each feature's weight (whilst keeping all other weights unchanged). The feature producing the largest such change is selected for weight adjustment during that iteration. This boosting-like approach permits the selection of the best target feature from a potentially very large number of features, so allowing the use of very large feature spaces (Dudík, Phillips et al. 2004, Phillips, Dudík et al. 2004).

Regularisation

Regularisation is a technique often used in optimisation problems to obtain a solution that is less likely to be the result of overfitting. Maxent uses a form of L_1-regularisation to lower gain by an overfitting penalty equal to a weighted sum of the absolute values of the model coefficients. This is the sum $\sum_i \beta_i |c_i|$ where each β_i is a non-negative constant corresponding to a relaxed constraint for feature f_i. The relaxed constraint specifies that the expectation of f_i over all background points need only to be within β_i of the empirical mean of the presence points, rather than equal to it.

The use of L_1-regularisation has several significant benefits for model building. Firstly, regularisation enhances model accuracy by reducing model variance. Secondly, the use of the L_1

[16] For a probability distribution p over domain X, the entropy of p is defined as $\mathrm{H}(p) = -\sum_{x \in X} p(x) \ln p(x)$

[17] This maximum likelihood exponential model can also be obtained from an inhomogeneous Poisson process (IPP). A brief overview of the IPP formulation followed by discussion on its implications for interpretation of model inputs and outputs is given in Phillips, Anderson et al. 2017.

norm (sum of absolute values) in solving optimisation problems promotes sparser solutions (making use of fewer features) with model coefficients that are either zero or not near zero. In addition, the inclusion of this weighted sum of coefficients in the regularised gain penalises the use of large feature weights and so encourages smaller model coefficients.

> *"Such models are less likely to overfit, because they have fewer parameters; as a general rule, the simplest explanation of a phenomenon is usually best (the principle of parsimony, Occam's razor)."* (Phillips, Anderson et al. 2006)

The feature error bounds (betas) are intended to reflect the sample variation. However, sample data are often biased and so, for this reason and to simplify model fitting, Maxent instead uses pre-tuned default beta values based on sample size and feature class (Elith, Phillips et al. 2011). The empirical tuning of these values is described in Phillips and Dudík (2008) and the resulting regularisation parameters used by Maxent are shown in Table 3-2.

Table 3-2 Regularisation parameter settings (reproduced from Phillips and Dudík (2008) Table 3)*

	Number of occurrence records					
Feature classes	**0**	**6**	**10**	**17**	**30**	**100**
Linear features: β_L	*1.0*	**1.0**	**1.0**	0.72	**0.2**	**0.05**
Linear and quadratic features: β_L, β_Q	*1.3*	**1.0**	**0.8**	**0.5**	**0.25**	**0.05**
Linear, quadratic and product features: β_L, β_Q,β_P	*2.6*	**2.0**	**1.6**	**0.9**	**0.55**	**0.05**
Threshold features: β_T	*2.0*	*1.94*	**1.9**	*1.83*	*1.7*	**1.0**
Category indicators, a single categorical variable, "low": β_C	0.2	**0.2**	**0.2**	**0.1**	**0.05**	*0.05*
Category indicators, "intermediate": β_C	0.65	0.53	0.45	0.25	0.15	0.05
Hinge features: β_H	**0.5**	**0.5**	**0.5**	**0.5**	**0.5**	**0.5**

** Parameter settings tuned using presence-only data. Values in boldface were determined exactly, values in italics are linearly interpolated or extrapolated, with the exception that the values to the right of the listed ranges remain constant. For category indicator features, the "low" settings were determined using a single (categorical) variable, while the intermediate settings were chosen to approximate the geometric average of the "low" setting and β_L. (Phillips and Dudík 2008)*

The default beta values can be scaled up or down together by multiplying all values by a user-specified constant referred to as the regularisation multiplier. Scaling of the beta values will result in models having more/less generality or complexity. Maxent's experimental settings also allow the beta values for feature classes to be explicitly specified, with separate beta values available for the threshold, categorical and hinge feature classes, and a fourth beta value applying to the linear, quadratic and product feature classes. However, Merow, Smith et al. (2013) suggest that the default regularisation values (having been chosen based on performance across a wide range of taxonomic groups) may be more useful when building models for multiple species simultaneously.

3.3.4 Model calibration

Model calibration is done with the goal of producing predictive models that fit well to known data but do not overfit in ways that cause low predictive ability for independent data. The calibration and evaluation of ecological niche modelling algorithms is discussed in detail in Chapters 7 and 9 of Peterson, Soberón et al. (2011), with particular attention to the issues of data splitting, variable

selection, model complexity and overfitting. Aspects of model calibration relevant to Maxent are discussed below.

Data splitting

An important calibration task is the selection of appropriate data on which to train a model and suitable test data for evaluating its predictive performance. Data splitting is the process of partitioning the known presence locations into two sets used respectively for model calibration (training) and model evaluation (testing). Ideally, evaluation data are fully independent from calibration data (points are not spatially or temporally correlated). However, fully independent datasets seldom exist and so other techniques are used to split the data. For example:

- Random test percentage – the algorithm randomly selects a specified percentage of presence points to use as test data and uses the remaining presences to train the model.
- Spatially structured partitions – presence points are separated into calibration and evaluation datasets by geographic area. This technique can be useful in identifying overfitting due to sampling bias.

Replication involves building multiple models using different calibration/evaluation datasets and allows for cross-validation of results. Common replication techniques (all offered by Maxent) include:

- *K*-fold cross-validation – data is randomly divided into *K* pools. *K* models are then built by setting aside one of the pools and using it for evaluation.
- Subsampling – sample the data points *without* replacement to create a validation set then use the remaining data points to calibrate a model. Typically many iterations (500-1000) are done.
- Bootstrap – sample the data points *with* replacement to create multiple calibration and validation datasets from the same set of presence points.

K-fold cross-validation improves on the single split approach because each presence point will appear in one evaluation dataset, thus making better use of small datasets (Phillips 2017). The subsampling and bootstrap techniques are typically used when very large numbers of presence points are available.

Variable selection

The selection of the environmental variables to use for a model should be based on biological reasoning and, where possible, include specific variables known (or suspected) to affect a species distribution. In general, using larger numbers of environmental variables provides more information on which a model can be based and enables more complex models to be built, whereas smaller numbers of variables can help to avoid overfitting by limiting model complexity (Peterson, Soberón et al. 2011, p. 115).

Many data-driven modelling algorithms (refer Section 1.4.3) require the predictors to be uncorrelated, with multi-collinearity being a problem particularly for statistical methods (Merow, Smith et al. 2013). This can necessitate the removal of correlated variables from the predictor set, or the use of uncorrelated predictor variables created using techniques such as principal component analysis (PCA).

In contrast, Maxent can handle unlimited numbers of environmental predictors and its algorithm is robust to correlated variables. The sequential optimisation routine, together with the

regularisation method, decide which predictors are most important in relation to the presence data and the less important variables from the calibration set are excluded from the final model.

Although undemanding with regard to prior knowledge of the importance of environmental variables and interactions between them, Maxent can produce many types of output useful for evaluating environmental predictor importance (refer Section 3.3.5, below). Selecting these outputs in preliminary models can provide insight into which environmental layers are most significant and whether some make no contribution. Reducing the number of environmental layers can improve runtimes, particularly where replication is performed; and evidence for the significance of layers can inform decisions about data collection for similar modelling at other locations.

Tuning model parameters

All algorithms have their own calibration methods with sets of parameters and constants that can be adjusted to influence model performance. In the case of Maxent, the most important parameters affecting the structure of the final models are: (1) the allowable feature classes; and (2) the regularisation parameters (both discussed above). Other parameters that can be specified when running Maxent to influence the model building process include: the maximum number of background points; the random test percentage for data splitting; whether a different random seed is used for each run; the maximum number of iterations of the algorithm; and the convergence threshold for the change in gain.

Phillips and Dudík (2008) who undertook the empirical tuning for Maxent's default parameter settings (e.g. Table 3-2) observe that the resulting default settings achieve almost as good performance as tuned parameters sets. This is supported by Radosavljevic and Anderson (2014) who tested nine regularisation multipliers ranging from 0.25 to 10 for several different data splitting regimes and found a performance peak for the default value of 1. In a systematic literature review of Maxent parameter configuration for small sample sizes Morales, Fernández et al. (2017) found that less than 20% had user-defined feature classes and less than 10% had user-defined regularisation multipliers. However, Peterson, Soberón et al. (2011) warn:

> *"… although theoretical and empirical research may have led to the suggestion of default settings by the researchers developing the software, it is generally poor practice to use default settings provided by software without justification, testing and exploration of these values for a particular application."* (Peterson, Soberón et al. 2011, p. 113)

3.3.5 Outputs

The Maxent software package produces a suite of outputs that includes details for the generated model, statistical analysis of model performance, and a range of optional analyses that allow insight into the predictive importance of environmental variables. These outputs and the run-time parameters to generate them are well described in the Maxent tutorial written by Steven J. Phillips (Phillips 2017). The statistical analysis of model performance will be discussed later in Section 3.3.6. The other outputs are briefly explained here.

Environmental niche model

The primary output from Maxent is the map of predicted suitability for the modelled species across the study area. The model values correspond to the suitability scores and can be expressed using four alternative output formats. These formats are monotonically related (so they rank sites in the same order) but are scaled differently and have different interpretations, as follows:

1. Raw – this is the output from the Maxent exponential model. These values are probabilities in the range 0 to 1 such that the values for all cells used in training sum to 1 (making raw values very small).
2. Cumulative – for a particular raw value *r* the corresponding cumulative value *c* is equal to the percentage of the Maxent distribution that has a raw value less than or equal to *r*. The values for *c* are in the range 0 to 100.
3. Logistic - for a particular raw value *r* the corresponding logistic value is equal to $\frac{e^H r}{1+e^H r}$ where *H* is the entropy of the distribution. These values are also probabilities in the range 0 to 1 but are scaled up in a non-linear way for easier interpretation as a probability of presence.
4. Cloglog – for a particular raw value *r* and distribution entropy of *H* the corresponding cloglog value is equal to $1 - e^{-e^H r}$. Cloglog values are the default output format and can also be interpreted as estimates between 0 and 1 of probability of presence.

Both logistic and cloglog values can be interpreted as approximations of the probability of presence and so are useful for comparing models with different spatial scales (Merow, Smith et al. 2013).[18].Cumulative values are best interpreted in terms of predicted omission rates since a binary prediction generated using a cumulative threshold of *c* would result in an omission rate of *c%* on samples from the Maxent distribution and we would predict a similar omission rate for samples from the species distribution.

When replication is performed, the results of the cross-validation runs are averaged to obtain a single probability score for each location on the map.

Evaluating environmental predictor importance

A natural question to ask in relation to ecological niche modelling is: which environmental variables are most important for the species being modelled? Maxent can produce a range of outputs that allow insight into the predictive importance of environmental variables, as described below.

Response curves

Maxent produces two sets of response curves to show how each environmental variable affects the Maxent prediction. The first are *marginal response curves* that each show how predicted suitability changes when only a single environmental variable is changed and all other environmental variables are fixed at their sample mean value. These can be hard to interpret if there are strongly correlated variables so Maxent also creates a set of *single predictor response curves* derived from a set of new Maxent models each trained using a single environmental variable.

When replication is performed, error bars of one standard deviation are shown on summary response curves for the set of models.

Analysis of variable contributions

Percent estimates of the relative contribution of environmental variables to the final model are provided by the algorithm and accumulated during model training. Each iteration of the Maxent algorithm increases the gain by modifying the coefficient of a single feature. The program assigns

[18] For a more detailed explanation comparing the logistic and cloglog outputs refer to Phillips, Anderson et al. 2017.

the change in gain to the environmental variable that feature depended on and so keeps track of the contribution each environmental predictor makes to the overall gain.

Jackknife test of variable importance

When the jackknife test of variable importance is selected for a run, Maxent builds multiple models trained on different sets of environmental variables to compare predictive performance. For each environmental variable Maxent builds one model trained using all environmental variables *except* this variable, and another model trained using *only* this variable. The results are displayed in a set of three bar charts (for regularised training gain, test gain and AUC) that each compare the results of all these models to the model created using all variables.

Interactive exploration of predictions

The Explain tool is an interactive interface provided in Maxent for investigating how Maxent's prediction is determined by the predictor variables across the study area. The tool displays the map of the prediction model or of any environmental layer alongside the response curves for all environmental variables. Clicking on a point in the map shows its location in each response curve and the contribution of each variable to the prediction at that point.

The Explain tool assumes the model is additive (without interactions between variables) and so is only available for examining models created without product features.

3.3.6 Running the Maxent software package

The Maxent Java software package can be run either interactively or in batch mode under both the Windows and Linux operating systems. When run interactively, the Java program *maxent.jar* presents a graphical user interface (GUI) that allows the user to specify the input files, the location for the output files, and various run-time parameters. Alternatively, Maxent can be run in batch mode by invoking *maxent.jar* and specifying all run-time parameters at the command prompt, or by calling *maxent.jar* from within a shell script that is executed from the command prompt.

Input files

The inputs for Maxent are species presence points and raster environmental layers:

- Presence points: The presence data input files are text files (CSV format) containing the geographical locations (defined using either a projected or geographic coordinate system) for each known occurrence of the species to be modelled. For multi-species presence data the species names are also required. All presence data can be included in a single file with Maxent randomly partitioning the points into training and test data for the models. Alternatively, the presence data can be explicitly partitioned into separate training and test files.
- Environmental layers: Any number of environmental layers can be used in the models, but Maxent requires all environmental layers to be in identical ASCII raster grid format using the same coordinate system as the presence data.

Output files

Each run produces a log file and a summary spreadsheet showing key results for all species modelled in the run. Multiple output files are generated for each species model, as follows:

- Suitability map: an ASCII raster grid of the modelled region showing the environmental suitability ratings for the species across the region.

- Model analysis: an informative HTML file providing a description of the Maxent model for the species and containing embedded plots and pictures showing:
    - Analysis of omission/commission
    - Pictures of the model
    - Response curves for the environmental variables
    - Analysis of variable contributions
    - List of processing parameters
- Model coefficients: the coefficients (lambdas) and minimum and maximum values of the features used by the model.
- Threshold analysis**:** the fractional predicted area and rates of training and test omissions for various suitability thresholds.
- Prediction strength: the values predicted by the model for all training and test presence points.

3.4 Evaluation of species distribution models

Visual inspection of model outputs offers an intuitive means of assessing the credibility of models – *do the distribution maps of predictions seem plausible? do the response curves conform to known tolerances?* – but, to establish whether a model is statistically valid and of sufficient quality to meet the needs of a project, it must be formerly evaluated.

The predictive performance of a model is measured using evaluation statistics of two types: (1) measures of performance that characterise how well or poorly the model achieves its predictive goal; and (2) tests of significance that determine the level of probabilistic confidence with which the model predictions on evaluation data differ from random predictions. The evaluation of ecological niche models is discussed in detail in Chapter 9 of Peterson, Soberón et al. (2011). The authors present key concepts and commonly used quantitative measures for model performance and significance but advise that there is no single "best" approach to evaluation.

Maxent's predictions are reported using a continuous score in the range 0 to 1. In order to generate many evaluation statistics these scores must be converted to binary predictions of 1 or 0. This is done by choosing a threshold value for suitability (or presence) and classifying localities as either suitable (present, value of 1) when the prediction value is greater than or equal to the threshold, or unsuitable (absent, value of 0) when the prediction value is less than the threshold. Evaluation statistics based on binary classifications are then referred to as *threshold-dependent* when they apply to a single threshold and *threshold-independent* when they report over all possible thresholds.

The predictive performance of models is evaluated by comparing the model predictions against verification data. Maxent uses the presence points from model training and testing as verification data to calculate standard evaluation statistics for models. Verification data for a particular project may also be available in map format, allowing direct comparison with model predictions in geographic space. When a verification map contains suitability ratings for a modelled species then statistical measures of similarity between the verification map and the model outputs can be calculated. Correlations can be performed between the maps of model predictions and continuously-valued verification maps. For verification maps with binary suitability ratings, interrater comparisons with thresholded maps of model predictions can be performed.

Land surface observations, in the form of high-resolution satellite imagery, present another form of verification data that can be assessed in geographic space. In this case, subjective assessment

as to what is observed in the image at any location can be directly compared to model predictions at that location. These subjective assessments are valuable when considering the plausibility of models.

Each of the evaluation techniques used in this thesis is explained in greater detail in the following sections.

3.4.1 Threshold-dependent evaluation

Once a threshold has been set, the binary classification of the model prediction for a locality can be compared to its observed value. These results, for any set of evaluation data, can be summarised in a confusion matrix, as shown in Table 3-3, showing the numbers of evaluation points correctly or falsely classified by the model. A falsely classified data point may be an omission error (false negative) or a commission error (false positive).

Table 3-3 Confusion matrix for evaluation of binary classifications

		Observation	
		Present	Absent
Prediction	Present	*A* true positive	*B* false positive (*commission error*)
	Absent	*C* false negative (*omission error*)	*D* true negative

Note: *A*+*B*+*C*+*D* = *N* where *N* is the total number of points in the evaluation dataset.

Various threshold-dependent evaluation statistics can be calculated from the numbers in the confusion matrix (A = true positives, B = false positives, C = false negatives, D = true negatives), as follows:

- **Sensitivity** (true positive rate) = $\frac{A}{A+C}$
- **Omission error rate** (false negative rate) = $\frac{C}{A+C}$

  Note: sensitivity + omission error rate = 1
- **Specificity** (true negative rate) = $\frac{D}{B+D}$
- **Commission error rate** (false positive rate) = $\frac{B}{B+D}$

  Note: specificity + commission error rate = 1
- **Accuracy** = $\frac{A+D}{A+B+C+D}$
- **Misclassification rate** = $\frac{B+C}{A+B+C+D}$

  Note: accuracy + misclassification rate = 1

Sensitivity measures the responsiveness of a model by how well it correctly detects presences, whereas *specificity* measures the discriminatory capacity of the model by how well it correctly detects absences. An effective model and well-chosen threshold would result in high values for both.

Presence only model evaluation

When only presence data are available for evaluation (as in the case of Maxent) the specificity and commission error rate statistics cannot be calculated (since B=D=0 in the confusion matrix). In this situation, sensitivity equals accuracy, and the omission error rate equals the misclassification rate.

In order to provide a measure of specificity for presence only model evaluation, the true negative rate is instead quantified by considering the fraction of the region predicted as being unsuitable for the species (Phillips, Anderson et al. 2006). The *fractional predicted area* for a given threshold is the fraction of all pixels that are predicted as suitable for the species using this threshold. Specificity for presence only models is then defined as follows:

- **Specificity** (true negative rate) = 1 – fractional predicted area

Significance testing

The accuracy statistic characterises the performance of the binary model for the given threshold. Its statistical significance is measured using binomial probabilities. Suppose there are N test localities for a model with an omission error rate r and a fractional predicted area a. A one-tailed binomial test can test whether the model predictions at the test localities are significantly better than random predictions. The null hypothesis for the test states that the test points are predicted no better than by a random prediction with the same fractional predicted area. The test calculates the probability P of having at least $N(1 - r)$ successes out of N random trials, each with probability of success of a. This probability, called the *p-value*, is the probability of the observed result or better assuming the null hypothesis was true. A small p-value (typically $P \leq 0.05$) is evidence to reject the null hypothesis and accept the alternative hypothesis that the predictions are not random.

Selecting thresholds

The choice of threshold affects results considerably, as different thresholds will result in different binary distribution maps, different omission rates and different levels of statistical significance. Published methods for selecting the threshold of presence that are commonly used in species distribution modelling are summarised by Peterson, Soberón et al. (2011, p. 119). The Maxent program produces evaluation statistics for the following thresholds:

- Fixed cumulative value N (for N = 1, 5, 10) – threshold for which N% of pixels are classified as unsuitable.
- Minimum training presence – the lowest predicted presence value corresponding to a training presence record.
- 10 percentile training presence – threshold for which 10% of training locations are rated unsuitable.
- Equal training sensitivity and specificity – the threshold at which sensitivity and specificity are equal, calculated from the training data.
- Equal test sensitivity and specificity – as above, but calculated from the test data.
- Maximum training sensitivity plus specificity – the threshold at which the sum of sensitivity and specificity is maximised, calculated from the training data.
- Maximum test sensitivity plus specificity– as above, but calculated from the test data.
- Balanced – minimizes the weighted sum: (6 x training omission rate) + (0.04 x cumulative threshold) + (1.6 x fractional predicted area)
- Equate entropy of thresholded and original distributions

For each of these thresholds Maxent lists the corresponding cumulative and logistic threshold values, the fractional predicted area, training and test omission rates and the p-value for the thresholded model. Many of the thresholds listed above can only be determined by comparing the evaluation statistics for all possible thresholds.

Although thresholds are necessary for calculating the evaluation statistics described above and the threshold-independent evaluations described below, thresholds are usually unnecessary when presenting model results. Merow, Smith et al. (2013) instead recommend embracing the continuous and probabilistic nature of Maxent's model predictions, observing that this also avoids placing undue confidence in predictions.

3.4.2 Threshold-independent evaluation

Threshold-independent evaluation compares model performance over all possible thresholds. Performance graphs can be created by plotting pairs of evaluation statistics for all thresholds, allowing threshold independent metrics to be calculated from the combined results. Maxent generates two threshold independent performance graphs for each model to facilitate the analysis of model performance and to inform decisions regarding the selection of individual thresholds for categorically symbolised distribution maps. These graphs are described below.

ROC Curve and AUC

The receiver operating characteristic (ROC) curve is a commonly used performance graph for quantifying the performance of predictive models and is commonly used to measure performance of species distribution models. The ROC curve plots the true positive rate (sensitivity) against the false positive rate (1 – specificity) for all possible thresholds on the model predictions. The area under the curve (AUC) provides an overall measure of the model's performance at all thresholds. AUC values less than or equal to 0.5 suggest the model is no better at distinguishing presence from absence than random prediction. Values greater than 0.5 indicate that model performance is better than random, with higher AUC values indicating better predictive performance, and a value of 1 indicating perfect prediction.

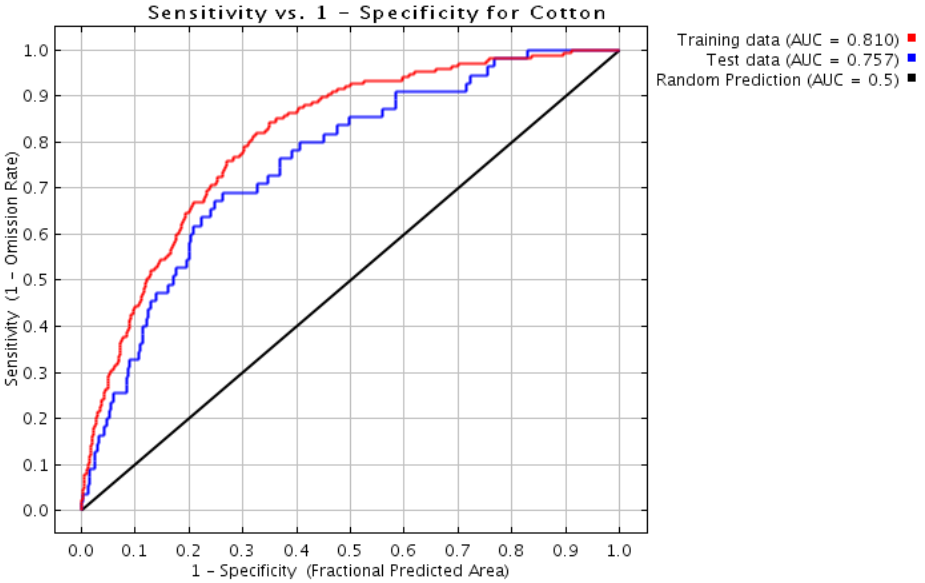

Figure 3-2 Receiver operating characteristic (ROC) curve – example from Maxent output

When only presence data are available the false positive rate cannot be calculated for the ROC plot. Instead, Phillips, Anderson et al. (2006) plot the true positive rate against the fractional predicted area (1 – specificity) to generate the ROC curve (refer Figure 3-2) and restate the classification problem as the task of distinguishing presence from random. The area under the curve (AUC) is then interpreted as the probability that a random positive instance is ranked above a random background instance.

A consequence of this reinterpretation of the ROC curve is that AUC values are lower for presence/background evaluations than for comparable presence/absence evaluations (Peterson, Soberón et al. 2011, pp 170-174). The maximum achievable AUC value for a Maxent ROC curve is less than 1 and is dependent on the prevalence of the species (Wiley, McNyset et al. 2003). In particular, Phillips, Anderson et al. (2006) state that "*If the species' true distribution covers the fraction α of the study area, then the maximum achievable AUC can be shown to be exactly 1 – α/2,*" but note that random prediction still corresponds to an AUC of 0.5, and that AUC scores higher than the maximum achievable AUC value may occur in practice. Table 3-4 shows the theoretical maximum achievable AUC value for species' ranges between 0% and 100% of the study area and demonstrates that higher AUC values are achievable for species with smaller ranges.

Table 3-4 Fractional predicted area and corresponding theoretical maximum achievable AUC

Fractional area	**α**	**0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1**
Maximum AUC	**$1 - \alpha/2$**	1	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.5

Although widely used as a measure of performance in species distribution modelling, the AUC has been criticised as being misleading. A review by Lobo, Jiménez-Valverde et al. (2008) questioned its reliability as a comparative measure of model results and did not recommend its use for five main reasons:

1. AUC scores ignore the actual probability values and the goodness-of-fit of the model.
2. The AUC summarises test performance over regions of the ROC space that are not useful in practice.
3. The AUC weights omission and commission errors equally, whereas misclassification errors frequently have different costs.
4. The ROC plots do not provide any information about the spatial distribution of errors so it is impossible to know if the errors are homogeneously distributed across the study area or if they reflect incapacity to predict in a specific region.
5. The geographical extent of the model highly influences AUC scores, with higher scores obtainable by increasing the extent outside the environmental domain of presence.

Jiménez-Valverde (2012) further observes that the AUC is not an appropriate performance measure when the goal of the research is to estimate the potential distribution of a species. Commission errors in relation to a realised niche for a species may not be errors in relation to its fundamental niche; however, they are weighted equally to omission errors in the AUC. In addition, the realised distribution for a species, being smaller, will have a higher achievable AUC score than the potential distribution.

The dependence of the AUC value on the prevalence of the species in presence/background evaluation invalidates it as a comparative measure for models of different species or over different geographical regions. However, AUC values can be validly used to compare the

performance of models for the same species and study area that were created using the same algorithm, e.g. that have been trained using different calibration datasets, environmental layers or parameterisations of the modelling algorithm (Peterson, Soberón et al. 2011).

Omission and predicted area

The omission rate and the fractional predicted area for a model are both threshold dependent, with higher threshold values resulting in smaller predicted ranges and higher omission rates. This is illustrated in Figure 3-3 which plots: (1) the fractional predicted area for the species against the cumulative threshold values; and (2) observed omission rates on the training and test data for each threshold value compared to the predicted omission rate.

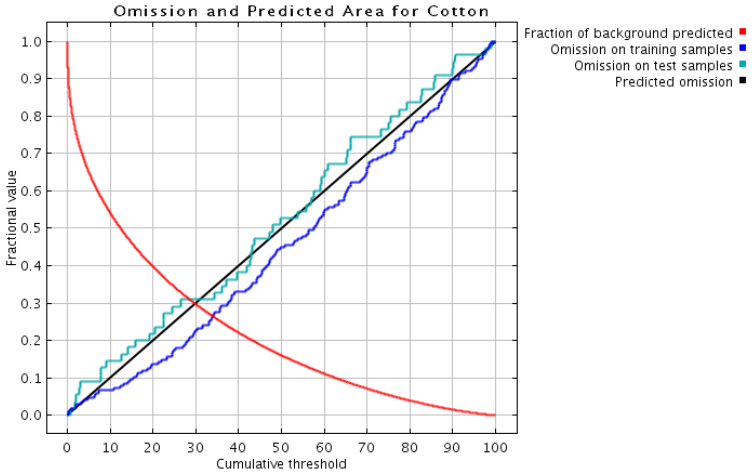

Figure 3-3 Graph of omission and predicted area – example from Maxent output

Note that the predicted omission rate is, by definition, equal to the cumulative value (refer Section 3.3.5) and deviations from it suggest sample bias.

3.4.3 Interrater reliability of binary suitability maps

A binary suitability map created by thresholding is referred to in statistics as a "rater". Each pixel in the map is rated as either suitable (present) or unsuitable (absent) for the target species. Interrater comparison of two thresholded suitability maps uses the confusion matrix of presences and absences to compare the predictions and allow calculation of evaluation statistics (refer Section 3.4.1).

The *accuracy* statistic measures the proportion of predictions that are the same in both maps. The accuracy rate is calculated by adding the numbers of true positives and true negatives from the confusion matrix and dividing by the total number of samples. The accuracy rate measures agreement of predictions due to both correct prediction and chance.

The *kappa coefficient,* proposed by Cohen (1960) and denoted using the lower case Greek letter κ, is a standardised measure of agreement between the categorical scores of two raters that corrects for chance agreement. Similar to the correlation coefficient, a kappa value of 1 indicates perfect agreement between the raters, and a kappa value of zero indicates random agreement (i.e. no more than would be expected by chance). Kappa values less than zero are possible and indicate interrater agreement that is worse than random.

The calculation of kappa for binary raters is presented in McHugh (2012) and summarised below using the terminology of the binary confusion matrix for raster map comparison.[19]

From the interrater comparison of two binary rasters, denote the numbers of pixels recorded in the confusion matrix by: A = true positives; B = false positives; C = false negatives; D = true negatives; and $N = A + B + C + D$ = total number of pixels. Let p_o denote the proportion of observed pixel agreement (accuracy) and p_e denote the proportion of pixel agreement expected due to chance. Then the kappa coefficient for the raster comparison is calculated as follows:

$$p_o = (A + D)/N \qquad \textit{= observed agreement (accuracy)}$$

$$p_e = (A + B)(A + C)/N^2 + (C + D)(B + D)/N^2 \qquad \textit{= expected chance agreement}$$

$$\kappa = \frac{p_o - p_e}{1 - p_e} \qquad \textit{= kappa coefficient}$$

When there is perfect agreement between the raster maps $(p_o = 1)$ then kappa will evaluate to 1, regardless of the value for p_e. When there is no more agreement than would be expected by chance $(p_o = p_e)$ then kappa will evaluate to zero. When the observed agreement is less than the expected agreement due to chance $(p_o < p_e)$ then the kappa value will be negative.

In the situation where one map represents a verification map against which to compare the predictions from alternative suitability models, the accuracy and kappa values from interrater comparisons with this map become performance measures that can be used to compare and rate the effectiveness of the models.

Comparing suitability maps with equal predicted areas

The threshold chosen for a continuous suitability map determines the fractional predicted area (FPA) for the resulting binary suitability map. When comparing two binary suitability maps with equal FPA, it is observed that:

- the true positive rate plus the false positive rate must equal FPA;
- the false positive rate must equal the false negative rate;
- the true positive rate cannot exceed FPA;
- the true negative rate cannot exceed (1 – FPA).

Table 3-5 tabulates the accuracy and kappa values for interrater comparisons of binary maps with equal FPA for a range of FPA values and different sensitivity values (true positive rate). Perfect agreement occurs when sensitivity equals FPA, always giving accuracy = 1 and kappa = 1.

From the table it is observed that the kappa value for a particular degree of accuracy varies depending on FPA. For example, for FPA of 50%, the kappa value for 80% accuracy is 0.6 (very good); but for FPA of 10% or 90% the kappa value for 80% accuracy is -0.11 (worse than random prediction).

[19] Accuracy and kappa values from interrater comparisons are calculated by the *confusionMatrix* function from the *caret* package in R.

*Table 3-5 Accuracy and kappa values from interrater comparisons of binary suitability maps with equal fractional predicted area; p_e is the expected accuracy due to chance agreement**

Fractional predicted area		Sensitivity = true positive rate										
		0.50	**0.45**	**0.40**	**0.35**	**0.30**	**0.25**	**0.20**	**0.15**	**0.10**	**0.05**	**0.00**
0.10	Accuracy									1.00	0.90	0.80
p_e = 0.82	*κ*									*1.00*	*0.44*	*-0.11*
0.20	Accuracy							1.00	0.90	0.80	0.70	0.60
p_e = 0.68	*κ*							*1.00*	*0.69*	*0.38*	*0.06*	*-0.25*
0.30	Accuracy					1.00	0.90	0.80	0.70	0.60	0.50	0.40
p_e = 0.58	*κ*					*1.00*	*0.76*	*0.52*	*0.29*	*0.05*	*-0.19*	*-0.43*
0.40	Accuracy			1.00	0.90	0.80	0.70	0.60	0.50	0.40	0.30	0.20
p_e = 0.52	*κ*			*1.00*	*0.79*	*0.58*	*0.38*	*0.17*	*-0.04*	*-0.25*	*-0.46*	*-0.67*
0.50	Accuracy	1.00	0.90	0.80	0.70	0.60	0.50	0.40	0.30	0.20	0.10	0.00
p_e = 0.50	*κ*	*1.00*	*0.80*	*0.60*	*0.40*	*0.20*	*0.00*	*-0.20*	*-0.40*	*-0.60*	*-0.80*	*-1.00*

Fractional predicted area		Sensitivity = true positive rate														
		0.90	**0.85**	**0.80**	**0.75**	**0.70**	**0.65**	**0.60**	**0.55**	**0.50**	**0.45**	**0.40**	**0.35**	**0.30**	**0.25**	**0.20**
0.60	Accuracy							1.00	0.90	0.80	0.70	0.60	0.50	0.40	0.30	0.20
p_e = 0.52	*κ*							*1.00*	*0.79*	*0.58*	*0.38*	*0.17*	*-0.04*	*-0.25*	*-0.46*	*-0.67*
0.70	Accuracy					1.00	0.90	0.80	0.70	0.60	0.50	0.40				
p_e = 0.58	*κ*					*1.00*	*0.76*	*0.52*	*0.29*	*0.05*	*-0.19*	*-0.43*				
0.80	Accuracy			1.00	0.90	0.80	0.70	0.60								
p_e = 0.68	*κ*			*1.00*	*0.69*	*0.38*	*0.06*	*-0.25*								
0.90	Accuracy	1.00	0.90	0.80												
p_e = 0.82	*κ*	*1.00*	*0.44*	*-0.11*												

* Missing data corresponds to invalid situations where the true positive rate would exceed FPA or the true negative rate would exceed (1 – FPA).

3.4.4 Correlations of suitability maps

The similarity of two continuously valued suitability rasters can be formally measured by calculating their correlation coefficient. The Pearson correlation coefficient measures the strength of the linear relationship between two variables and can take values between -1 and 1. Values close to 1 or -1 indicate a strong association, and values close to zero indicate a weak association. A correlation of 1 indicates a perfect positive correlation (i.e. the bivariate scatterplot is a straight-line with positive slope); a correlation of -1 indicates a perfect negative correlation; and a correlation of zero indicates that there is no relationship between the two variables (McMurray, Pace et al. 2004).

Pairwise correlations between multiple variables are often visualised as a correlation matrix chart. Figure 3-4 shows correlation matrix charts for pairwise correlations of five raster suitability maps produced using the *chart.Correlation* function from the *PerformanceAnalytics* package in R. The components of the correlation matrix chart are as follows:

- above the diagonal are the correlation coefficients for the map comparisons (with text size scaled according to the strength of the correlation), plus the significance level symbolised using stars:

p-values	Symbol
0.0 – 0.001	***
0.001 – 0.01	**
0.01 – 0.05	*
0.05 – 0.1	.
0.1 – 1	No symbol

- the diagonal shows the distribution of values in each map as a histogram;
- below the diagonal are bivariate scatterplots for each pair of maps, with fitted line.

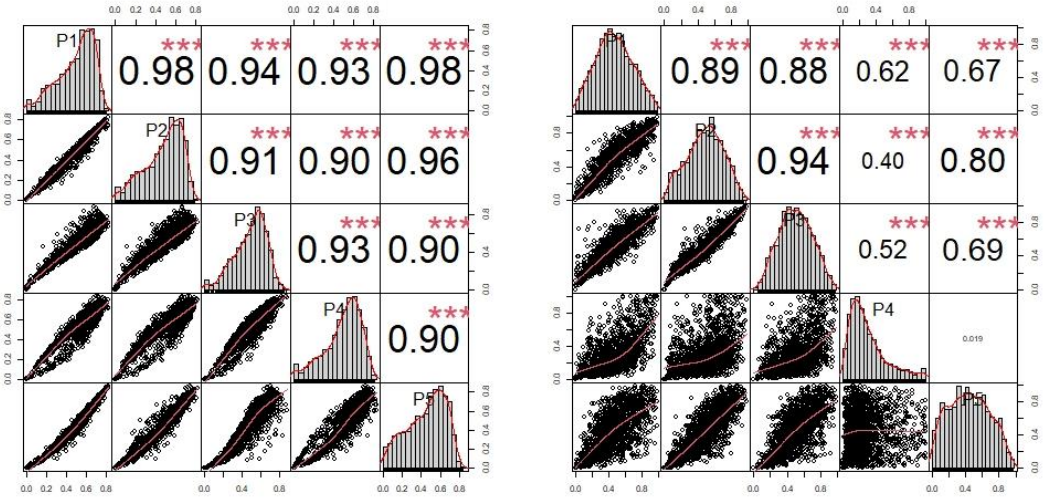

Figure 3-4 Example correlation matrix charts for five suitability rasters showing strong correlations (left) and weaker correlations (right)

When calculated for suitability maps, where the raster values are model predictions, a correlation coefficient close to 1 indicates very similar predictions. When comparing multiple alternative suitability maps for a species, strong correlations between the raster values provides evidence for the robustness of the predictions.

3.4.5 Information from high-resolution satellite imagery

The use of high-resolution Earth observation images is commonly used for validating research results in many geographical applications. For example, Google Earth and other services that provide high-resolution images can be used for the validation of land cover maps and to generate training and test data for the classification task.

Woodhouse (2021) in reviewing the use of the term “ground truth” quotes from an editorial in Nature Methods (2011):

> *“Researchers using satellite imaging to remotely observe features on the Earth enjoy the luxury of a simple solution for verifying the interpretation of their data with the truth on the ground or 'ground truth'. They or a surrogate can go observe it firsthand.”*

Examples of the visual detail available from Earth observation data in the Burkina Faso project region and in South Australia are provided in the Figure 3-5 and Figure 3-6, respectively. Individual buildings and trees are distinguishable and the presence of cultivated land is clearly recognisable.

In the Burkina Faso images, farmsteads with their clusters of round granaries are observable in landscapes of dryland agriculture. Areas of rice cultivation, distinguishable by small fields with raised boundaries to trap water, are also easily recognisable. The patterns of agriculture in the South Australian landscape are very different from those in Burkina Faso but are similarly very easy to recognise from high-resolution Earth observation data.

Figure 3-5 High-resolution satellite images of farmland in the Burkina Faso project region (Worldview2 true colour image April 2010 supplied by Gryphon Minerals Ltd)

Although not directly usable for quantitative measures of model performance, this imagery provides a rich source of qualitative data against which to compare predictions. The models developed for the Burkina Faso project region in this research use samples from a species' realised niche to estimate the fundamental niche for the species across the region. It was observed by Jiménez-Valverde (2012) that commission errors in relation to the realised niche are not necessarily errors in relation to the fundamental niche. So, predictions of agricultural suitability in areas where cultivation is not observed in the landscape do not necessarily correspond to commission errors by the models, but may instead identify potential areas of expansion for those crops.

Omission errors, on the other hand, can provide evidence of poor model performance in those areas where they occur. Although the Earth observation data cannot definitively identify

individual species growing at a location, it can reveal whether cultivation is occurring and the general style of agriculture that is being performed. In the Burkina Faso project region, dryland cropping and rice cultivation have distinctive landscape patterns, and orchards are readily distinguishable from bush due to the regularity of tree placement. In South Australia, broad-acre farming is unmistakeable in the imagery (often the tracks of the farm machinery are visible in the images), and vineyards and orchards are also easily recognisable.

Figure 3-6 High-resolution satellite images of farmland in South Australia (ESRI base maps, January 2022)

Comparing model predictions with the observable evidence on the ground is important for assessing the plausibility of the models. It would be expected that the predictions of the crop suitability models showed consistency with existing agricultural practices. The failure of suitability models for the major crops grown in an area to predict agricultural suitability at locations where this style of agricultural cultivation is clearly observable in the satellite imagery is strong evidence for omission errors occurring in the models and cannot be ignored.

3.5 Summary

This chapter has been primarily theoretical and has explained the theory underpinning maximum entropy modelling and various evaluation methods that can be used to assess generated models. The traditional approaches to land suitability modelling (FAO framework and GIS-based approaches) were first shown to be unsuitable for this project. However, experiments with species distribution modelling (using the BCCVL) demonstrated that correlative data-driven approaches could be readily applied to the research task. The presence/background algorithm Maxent seemed best suited for the modelling task, with species occurrences (but not absences) being derivable from the crop compensation maps and Maxent's output maps interpretable as models for the fundamental environmental niches for species and so aligning well with the research goals.

The maximum entropy modelling algorithm implemented in the Maxent software program was explained in detail, with descriptions given for how it works, calibration methods and the range of outputs available. The purpose of this detailed description is to inform readers unfamiliar with the algorithm so that they have appropriate context for interpreting the methodologies used and the

tests run that are described in the following two chapters. A similar approach was taken in explaining methods used for evaluating species distribution models.

There is no single “best” approach to evaluating model performance and the methods that can be applied in a particular project will depend on the type of evaluation data that is available for assessing model predictions. Continuously valued suitability predictions can be converted to binary predictions of species presence or absence by the use of thresholds to facilitate formal comparison with binary evaluation data (e.g. using a confusion matrix for point data or interrater reliability scores for binary verification maps). The AUC (area under the receiver operating characteristic curve) is a commonly used measure for assessing model performance that is calculated over all possible thresholds.

Most commonly used evaluation statistics in species distribution modelling are calculated using known occurrences and presumed absences that were not used in training the models; however, only occurrences can used in the case of Maxent as background points do not represent presumed absences. This affects the interpretation of the evaluation statistics and the ranges of possible values for them. In particular, the theoretical maximum AUC values achievable for Maxent’s predictive models depend on the species’ range, with higher AUC values achievable for species with narrow environmental niches than for species with less restricted habitats.

Clustered occurrence data presents problems with both the testing and training of species distribution model models. Where test data are clustered in small areas, the formal evaluation statistics will only measure the accuracy of the model in those areas and cannot be relied on as an indicator of model accuracy elsewhere. Spatially clustered training data can result in models that are distorted by spatial sampling bias. This bias is observable when very different patterns of prediction occur for models that were trained on spatially different subsets of the presence points using the same predictors. Correlating the predictions from such models is useful for detecting spatial bias.

Visual inspection of model outputs offers an intuitive means of assessing the credibility of models. High-resolution Earth observation imagery provides a rich source of qualitative data against which to compare model predictions and can be used to detect obvious omission errors that provide evidence of poor model performance in the areas where they occur.

In the modelling that is described in the following two chapters the configuration of the presence data presents a high risk of spatial sampling bias in generated models. Evaluation of these models will require the use of multiple complementary evaluation techniques to assess their feasibility and the accuracy and robustness of their predictions.

Chapter 4 Agricultural land suitability modelling: West Africa

4.1 Introduction

The goal of this research is to spatially predict suitable land for growing agricultural crops by estimating their fundamental ecological niches. The chosen methodology uses maximum entropy modelling (refer Chapter 3) to generate ecological niche models and resulting suitability maps using information on known occurrences for each species and environmental factors.

This chapter documents the development of suitability maps for crops grown locally in south-west Burkina Faso. The purpose of these maps is to assist farmers subject to compulsory land acquisition to identify other potential planting locations in the region for their crops. The reason for the compulsory acquisition of this land and resettlement of its farmers is the establishment of a new surface gold mine. The crop suitability models described in this chapter were developed using data acquired by the Australian company Gryphons Minerals Pty Ltd during the exploration phase for the mine (refer Chapter 2).

The community land use maps described in Section 2.2 are polygon shapefiles of surveyed fields where crops were recorded to be growing in 2014 when the Social Impact Assessment for the proposed mine was prepared. These data identify known occurrences for many locally grown crops in the form of presence locations (note that absence locations are not derivable from these maps). The presence points used by the modelling algorithm Maxent are sample points from the areal occurrence locations in these compensation maps.

Each ecological niche model created using Maxent is determined by the inputs to the algorithm. These inputs are: (1) presence points characterising known suitable locations for the target crop; (2) raster layers of predictor environmental variables; and (3) the algorithmic parameters specified for the run. Varying any of these inputs will change the suitability model that is generated, and some input combinations may produce unrealistic models. When the presence points used to train a model are not fully representative of the species' fundamental niche then spatial bias may be apparent in the generated model, with either significant areas of suitability omitted from the predictions or overprediction in areas unsuitable for the species.

The four agricultural communities from which the presence data in this study have been sampled grow similar crops. Many of the environmental conditions at these sites will be characteristic of land suitable for agriculture in general or these crops in particular; however, some may simply be incidental to geographic location and have no influence on suitability for cultivating crops (an example is longitude). When the presence data exhibits spatial sampling bias the contribution to the predictions of incidental environmental factors may be magnified and lead to the generation of unrealistic suitability models.

Which combinations of environmental predictors and Maxent parameters will generate realistic niche models for the crops is not known a priori and can only be determined by heuristic methods. The methodology used in this research generates multiple crop suitability models using different sets of environmental layers and model parameters, then compares the performance of the various models to identify those factors contributing most effectively to producing useful crop suitability models.

The configuration of the presence data for this research makes validation of results and assessment of model performance particularly challenging. All presence points occur within the four small areas where compulsory land acquisition will occur (refer Figure 1-3), so accuracy measures based on presence points alone (e.g. AUC) can only measure accuracy at these sites and not across the whole project region. The presence points exhibit geographical sampling bias and this can result in spatially biased ecological niche models. These issues of addressing spatial sampling bias and assessing model performance are discussed in more detail below, prior to the formal explanation of the methodology used in the research.

4.1.1 Spatial sampling bias

The presence data used in this modelling are an example of incidental occurrence data collected for a different purpose. They are a by-product of the formal application by Gryphon Minerals Ltd for a mining licence in the region and were gathered for the Social Impact Assessment. As such, the presence points do not represent a broad and representative sample across the target region but are spatially clumped in four small areas. Overrepresentation of some regions in the presence data can cause spatial bias and lead to environmental bias in resulting models, and so must be addressed (Hijmans 2012, Kramer-Schadt, Niedballa et al. 2013, Boria, Olson et al. 2014).

Model performance is commonly assessed using cross-validation data. Common cross-validation data splits of the presence data, such as random test percentage and those used in the replication methods built into Maxent, for this project, will result in test points taken from all four presence sites and in close proximity to training points. Such an approach cannot reveal the effects of sampling bias, and the spatial autocorrelation between training and test data can artificially inflate measures of model accuracy and perceived quality of predictions (Hijmans 2012).

There are a number of documented techniques for correcting sampling bias in Maxent species distribution models. Kramer-Schadt, Niedballa et al. (2013) compared two common strategies to cope with uneven sampling effort: *spatial filtering* addresses sampling bias by reducing the number of presence points in oversampled regions; and *background manipulation* uses a bias file that allows the user to choose background data with the same bias as the presence data. Both methods can produce substantial improvement in the quality of model predictions. However, spatial filtering increases the risk that the number of records will become too few to build statistically sound models, and background manipulation uses arbitrary values for the bias file that imply prior knowledge about the species' distribution that often is not available. Hijmans (2012) described *pairwise distance sampling* to remove spatial sorting bias but observed that this approach is not relevant when testing sites are very far away from training data. These techniques are not applicable to the configuration of presence data available to this project.

An alternative approach, identified by Merow, Smith et al. (2013), is to build a model using potentially biased samples and evaluate it against a spatially independent dataset. Accurate predictions from the model imply negligible sampling bias. As a general recommendation, these authors advise evaluating models based on their predictive accuracy on spatially independent cross-validation data using fit metrics based on sensitivity and avoiding thresholding whenever possible. However, they observe that it is usually challenging to obtain test data that is statistically independent of training data.

The configuration of the presence data for this project lends itself to this alternative approach. Partitioning of the samples to use three sites for training allows the fourth (spatially independent) site for cross-validation testing. Implementing this form of cross-validation allows four different combinations of spatially independent training and test data to be created from the same

presence data, and allows parallel models to be built and compared. This is the approach that has been taken in this project.

4.1.2 Assessing model performance

The exploratory ecological niche modelling performed using the BCCVL (refer Section 3.2) highlighted problems with the interpretation and comparison of suitability maps and demonstrated the capacity of the algorithms to generate highly unrealistic models at times. Maxent is a machine-learning algorithm that creates models based on the data presented to it in training; different sets of training data may result in very different suitability maps; or models may overfit to produce excellent predictions at the known locations used in training but very poor predictions elsewhere. These issues pose the questions: "what constitutes a good model?" and "how can we recognise and justify good results?" and so present the challenge of how to effectively validate results.

The purpose of ecological niche models is usually prediction rather that explanation or hypothesis testing, so estimates of predictive power are often more relevant that significance (Hijmans 2012). The predictive performance of a model is commonly evaluated through cross-validation whereby the model makes predictions for new data that were not used to estimate it. Typically, the available presence data are partitioned into sets for training and testing and the model's predictive performance is calculated from the accuracy of the test data predictions.

The available validation data for the models are the crop presence points at the four mine sites (used to calculate AUC measures) and high-resolution satellite imagery of the entire region that reveals existing land use. The detail in the imagery allows dryland agriculture to be distinguished from bas-fonds rice fields or uncleared land (refer Figure 2-12 and Figure 2-13 on page 35), but cannot distinguish particular dryland crops or definitively distinguish rangeland unsuitable for cultivation from fallow fields with bush regrowth.

In this chapter, multiple models are developed for each crop to explore the factors that contribute to good models or that may lead to production of poor models. These models differ in the combinations of predictors and partitions of presence data used to train and test the models. In particular, suites of similar models using the same predictors and Maxent settings but trained using different subsets of the presence points are compared to assess the model robustness and sensitivity to variations in the selection of training data.

In comparing the performance of different models and suites of models several formal and informal criteria are applied:

- **AUC** – the area under the curve (AUC) for receiver operation characteristic (ROC) curve is a formal accuracy measure that is valid for comparing alternative suitability models for the same crop. However, there are cautions against relying on AUC alone (refer the discussion of ROC curve in Section 3.4.2) and its accuracy is doubtful when there are very small numbers of test points from which to calculate it.
- **Suitability map correlations** – the similarity of map pairs is formally measured by computing their correlation coefficient. Well-correlated suitability maps for the same crop from similar models trained using the same predictors but different sets of presence points is evidence for model robustness.
- **Qualitative assessment using satellite imagery** – visual comparisons of suitability maps to the satellite imagery provides an informal method for assessing accuracy and may invalidate a model that has achieved good scores using the formal measures. For

example, a model that predicts very poor suitability for dry land crops or rice over large areas where that type of agriculture is clearly observed must be judged as a poor model.

In characterising good models, we would expect them to be realistic, with suitability ratings that did not contradict observable patterns of agricultural activity. Evidence for their accuracy and robustness would be high AUC scores and highly correlated suitability maps across the suite of similar models that varied only in the sets of presence points used in their training.

4.2 Methodology

Figure 4-1 provides an overview of methodology employed in this chapter to produce crop suitability models using the data sources described in Chapter 2. The methodology used takes advantage of the geographical separation of the four presence data sites to address the inherent problem of sampling bias in the presence data.

In order to detect the presence of sampling bias in relation to the environmental layers used as predictors, each test of environmental predictors or Maxent parameters is replicated five times to produce five parallel models for each crop. Each of the five models is trained and tested using different partitions of the presence data, four of which have test data that is spatially independent of the training data. These partitions of the presence data are referred to in this thesis as *model scenarios*. The accuracy of the test predictions and resulting suitability maps from the five model scenarios for a particular set of predictors are compared to assess model accuracy and robustness and to identify sensitivity to sampling bias.

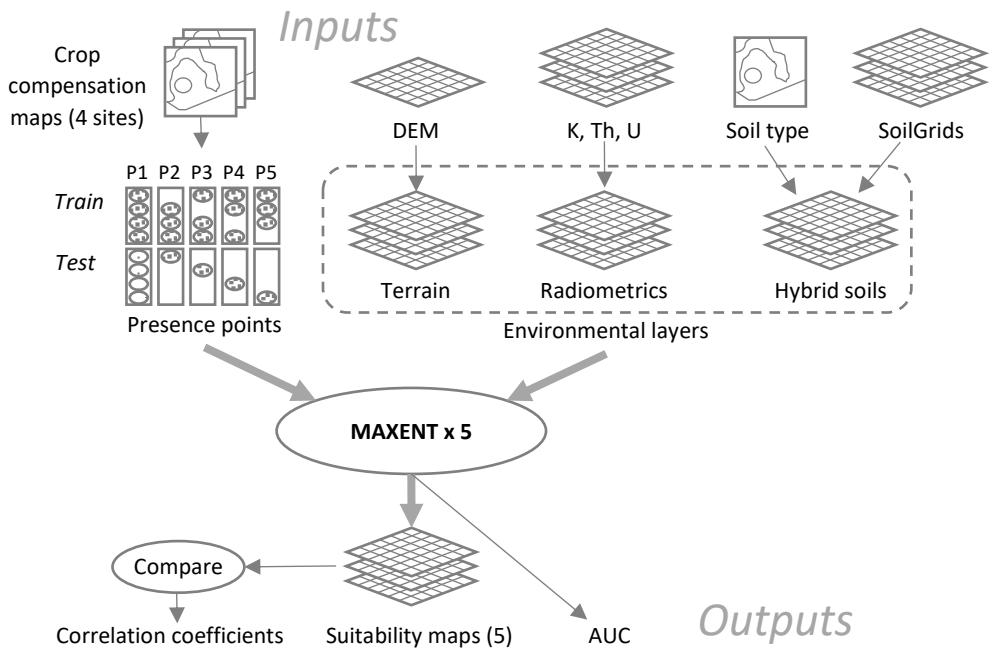

Figure 4-1 Data flow diagram showing inputs and outputs used to generate parallel crop suitability models in five model scenarios

4.2.1 Model scenarios

Table 4-1 documents the five model scenarios used in this research. The first model scenario uses a random presence data split of 75% for training and 25% for testing and so contains training and test presence points from each of the four mine sites. The four other model scenarios are

designed using a leave-one-out cross-validation partition of the mine sites and have spatially independent training and test sets of presence points. The model scenario names explicitly identify the set of presence points used to *test* the models (and, by inference, the set of presence points used to train the models).

Table 4-1 Model scenario definitions

Scenario	Name	Test data for models	Training data for models
P1	Random 25%	25% of presence points – randomly selected	Remaining 75% of presence points
P2	Fourkoura	Presence points at Fourkoura	Presence points at the other 3 sites
P3	Nogbele	Presence points at Nogbele	Presence points at the other 3 sites
P4	Samavogo	Presence points at Samavogo	Presence points at the other 3 sites
P5	Stinger	Presence points at Stinger	Presence points at the other 3 sites

4.2.2 Presence data

The presence data for the crop models are derived from the crop compensation maps that were discussed in Section 2.2 and displayed in Figure 2-1 and Figure 2-2. The centroids of all fields recorded as growing a particular crop were used as presence data for that crop. The number of presence points created for each of the twelve crops is shown in Table 4-2, as well as the number of training and test points used for each crop in each model scenario.

Table 4-2 Number of training and test presence points used for each crop in each model scenario

Model scenario	Type	Beans	Cashew	Cotton	Cowpea	Earth pea	Maize	Millet	Peanut	Potato/yam	Rice	Sesame	Sorghum	Total
P1 Random 25%	Training	48	24	168	31	29	260	53	99	30	506	155	32	1434
Random 25%	Test	16	8	56	10	10	87	18	33	10	169	52	11	478
P2 Fourkoura	Training	47	32	199	36	32	296	64	109	21	583	187	39	1645
Fourkoura	Test	17		25	5	7	50	7	23	19	91	20	3	267
P3 Nogbele	Training	30	14	67	24	31	152	30	71	26	345	116	37	943
Nogbele	Test	34	18	157	17	8	194	41	61	14	329	91	5	969
P4 Samavogo	Training	63	20	211	26	21	289	64	102	38	475	162	15	1486
Samavogo	Test	1	12	13	15	18	57	7	30	2	199	45	27	426
P5 Stinger	Training	52	30	195	37	33	301	55	114	35	619	156	35	1662
Stinger	Test	12	2	29	4	6	45	16	18	5	55	51	7	250
	Total	**64**	**32**	**224**	**41**	**39**	**346**	**71**	**132**	**40**	**674**	**207**	**42**	**1912**

4.2.3 Environmental layers

Chapter 2 described the development of the environmental layers for the modelling. Multiple terrain, soil and radiometric layers were produced, although no meaningful climate layers could be developed because there is only slight variation in climate over the region. All environmental layers were created as ASCII raster grids with cell size of 30metres by 30metres.

Terrain layers

The terrain layers include the DEM and five further layers derived from it: slope, wetness index, solar radiation, index of valley bottom flatness, and index of ridgetop flatness. These layers are described in Section 2.4 and shown in Figure 2-7.

Soil layers

The development of effective soils layers for use in suitability modelling proved the most challenging aspect of this research project. Two sources of soil data were available to the project. The supplied Gryphon Minerals map of soil types had very fine spatial detail. However, it was simply a categorical polygon map with 17 soil categories. The raster soil maps from SoilGrids quantified multiple soil qualities, but had coarse spatial resolution and their accuracy was not verified for the project region. Early experimentation using the BCCVL (described in Section 3.2) showed that neither were useful as environmental predictors for suitability models. This led to the development of sets of hybrid soil layers that combined these two data sources in various ways (documented in Section 2.5.3).

The hybrid soil layers had both the fine spatial detail from the Gryphon Minerals soil map and the multi-dimensionality of SoilGrids. Four sets of hybrid soil layers were developed as follows:

- Mean SoilGrids values for soil types
- Mean SoilGrids values for soil types plus Gaussian noise
- Mean SoilGrids values for soil types (50%) plus SoilGrids value (50%)
- Mean SoilGrids values for soil types (75%) plus SoilGrids value (25%)

The twenty soil property layers included some redundancy (e.g. two layers for pH, three layers for carbon content and five layers for water holding capacity). Nine soil properties were selected from each set of hybrid soil layers for use in the modelling, as follows:

- BLDFIE – Bulk density (fine earth) in kg per cubic metre
- OCDENS – Soil organic carbon density in kg per cubic metre
- CECSOL – Cation exchange capacity of soil in cmolc per kg
- PHIHOX – Soil pH x 10 in water
- CLYPPT – Clay content (0-2 micrometre) mass fraction in %
- CRFVOL – Coarse fragments volumetric in %
- SNDPPT – Sand content (50-2000 micrometre) mass fraction in %
- SLTPPT – Silt content (2-50 micrometre) mass fraction in %
- WWP – Available soil water capacity (volumetric fraction) until wilting point

Radiometric layers

Four radiometric layers, for radiometric dose rate and ground surface concentrations of potassium, thorium and uranium, were created for use in the modelling. These layers are described in Section 2.6 and shown in Figure 2-15.

Correlated layers and principal component analysis

Principal component analysis (PCA) is often used in predictive modelling to reduce the dimensionality of a problem by converting large sets of potentially correlated predictors into smaller sets of uncorrelated predictors. The correlation coefficients for pairwise comparisons of the environmental layers described above are presented in Appendix E and demonstrate highly correlated soil and radiometric layers.

Although the Maxent algorithm has no requirement that predictor layers be independent, some authors still discuss the detection and removal of correlated layers in the context of Maxent. For example, The Banta Lab[20] for integrative evolutionary and conservation biology offers a tutorial on its website for this process in preparation for running Maxent. Further, Dan Warren's ENMTools software for quantifying the similarity of environmental niche models generated by Maxent (Warren, Glor et al. 2010) offers a correlation function that is described in the user manual as "useful for eliminating spatially correlated variables prior to the modelling process, or as another measure of similarity between models." The use of principal component analysis (PCA) to generate sets of uncorrelated environmental layers has also been used for Maxent, e.g. as described in Cruz-Cárdenas, López-Mata et al. (2014).

The use of PCA was explored in this research to evaluate whether this technique would improve the measured accuracy and robustness of the models. Standardised PCA was applied to sets of environmental layers to derive sets of principal component predictors for use in modelling, as follows:

- **PCA of all environmental layers** – seven PCA layers derived from the complete set of six terrain layers, nine soil layers (hybrid soil layers with added Gaussian noise) and four radiometric layers.
- **PCA of 17 environmental layers** – seven PCA layers derived from 17 environmental layers, excluding the DEM and MRRTF.
- **PCA of soil layers** – four PCA layers derived from the nine soil layers; repeated for three sets of hybrid soil layers.

4.2.4 Software implementation

The Maxent software was run on an Ubuntu Linux virtual machine implemented using Oracle Virtual Box on a Windows10 workstation. Perl scripting was used to automate the generation of suites of Maxent crop suitability models based on specified sets of environmental layers. Each separate run created suitability models for all 12 crops in each of the five model scenarios, producing potentially 60 individual crop suitability models for that set of predictors. Maxent failed to generate a suitability model for Cashew in scenario P2 due to the absence of test data and so only 59 models were generated during each run.

Presence data for model scenario P1 consisted of all presence points. The presence points for each crop were randomly partitioned by Maxent at run-time into 75% for training and 25% for testing. Note that the seed for Maxent's random number generator was fixed so that the same partition of presence and background points was created in all runs. The training and test presence datasets for the four other model scenarios were explicitly created, as per Table 4-1 and Table 4-2.

Each environmental layer was assigned an alphanumeric code, as shown in Table 4-3. The user interface for the automated processing presented the list of codes for selection of environmental layers, and the codes were used in labelling the output directories to explicitly identify the model predictors used for each run.

[20] See https://sites.google.com/site/thebantalab/home

Table 4-3 Environmental layers used in modelling: codes and descriptions

Code	Name	Type	Layer description
0	DEM	Terrain	Digital elevation (meters)
1	Slope	Terrain	Slope (degrees)
2	Wetness	Terrain	Wetness index
3	Solar	Terrain	Solar radiation (J/cm^2)
A	BLDFIE	Soil	Bulk density (fine earth) in kg / cubic metre
B	OCDENS	Soil	Soil organic carbon density in kg / cubic metre
C	CECSOL	Soil	Cation exchange capacity of soil in cmolc / kg
D	PHIHOX	Soil	Soil pH x 10 in water
E	CLYPPT	Soil	Clay content (0-2 micrometre) mass fraction in %
F	CRFVOL	Soil	Coarse fragments volumetric in %
G	SNDPPT	Soil	Sand content (50-2000 micrometre) mass fraction in %
H	SLTPPT	Soil	Silt content (2-50 micrometre) mass fraction in %
I	WWP	Soil	Available soil water capacity (volumetric fraction) until wilting point
M	MRVBF	Terrain	M = Terrain - valley bottom flatness index
N	MRRTF	Terrain	N = Terrain - ridge top flatness index
P	PC1	PCA	Principal component 1
Q	PC2	PCA	Principal component 2
R	PC3	PCA	Principal component 3
S	PC4	PCA	Principal component 4
T	PC5	PCA	Principal component 5
U	PC6	PCA	Principal component 6
V	PC7	PCA	Principal component 7
W	Dose rate	Radiometric	Dose rate (nGy/hr)
X	Potassium	Radiometric	Potassium ground concentration (%)
Y	Thorium	Radiometric	Thorium ground concentration (ppm)
Z	Uranium	Radiometric	Uranium ground concentration (ppm)

Run-time parameters

Run-time parameters are used by Maxent to specify the location of input files, tune the modelling algorithm and select desired outputs. The run-time parameters to be used were encoded in the Perl script and edited when necessary to change parameters for particular runs.

The following run-time parameters were used for all, or some, runs:

- Output format of Cloglog was selected for all runs to produce an estimate between 0 and 1 of probability of presence (refer Section 3.3.5).
- Number of background points = 10,000.
- Random seed for creation of background points was set to FALSE for all runs, and writing of background predictions to a .csv requested for one run.
- Response curves produced for all runs.
- Jackknife results produced for some runs.
- Features classes were specified for some runs.
- The beta multiplier parameter was specified for some runs.

Replication

Maxent offers three methods of model replication: *K*-fold cross-validation, subsampling and bootstrap (refer discussion on data splitting in Section 3.3.4). Experimentation with each of these methods was performed interactively using the Java software package. However, as these methods are unable to address the problem of sampling bias inherent in the input data, these results are not reported in the body of this thesis.

Summary results from the replication runs are included in Appendix F for comparison with the results of other runs.

Post-processing

Post processing of the Maxent outputs was performed using R programs for the following tasks:

- Consolidate and plot suitability maps from each run (59 models)
- Consolidate and plot AUC results from each run (59 models)
- Consolidate and plot omissions from each run (59 models)
- Consolidate and plot environmental contributions from each run (59 models)
- Calculate correlation coefficients for pairwise comparisons of suitability maps for each crop in each run (10 map comparisons for most crops, 6 for Cashew; total of 116 per run)
- Summarise AUC and correlation results for all runs to allow comparisons between the different tests of predictors and parameters used by each run
- Smooth suitability maps and recalculate AUC and omissions for tests that used jittered hybrid soil layers.

4.2.5 Outputs

In total, 47 separate tests of predictors and parameters were run using the methodology described above. These runs resulted in the generation of 2,773 individual crop suitability models. Appendix F lists the 47 runs, identifying the environmental predictors and Maxent parameters that were used for each run and providing some summary test AUC results.

In addition to the standard outputs generated by Maxent for each model, summary plots were produced from each run to allow visual comparison of the results from all 59 models at once, as follows:

- Facet grid plots indexed by crop and model scenario for the following outputs:
    - Suitability maps
    - ROC curves
    - Omission and predicted area plots
    - Environmental layer contributions (histograms)
- Correlation chart matrices of suitability map comparisons for each of the 12 crops to document the similarity of the crop suitability maps from the different model scenarios.

These summary plots were copied to an Excel spreadsheet with 47 data rows to facilitate comparisons of outputs from different runs and inform qualitative assessments of relative model performance.

4.3 Results and discussion

The results of 47 runs, each representing a separate test of environmental predictors or Maxent parameters, are reported and discussed in this chapter. Figure 4-2 and Figure 4-3 show the distributions of the test AUC scores for all models and the correlation coefficients for pairwise comparisons of suitability maps for the same crop in all model scenarios from all 47 runs.

Test AUC scores

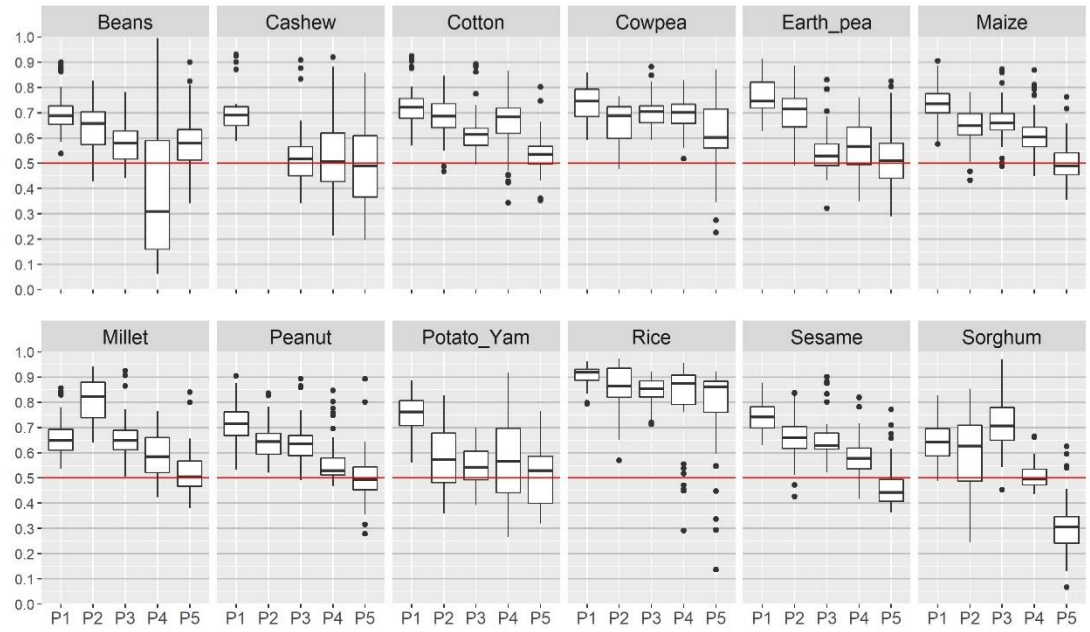

Figure 4-2 Distribution of test AUC scores for models from all runs, by crop and model scenario (red line shows AUC = 0.5, random prediction)

Correlations between suitability maps

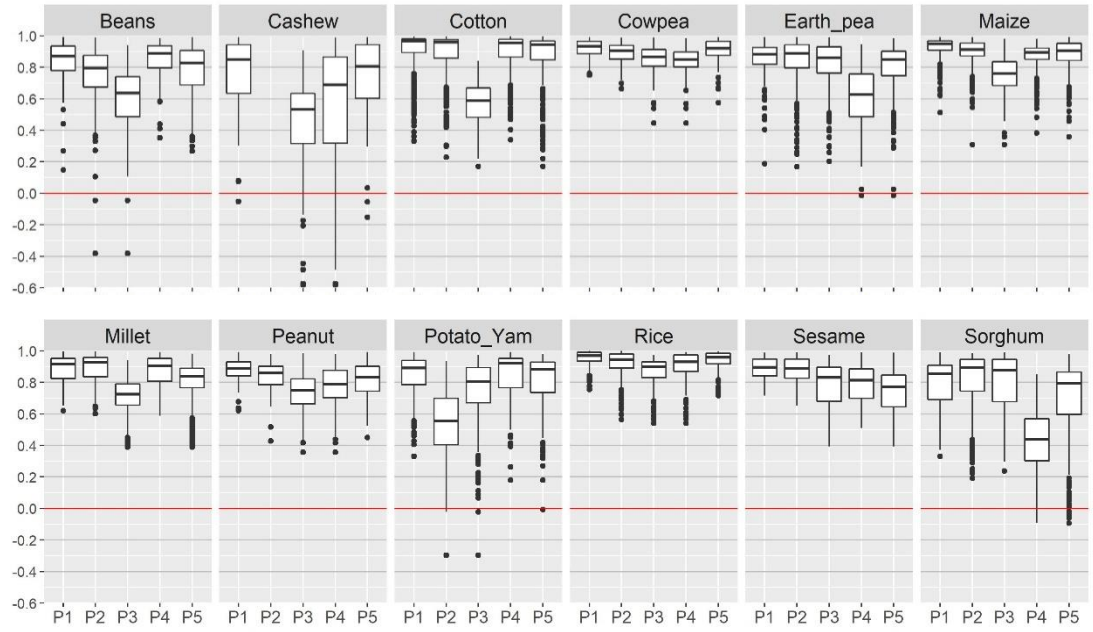

Figure 4-3 Distribution of correlation coefficients for pairwise comparisons of parallel suitability maps, by crop and model scenario (red line shows correlation coefficient = 0)

The 47 runs generated 235 (= 47 x 5) models for each crop (188 models for Cashew), resulting in 2,773 crop suitability models in total. Each box in the Figure 4-2 boxplots summarises the 47 test AUC scores for one crop in one model scenario.

Each run produced five parallel suitability maps per crop (four for Cashew). In order to test the similarity of predictions from the five model scenarios, each suitability map was compared with each of the other four maps for the same crop from the same run (three others for Cashew) and the correlation coefficients calculated. Figure 4-3 shows the correlation coefficients from these comparisons for all runs. Each box in the boxplots summarises 188 (= 47 x 4) correlation coefficients for comparisons with the maps for one crop in one model scenario (141 = 47 x 3 coefficients for Cashew).

The crop suitability models were built using different sets of environmental layers and model parameters with the purpose of identifying those factors contributing to good or poor models. Multiple similar models trained using different partitions of the presence data were generated to reveal the potential for sampling bias.

The numbers of presence points varied greatly between crops (refer Table 4-2) and, when partitioned for training and testing according to the model scenarios, resulted in some small test sets. Seven of the crops had less than 100 presence points: Beans (64), Cashew (32), Cowpea (41), Earth pea (39), Millet (71), Potato/yam (40) and Sorghum (42); and 15 of their test datasets had less than 10 presence points from which to calculate the test AUC. The plots in Figure 4-2 and Figure 4-3 show large variation in AUC scores and correlation coefficients for many of these crops.

Only five crops had more than 100 presence points: Cotton (224), Maize (346), Peanut (132), Rice (674) and Sesame (207). These are the main crops grown by the farmers and the discussion of results will focus on these five crops.

4.3.1 AUC scores

Models with very low test AUC scores (very poor models)

The plots in Figure 4-2 reveal some very low test AUC scores for some crop models. Table 4-4 lists the number of models for each of the five main crops that achieved test AUC scores of less than 0.5 (i.e. performed worse than random prediction). These 149 crop models make up 12.7% of the total 1,175 models for the five crops in the five scenarios.

Table 4-4 Number and percentage of models with test AUC scores less than 0.5 for the five main crops, by model scenario

Model scenario	Cotton	Maize	Peanut	Rice	Sesame	*Total*	*%*
P1 Random 25%	0	0	0	0	0	***0***	*0.0%*
P2 Fourkoura	2	2	0	0	2	***6***	*2.6%*
P3 Nogbele	1	1	2	0	0	***4***	*1.7%*
P4 Samavogo	7	7	6	4	9	***33***	*14.0%*
P5 Stinger	13	27	27	4	35	***106***	*45.1%*
Total	***23***	***37***	***35***	***8***	***46***	***149***	*12.7%*
%	*9.8%*	*15.7%*	*14.9%*	*3.4%*	*19.6%*	*12.7%*	

All the models in scenario *P1 Random 25%* performed better than random prediction. Very low test AUC scores occurred only in the model scenarios that used spatially independent training and

test data. Inspection of the models with very low test AUC scores provides insight into some of the factors that can contribute to poor models:

- Only linear features were used in building most of the very poor models in scenarios *P2 Fourkoura* and *P3 Nogbele* (4 of 6 and 3 of 4 models, respectively).
- In scenario *P4 Samavogo*, 27 (82%) of the 33 poor models used environmental layers whose values are influenced by elevation (DEM, MRVBF or MRRTF). Samavogo has the highest altitude of the four sites and sampling bias is apparent when samples from this site are not used for training models built using these layers.
- In model scenario *P5 Stinger*, 45% of the 235 crop models performed worse than random prediction. This is almost certainly due to the effect of sampling bias as a consequence of Stinger's distinctive soil profile and lower altitude.

Stinger, with 50% cuirasse soils, has a very different soil profile from the other mining lease sites that each have between 2% and 5% curaisse soils (refer Table 2-6 in Section 2.5.1). In scenario P5, the crop models are trained using data from the other three sites where residual soils predominate, but the test AUC score for scenario P5 is calculated using presence data from Stinger alone where 50% of cultivation is on curaisse soils. The extreme sampling bias evident in scenario P5 is not evident in the other four scenarios where training data for models includes presence data from Stinger.

Fractional predicted area and test AUC values

This research aims to model the ecological niches for crops and so identify potential planting locations across the study area. The theoretical maximum AUC value for Maxent ecological niche models corresponds to $1 - \alpha/2$ where α is the fraction of the study area corresponding to the species' true distribution (refer Table 3-4 and discussion in Section 3.4.2). The fractional predicted areas for the crops in this study are not known, but analysis of the tables of soil type in Section 2.5.1 allows some estimates to be made.

At the four known sites (mining leases) approximately 60% by area of rice cultivation is on depositional soils (codes Da and Dc), and it is observed that rice is planted in preference to other crops on these soils. The remaining 40% of rice cultivation is on nearby residual soils and curaisse. With depositional soils (codes Da and Dc) making up 9% of the project area, we can extrapolate to estimate a fractional predicted area for rice of approximately 15%, and hence a theoretical maximum AUC for Rice suitability models of 0.93.

In contrast, over 97% (by area) of cotton, peanut and sesame and over 93% of maize grown at the four mining leases is on residual soils or curaisse. These crops thrive in similar environments and are typically planted next to each other in patchworks of mixed fields at the discretion of farmers (refer Figure 2-1 and Figure 2-2 on page 18). Residual soils and curaisse make up 85% of the project region. If 100% of these soils were suitable for cultivation, the theoretical maximum achievable AUC value for suitability models for these crops would be 0.58. If 90%, 80% or 70% of these soils were suitable, the corresponding maximum achievable AUC values would be 0.62, 0.66 and 0.70, respectively.

Models with high AUC scores (good models)

Table 4-5 tabulates the numbers of suitability models achieving high test AUC scores by crop and model scenario, where a high score for Rice models is defined to be greater than 0.9, and a high score for other crop models is defined to be greater than 0.7.

Overall, almost 30% of the 1,175 suitability models achieved high test AUC scores, including over two thirds of models for scenario *P1 Random 25%*. However, only 15 of the models from *P5 Stinger* (6% of 235) achieved high test AUC scores.

Table 4-5 Number and percentage of models with high test AUC scores for the five main crops, by model scenario: high test AUC scores are > 0.9 for Rice and > 0.7 for Cotton, Maize, Peanut, Sesame

Model scenario	Cotton	Maize	Peanut	Rice	Sesame	*Total*	*%*
P1 Random 25%	29	35	30	31	34	***159***	*67.7%*
P2 Fourkoura	18	10	8	17	13	***66***	*28.1%*
P3 Nogbele	5	11	6	7	7	***36***	*15.3%*
P4 Samavogo	20	8	5	21	6	***60***	*25.5%*
P5 Stinger	2	2	2	7	2	***15***	*6.4%*
Total	***74***	***66***	***51***	***83***	***62***	***336***	*28.6%*
%	*31.5%*	*28.1%*	*21.7%*	*35.3%*	*26.4%*	*28.6%*	

4.3.2 Variable contributions and permutation importance

The relative importance of the environmental variables to a Maxent model is measured using two metrics (Phillips 2017):

- **Permutation importance** measures the drop in training AUC for the final model when the values for an environmental layer are randomly permuted among the training points (both presence and background) and the suitability scores recalculated. A large decrease in training AUC indicates that the model depends heavily on that variable.
- **Variable contribution** is accumulated during model training by assigning the increase in gain at each iteration to the environmental variable that the feature being modified depends on.

Both metrics are normalised to percentages to allow comparison between environmental variables. However, Phillips (2017) advises that variable contributions should be interpreted with caution when predictor variables are correlated.

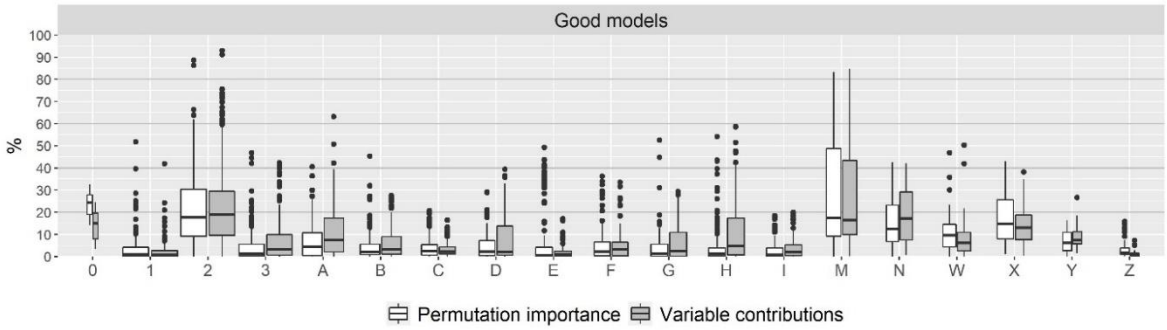

Figure 4-4 Distribution of permutation importance and variable contribution values for environmental layers in good suitability models (test AUC > 0.9 for Rice and > 0.7 otherwise) for the five main crops

The environmental layers used for training the models are listed in Table 4-3. Figure 4-4 uses boxplots to show the distribution of values for permutation importance and variable contribution from the 336 good models (refer Table 4-5) for 19 environmental predictor layers. The models

used different numbers and combinations of predictors and the width of each boxplot is indicative of the number of models that used that predictor. Note that, although some of the models also used PCA layers (PQRSTUV), these layers are not shown in the plot.

Reasonably similar patterns for permutation importance and variable contribution can be observed in Figure 4-4 for the terrain layers (0123MN) and radiometric layers (WXYZ). For the highly correlated soil layers (ABCDEFGHI) the permutation importance of each layer is typically very low, whereas the variable contributions exhibit much more variation.

Figure 4-5 compares the permutation importance of these environmental layers for good and very poor models. For Rice, the MRVBF (M) layer has very high permutation importance in both sets of models. In contrast to the good models, the importance of the DEM (0) and MRRTF (N) terrain layers was also high for the very poor Rice models. The only soil layer with non-negligible permutation importance for Rice was soil organic carbon density (B) in the good models.

For the other four crops, the permutation importance of the correlated soil layers (CDEFGHI, see Appendix E) is higher for the good models, whereas wetness (2) has higher importance in the very poor models.

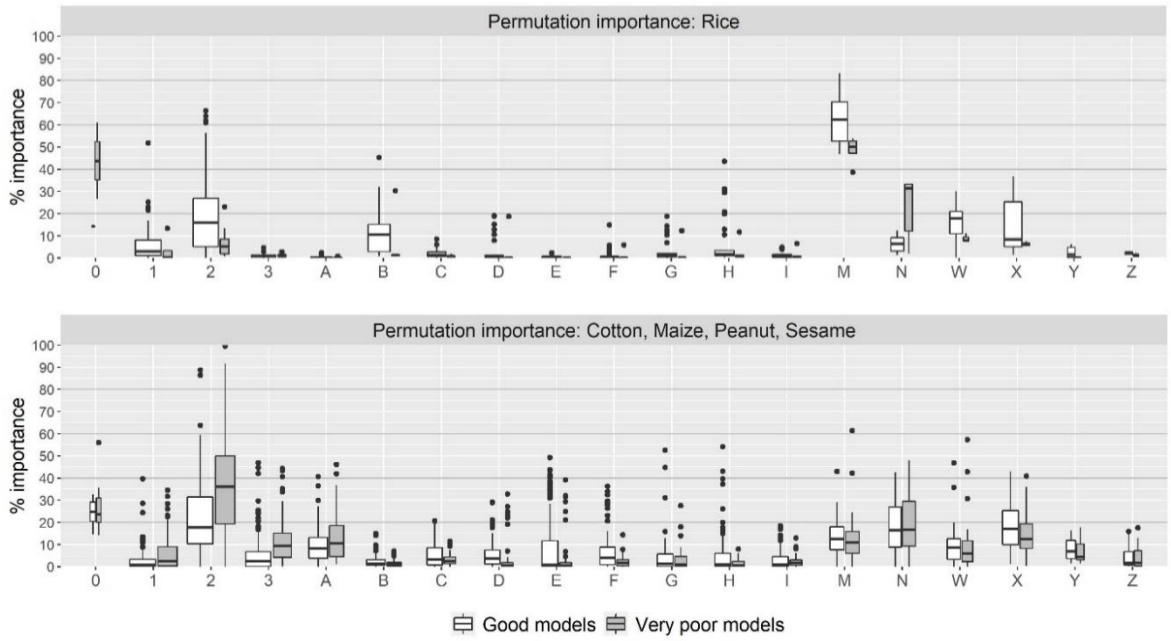

Figure 4-5 Distribution of permutation importance values for environmental layers in Rice suitability models (top) and Cotton, Maize, Peanut and Sesame suitability models (bottom): good models have test AUC > 0.9 for Rice and > 0.7 otherwise; very poor models have test AUC < 0.5

4.3.3 Model parameters

The most significant Maxent run-time parameters are the regularisation parameter (referred to as the beta multiplier) and the choice of feature classes to use in the modelling. The default beta multiplier value is one, and the default set of feature classes used by a model depends on the number of training samples (refer Table 3-1). To assess the effect of changing these parameters, sets of tests were performed using identical environmental layers and varying just one of these parameters. The results for each test are summarised over the five model scenarios for each crop, with the mean and minimum test AUC values reported for comparison between tests. Two summary statistics are used to compare overall accuracy (mean test AUC over all five crops) and

similarity of model results (mean of all pairwise map correlations from the five model scenarios for crops) between tests.

Regularisation

Table 4-6 reports the effect of using beta multiplier values of 1, 2 and 4 for models built from identical environmental layers using default feature classes. Small changes in model performance are observable, with both the mean and the minimum test AUC values for crops reducing as the beta multiplier value increases.

Table 4-6 Summary test AUC values for models built from uncorrelated environmental layers (123ABWX) using different beta multipliers, by crop

Beta	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
multi-plier	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	***Mean map correlations***
1	0.66	*0.51*	0.65	*0.49*	0.60	*0.43*	0.90	*0.88*	0.61	*0.45*	0.68	*0.84*
2	0.66	*0.50*	0.65	*0.48*	0.59	*0.42*	0.90	*0.88*	0.60	*0.43*	0.68	*0.83*
4	0.65	*0.50*	0.64	*0.48*	0.57	*0.40*	0.90	*0.88*	0.59	*0.42*	0.67	*0.81*

These results are consistent with the findings of Radosavljevic and Anderson (2014) whose systematic testing found peak model performance for the default beta multiplier value of one. The similarity of suitability maps also decreased as the beta multiplier value increased. No further experiments were performed in relation to regularisation and the default beta multiplier of one was used in all other tests.

Feature classes

Table 4-7 and Table 4-8 report the effect of increasing the complexity of the feature classes used in models built from identical sets of environmental layers. (Note: as there were fewer than 80 training samples for Cotton and Peanut in scenario *P3 Nogbele* product features were not used for these two models.)

Table 4-7 compares the effect of using different sets of feature classes for models built from just terrain layers (1=slope, 2=wetness, 3=solar radiation, M=MRVBF and N=MRRTF). The results show that the combination of linear, quadratic and product features (LQP) performed better than linear features (L) alone, giving higher mean and minimum test AUC values across the five model scenarios for all crops. Adding the hinge feature (H) degraded test AUC results for Cotton, Rice and Sesame, but improved results for Peanut. Significantly, the addition of the hinge feature resulted in one very poor Rice model that had test AUC value of just 0.29.

Table 4-7 Summary test AUC values for models built from terrain layers only (123MN) using different combinations of feature classes (L=linear, Q=quadratic, P=product, H=hinge), by crop

	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
Features used	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	***Mean map correlations***
L	0.55	*0.35*	0.52	*0.36*	0.49	*0.35*	0.77	*0.57*	0.60	*0.37*	0.59	*0.80*
LQP	0.62	*0.52*	0.58	*0.45*	0.54	*0.45*	0.80	*0.65*	0.63	*0.44*	0.63	*0.78*
LQPH	0.61	*0.42*	0.58	*0.46*	0.58	*0.49*	0.73	*0.29*	0.59	*0.42*	0.62	*0.80*

Table 4-8 compares the use of different sets of feature classes in models built from three terrain layers (123) and all nine soil layers. The results show that the use of linear feature alone produced the lowest overall accuracy (as measured by mean test AUC), but the highest overall similarity of resulting suitability maps (as measured by the mean of the correlation coefficients for pairwise comparisons of suitability maps for the same crop). The best overall model performance occurred with the use of linear, quadratic and product features (LQP), and the further addition of the hinge feature (H) degraded overall performance slightly. Inspection of the variable contributions for all the models in this set of tests showed very similar variable contributions for corresponding LQ and LQP models (i.e. same crop and model scenario). However, the contribution of the wetness environmental layer was far higher in the LQPH models.

Table 4-8 Summary test AUC values for models built from terrain and soil layers (123ABCDEFGHI) using different combinations of feature classes (L=linear, Q=quadratic, P=product, H=hinge), by crop

	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
Features used	**Mean AUC**	***Min AUC***	**Mean AUC**	***Min AUC***	**Mean AUC**	***Min AUC***	**Mean AUC**	***Min AUC***	**Mean AUC**	***Min AUC***	**Mean AUC**	***Mean map correlations***
L	0.59	*0.50*	0.57	*0.39*	0.62	*0.54*	0.86	*0.77*	0.59	*0.41*	0.64	*0.89*
LQ	0.64	*0.55*	0.61	*0.43*	0.62	*0.51*	0.88	*0.82*	0.58	*0.39*	0.67	*0.87*
LQP	0.65	*0.57*	0.63	*0.47*	0.64	*0.57*	0.89	*0.83*	0.62	*0.40*	0.69	*0.88*
LQPH	0.65	*0.57*	0.61	*0.47*	0.62	*0.48*	0.89	*0.84*	0.59	*0.40*	0.67	*0.86*

4.3.4 Environmental layers

The next two sections investigate how particular environmental layers influence the predictive performance of models. This section explores the effect of using terrain, radiometric and PCA layers in developing crop suitability models. The development of suitable soil layers was the most problematic element of this research and so discussion on the soil layers is treated separately in Section 4.3.5.

All the mine sites are agricultural communities that grow similar crops. Many of the environmental conditions at these sites are characteristic of land suitable for agriculture and some are incidental due to geographic location. Inspection of the distributions of values from environmental layers gives insight into the potential for spatial bias in some tests and model scenarios.

Terrain layers

Figure 4-6 shows the distributions of values from the six terrain layers at the four mine sites. Violin plots show the shapes of the distributions, with black bars identifying mean values. The red line in each layer plot represents the mean value across the project region for that environmental predictor.

It is observed that the average slope at all four sites is slightly lower than the region average, as might be expected for agricultural communities, and the amount of solar radiation received at all sites is similar to the region mean. Site variations are noticeable in the plots for the 0-DEM and 2-Wetness. Overall, site *P5 Stinger* is lower in altitude and wetter than the other sites, whereas site *P4 Samavogo* is higher in altitude and has lower wetness values. The distribution of MRVBF and MRRTF values varies greatly between the sites, with *P2 Fourkoura* having no values for MRVBF greater than five and *P4 Samavogo* being the only site with MRRTF values greater than six.

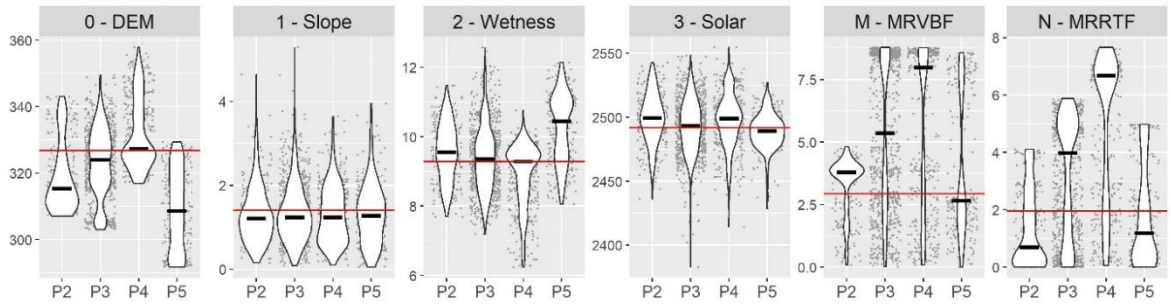

Figure 4-6 Violin plots of terrain layer values at the four mine sites (P2 Fourkoura, P3 Nogbele, P4 Samavogo and P5 Stinger) with mean value (black bar) compared to the mean for the project region (horizontal red line)

Table 4-9 presents results for models trained using terrain layers only. All the very poor models (test AUC < 0.5) occurred for model scenarios *P5 Stinger* and *P4 Samavogo* where the terrain values for the test data were less similar to the values used in training.

Table 4-9 Summary test AUC values for models built from terrain layers only (0=DEM, 1=Slope, 2=Wetness, 3=Solar, M=MRVBF, N=MRRTF) using default feature classes and beta multiplier value

Layers	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
	Mean AUC	*Min AUC*	**Mean AUC**	*Min AUC*	**Mean AUC**	*Min AUC*	**Mean AUC**	*Min AUC*	**Mean AUC**	*Min AUC*	**Mean AUC**	***Mean map correlations***
0123	0.65	*0.56*	0.62	*0.49*	0.60	*0.47*	0.66	*0.14*	0.64	*0.42*	0.64	*0.87*
12	0.63	*0.54*	0.60	*0.47*	0.57	*0.48*	0.75	*0.65*	0.60	*0.38*	0.63	*0.87*
123	0.66	*0.53*	0.61	*0.46*	0.58	*0.45*	0.75	*0.66*	0.63	*0.40*	0.65	*0.88*
123M	0.65	*0.54*	0.60	*0.44*	0.55	*0.47*	0.83	*0.67*	0.61	*0.41*	0.65	*0.84*
123MN	0.61	*0.42*	0.58	*0.46*	0.58	*0.49*	0.73	*0.29*	0.59	*0.42*	0.62	*0.80*

The effect of particular terrain layers was most pronounced for the Rice suitability models. The use of M=MRVBF as a predictor substantially improved the performance of Rice models in terms of test AUC (but also slightly degraded the performance of models for the dryland crops). Using 0=DEM as a predictor resulted in a very poor Rice model with test AUC of 0.14 for scenario *P5 Stinger*, and the use of N=MRRTF resulted in a very poor Rice model with test AUC of 0.29 for scenario *P4 Samavogo*.

Figure 4-7 shows correlation results for pairwise comparisons of Rice suitability maps produced by the five model scenarios in each run. The highest measures of map similarity occurred for models using layers 12 or 123 only, and the lowest similarity scores occurred for the models using layer N=MRRTF.

Examination of the scatterplots for the 123M, 123MN and 0123 models reveals many divergences in model predictions between the suitability maps from different model scenarios. This means that, for large parts of the project region, models developed using the same set of environmental predictors have produced contradictory results – one model predicting very poor suitability and another predicting high suitability for the same location. As the suitability maps compared in each of the scatterplots differ only in the samples used for training, these results demonstrate that these models exhibit substantial spatial bias derived from the training samples.

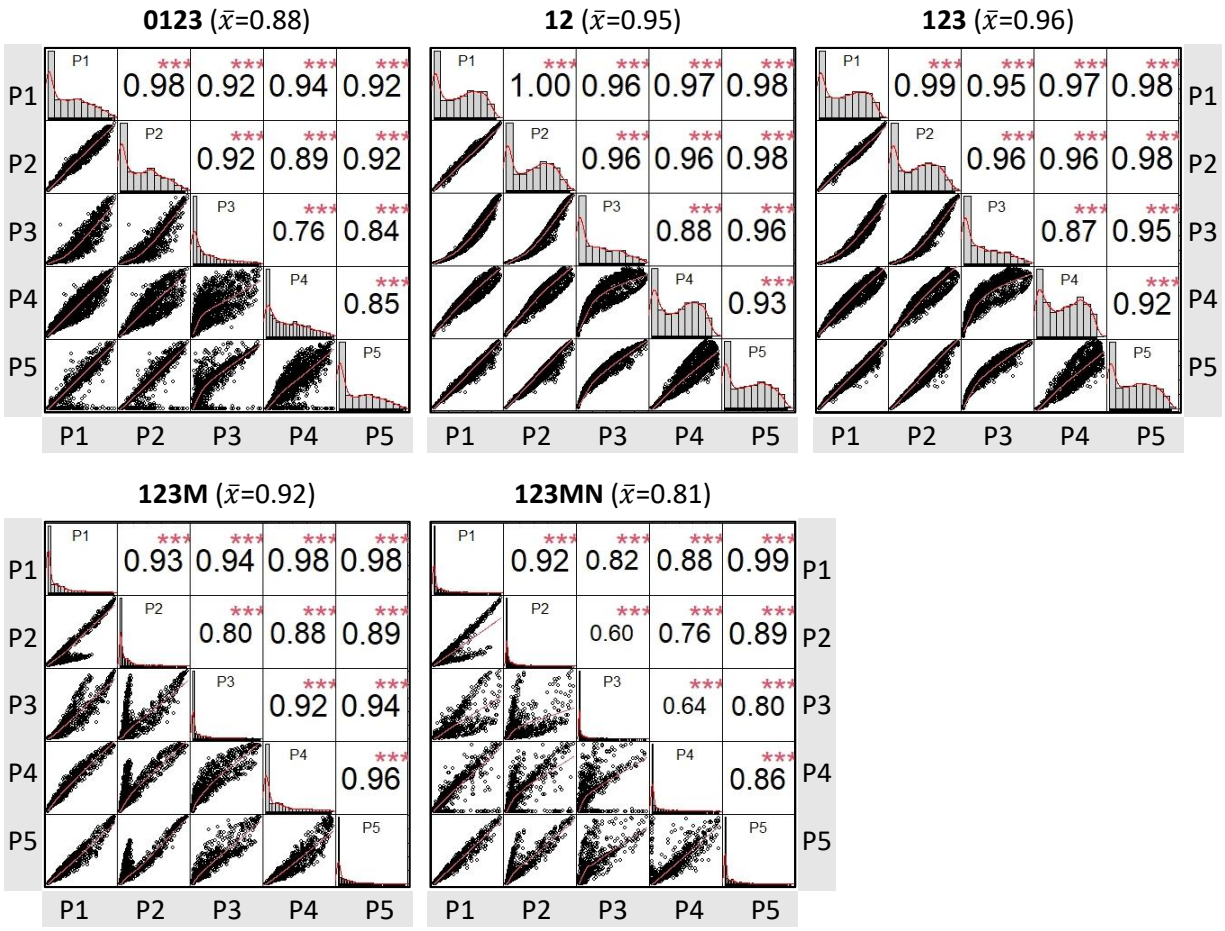

Figure 4-7 Rice suitability map correlations between the five model scenarios, from models built using combinations of terrain layers only: 0=DEM, 1=Slope, 2=Wetness, M=MRVBF, N=MRRTF

Radiometric layers

Figure 4-8 shows the distribution of radiometric layer values at the four mine sites. All sites had higher mean potassium values than the project region as a whole. As potassium is usually related to clay content in soils, this suggests that soils with higher clay content may be preferred for cropping.

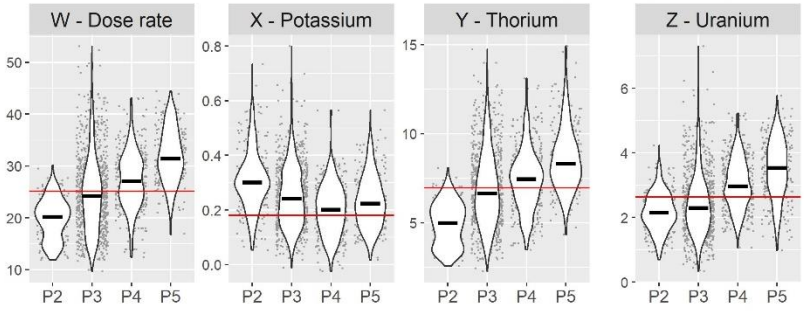

Figure 4-8 Violin plots of radiometric layer values at the four mine sites (P2 Fourkoura, P3 Nogbele, P4 Samavogo and P5 Stinger) with mean value (black bar) compared to the mean for the project region (horizontal red line)

Table 4-10 provides summary test AUC results for models built using radiometric layers. Run WXYZ used only radiometric layers as predictors. The addition of terrain layers 1=Slope, 2=Wetness and 3=Solar as predictors increased mean AUC for all crops. The further addition of two soil layers for run 123ABWX of uncorrelated predictors increased test AUC for Rice, but made little change to AUC results for the dryland crops.

Table 4-10 Summary test AUC values for models built using radiometric layers: W=Dose rate (nGy/hr), X=Potassium (%), Y=Thorium (ppm), Z=Uranium (ppm)

Layers	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	***Mean map correlations***
WXYZ	0.57	*0.43*	0.61	*0.54*	0.57	*0.53*	0.80	*0.69*	0.55	*0.43*	0.62	*0.85*
123WXYZ	0.67	*0.51*	0.66	*0.52*	0.60	*0.47*	0.85	*0.82*	0.62	*0.47*	0.68	*0.85*
123ABWX	0.66	*0.51*	0.65	*0.49*	0.60	*0.43*	0.90	*0.88*	0.61	*0.45*	0.68	*0.84*

Principal component layers

Standardised PCA was applied to sets of terrain, soil and radiometric environmental layers to derive sets of seven principal component predictors (PQRSTUV) for use in modelling. Two sets of PCA layers were generated from these layers, as follows:

- **PCA-19** – based on all 19 environmental layers (0123ABCDEFGHIMNWXYZ)
- **PCA-17** – based on 17 layers, excluding the DEM and MRRTF (123ABCDEFGHIMWXYZ)

Suitability models were built using these sets of principal component layers as predictors and the test AUC results from these runs are presented in Table 4-11. Comparing the PCA-19 and PCA-17 results shows that the exclusion of the DEM and MRRTF from the source layers for PCA resulted in a small improvement in predictive performance overall.

Table 4-11 Summary test AUC values for models built using seven standardised PCA layers (PQRSTUV) derived from 19 and 17 source environmental layers

PCA version	Cotton		Maize		Peanut		Rice		Sesame		5 main crops	
	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	*Min AUC*	Mean AUC	***Mean map correlations***
PCA-19	0.62	*0.53*	0.64	*0.54*	0.61	*0.50*	0.87	*0.83*	0.60	*0.50*	0.67	*0.84*
PCA-17	0.62	*0.51*	0.66	*0.57*	0.60	*0.51*	0.89	*0.83*	0.65	*0.59*	0.68	*0.85*

The PCA models had test AUC values of at least 0.5 for the five main crops in all model scenarios, indicating mitigation of the effects of spatial sampling bias. However, the use of PCA obscures the contributions of individual environmental layers to the final models and so provides little insight as to which environmental layers are more significant as predictors.

Application of PCA to the soils layers only was also tested, and the results of these tests are reported in the next section.

4.3.5 Soil layers

The development of four sets of hybrid soil layers is documented in Section 2.5.3, and the nine soil properties used for suitability modelling are listed in Section 4.2.3 and in Table 4-3. For clarity in the presentation of results, the sets of hybrid soil layers are named in this section as follows:

Means = mean SoilGrids values for soil types

Means + jitter = mean SoilGrids values for soil types with the addition of Gaussian noise

Means + SG (50:50) = mean SoilGrids values for soil types plus SoilGrids values in ratio 50:50

Means + SG (75:25) = mean SoilGrids values for soil types plus SoilGrids values in ratio 75:25

The effectiveness of each set of hybrid soil layers was tested by generating suitability models with these layers as environmental predictors. Principal component analysis of the nine hybrid layers in each set was also performed, and models built using the first two principal components as environmental predictors.

Hybrid soil layers

Figure 4-9 illustrates the differences between the sets of hybrid soil layers by showing corresponding pixel-level map detail for layer G=SNDPPT (sand content mass fraction in %) from the SoilGrids (top 30cm) raster and the hybrid soil layers.

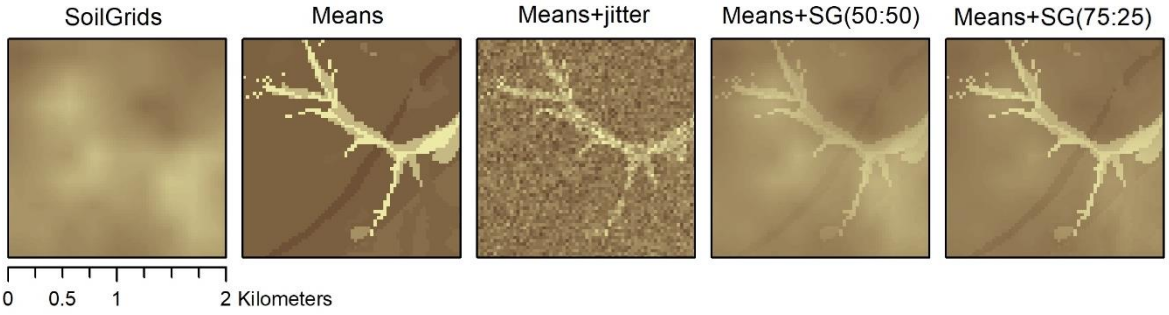

Figure 4-9 Corresponding map detail from the SoilGrids (top 30cm) and hybrid soil layer rasters for the sand content soil layer G=SNDPPT (darker colours indicate higher values)

Figure 4-10 illustrates the distribution of values in the SoilGrids (top 30cm) rasters and the four hybrid soil layers for the soil predictors A=BLDFIE (bulk density fine earth in kg per cubic metre) and B=OCDENS (soil organic carbon in kg per cubic metre). Violin plots show the shapes of the distributions for the background points and for the presence points at the four mine sites, using black bars to identify mean values.

The continuous values from the SoilGrids raster are resolved into seventeen discrete values for the soil type means (see *Means* plots). All seventeen values occur in the background distribution; however, fewer values occur at each mine site and some values do not occur at any of the mine sites. The distributions of values where Gaussian noise was added to the means (*Means + jitter*) shows greater similarity across the four mine sites when compared to the original SoilGrids values or to the hybrid soil layers where the SoilGrids values were added to give variation about the mean values (*Means + SG (50:50)* and *Means + SG (75:25)*).

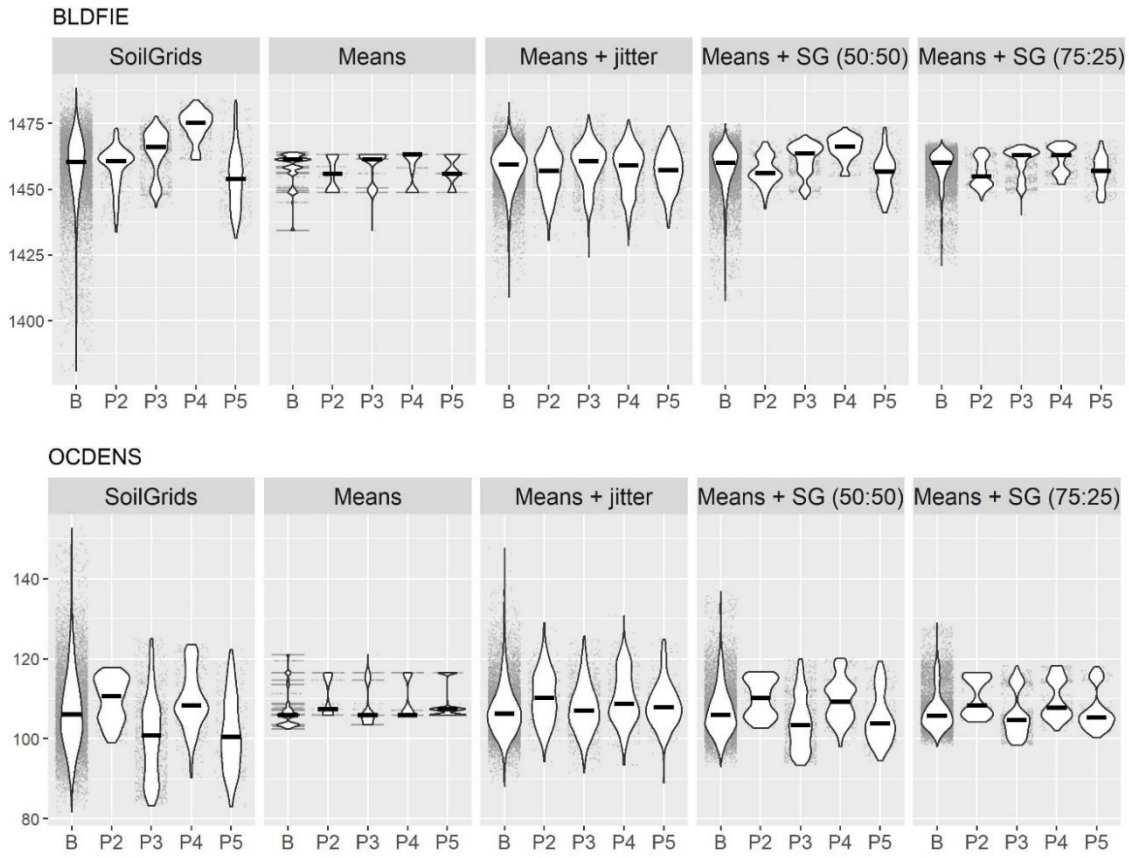

Figure 4-10 Violin plots of values with mean (black bar) from SoilGrids and the hybrid soil layers at background points (B) and at the four mine sites (P2=Fourkoura, P3=Nogbele, P4=Samavogo, P5=Stinger) for layers A=BLDFIE and B=OCDENS

Models were trained using each of the sets of hybrid soil layers plus three terrain layers (1=Slope, 2=Wetness, 3=Solar) and summary AUC results for these models are given in Table 4-12. All models used linear, quadratic and product feature classes (LQP) and the default beta multiplier of one.

Note that the suitability maps for the *Means + jitter* models all exhibited random noise, consequent of the noisy predictor layers used in the models. Two passes of a simple smoothing filter of radius 3 were applied to these suitability maps to remove this random noise (refer Figure 4-11) and AUC values were recalculated from the smoothed maps.

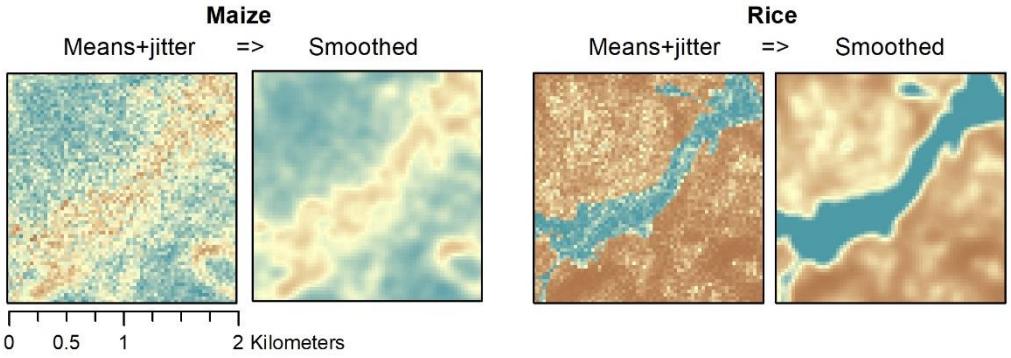

Figure 4-11 Corresponding suitability map detail showing the results from smoothing Maize and Rice models built using the Means + jitter hybrid soil layers

Table 4-12 Summary test AUC values for models built from terrain and soil layers (123ABCDEFGHI) for four sets of hybrid soil layers constructed using different methods

Soil layers = ABCDEFGHI	Mean test AUC (5 model scenarios)					5 main crops	
Method	**Cotton**	**Maize**	**Peanut**	**Rice**	**Sesame**	**Mean AUC**	***Mean Map Correlations***
Soil type means	0.71	0.70	0.66	0.91	0.66	0.73	*0.87*
Means + jitter	0.65	0.63	0.64	0.89	0.62	0.69	*0.88*
Means + jitter: smoothed maps	0.68	0.65	0.67	0.91	0.64	0.71	*0.89*
Means (50%) + SG (50%)	0.76	0.74	0.74	0.81	0.72	0.76	*0.81*
Means (75%) + SG (25%)	0.74	0.72	0.70	0.86	0.71	0.75	*0.81*

From the test AUC results reported in Table 4-12, the best performing models would appear to be those using soil layers created by adding the SoilGrids raster values to the means by soil type, i.e. using the *Means + SG (50:50)* and *Means + SG (75:25)* hybrid soil layers. These models achieved mean test AUC scores over 0.7 for dryland crops – higher test AUC scores for these crops than all other suitability models reported so far; however, the mean test AUC values for the Rice models were lower compared to models trained using the other hybrid soil layers. The models for the hybrid soil layers with Gaussian noise (*Means + jitter*) had the lowest test AUC scores overall, even after smoothing of the suitability maps. However, their suitability maps achieved the highest mean map correlation score indicating greater similarity of predictions across the five model scenarios for each crop.

Figure 4-12 shows the suitability maps for Maize (model scenario P1) from the sets of models listed in Table 4-12. Contrasting geographical patterns of suitability for maize across the region are apparent in the maps:

- The predictions from the *Means* model are closely related to the soil types from the Gryphon Minerals soil map, with shapes of soil polygons detectable in the suitability map (compare Figure 2-9).
- The *Means + jitter* model predicts large areas of suitability across the landscape, and comparison with the terrain layers (Figure 2-7) reveals that the areas rated unsuitable by this model correspond to locations with extreme values for elevation and wetness (very high or very low values).
- The suitability maps from the *Means + SG (50:50)* and *Means + SG (75:25*) models are quite similar to each other, predicting small pockets of suitability across a landscape that is mostly rated as unsuitable for maize.

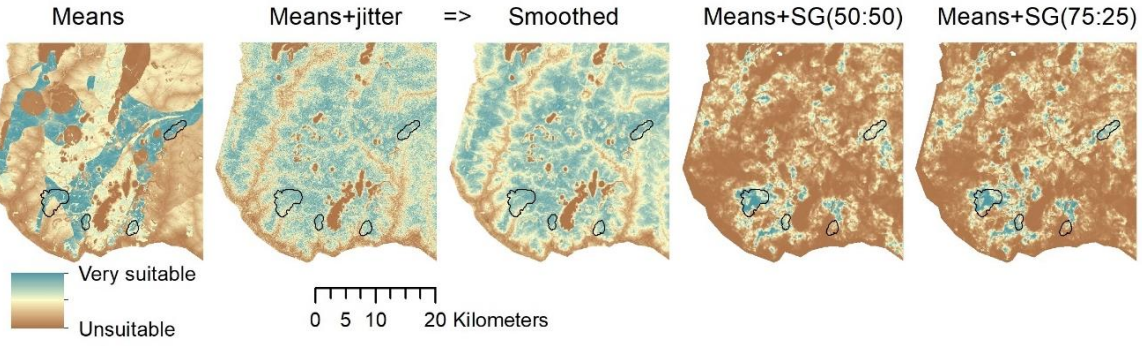

Figure 4-12 Suitability maps for Maize across the Burkina Faso project region from models built using terrain and soil layers (123ABCDEFGHI) for each set of hybrid soil layers

PCA of hybrid soil layers

Principal component analysis was applied to the nine soil layers in three sets of hybrid soil layers to generate corresponding sets of PCA soil layers. Figure 4-13 uses violin plots to show the distributions of values for the first two principal components (PC1, PC2) in each set of PCA soil layers.

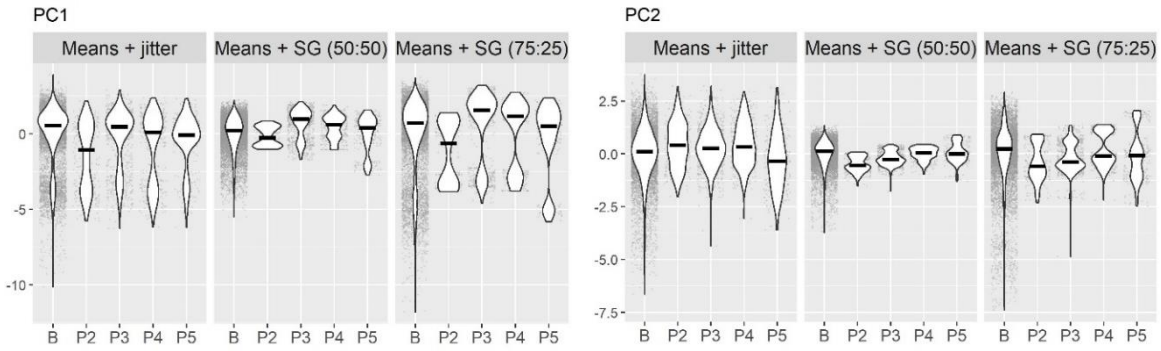

Figure 4-13 Violin plots of PCA soil values with mean (black bar) at background points (B) and at the four mine sites (P2=Fourkoura, P3=Nogbele, P4=Samavogo, P5=Stinger) for layers P=PC1, Q=PC2

Suitability models were trained using three terrain layers (123) and layers P=PC1 and Q=PC2 from the sets of PCA soil layers (all used feature classes LQP and beta multiplier = 1). Summary test AUC results for these models are presented in Table 4-13.

Table 4-13 Summary test AUC values for models built from terrain and PCA soil layers (123PQ) for three sets of PCA soil layers constructed using different methods

Soil layers = PQ	**Mean test AUC (5 model scenarios)**					**5 main crops**	
Method	**Cotton**	**Maize**	**Peanut**	**Rice**	**Sesame**	**Mean AUC**	***Mean Map Correlations***
PCA of Means + jitter	0.62	0.60	0.54	0.89	0.59	0.65	*0.86*
PCA of Means (50%) + SG (50%)	0.84	0.77	0.81	0.77	0.79	0.80	*0.92*
PCA of Means (75%) + SG (25%)	0.85	0.80	0.84	0.87	0.82	0.84	*0.94*

Comparing Table 4-12 and Table 4-13 it can be observed that the use of PCA on the hybrid soil layers reduced measured accuracy when used on the *Means + jitter* layers. However, it noticeably improved overall measured accuracy when used on the *Means + SG (50:50)* and *Means + SG (75:25)* layers, with the latter resulting in mean test AUC scores of over 0.8 for the four dryland crops. The very best overall model performance for the dryland crops from all 47 runs, in terms of the highest mean test AUC scores and the highest mean map correlation scores, occurred for the run that used PCA layers derived from the *Means + SG (75:25)* hybrid soil layers.

Figure 4-14 shows the suitability maps for Maize (model scenario P1) from the sets of models using PCA layers listed in Table 4-13. When comparing the PCA suitability maps to the corresponding maps in Figure 4-12 it is observable that the predictions are very similar for the *Means + jitter* maps, but that using PCA layers for *Means + SG (50:50)* and *Means + SG (75:25)* models has resulted in greater areas of predicted suitability for maize across the region.

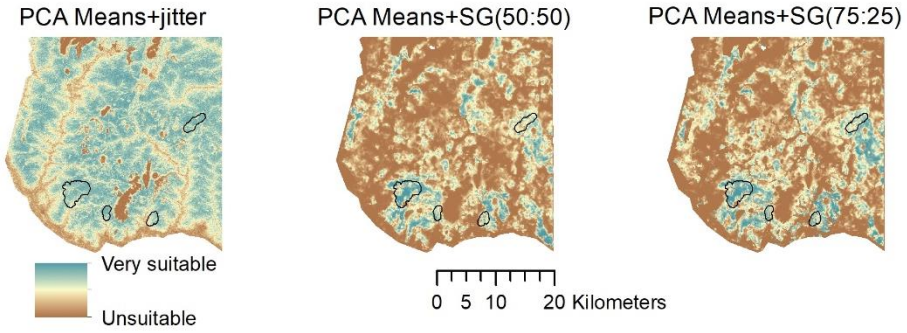

Figure 4-14 Suitability maps for Maize across the Burkina Faso project region from models built using terrain and two PCA soil layers (123PQ) for each set of hybrid soil layers

Qualitative assessment using satellite imagery

The suitability maps illustrated in Figure 4-12 and Figure 4-14 present differing geographical patterns of predicted suitability from models for the same crop across the region and pose the question: which models (if any) are feasible as predictors of agricultural land suitability for the crop?

The use of some of the hybrid soil layers and their corresponding PCA layers has resulted in very high test AUC scores for many crops, possibly exceeding the theoretical maximum achievable AUC for these widely grown crops. Test AUC values are calculated solely from the predictions for presence points; however, in this project the presence points have not been sampled across the project region but are clustered in four small areas. Measures of accuracy based on these points cannot, on their own, provide evidence of good models. It is also necessary to be able to demonstrate that the models are realistic; i.e. that they do not contradict visual interpretations of land use observable in high-resolution Earth observation imagery such as that available in Google Maps, Bing Maps and ESRI's base maps for ArcGIS.

The Worldview2 satellite image of the project region supplied by Gryphon Minerals is used in this chapter to provide visual comparisons between observable land use patterns and modelling results. The image has 0.5 metre pixel resolution and was taken in April 2010 during the dry season and after harvest. Agroforestry is widely practiced in this region, so visible soil surrounding trees in post-harvest images typically indicates areas of cultivation. Note that white areas in any parts of the image indicate patches of cloud cover and that shadows from clouds may also visible.

The next three figures compare suitability maps (from model scenario *P1 Random 25%*) produced by the models reported in Table 4-12 and Table 4-13 to observed land use in three different 4km by 4km areas of the project region. In each figure the suitability maps are displayed next to the satellite image of the area, with corresponding maps of the DEM (altitude), wetness index values and soil polygons from the Gryphon Minerals soil provided below to help illustrate factors affecting suitability for cropping.

Figure 4-15 compares the suitability maps for Maize and Rice in a 4km by 4km area that includes the Stinger mine site. The presence points for the two crops are overlaid on the satellite image and on all relevant suitability maps. The reader is reminded that the models from scenario *P1 Random 25%* do not use spatially independent training and test data but use 75% of the presence points from all sites (including Stinger) for training and the remaining 25% for testing and calculation of AUC scores.

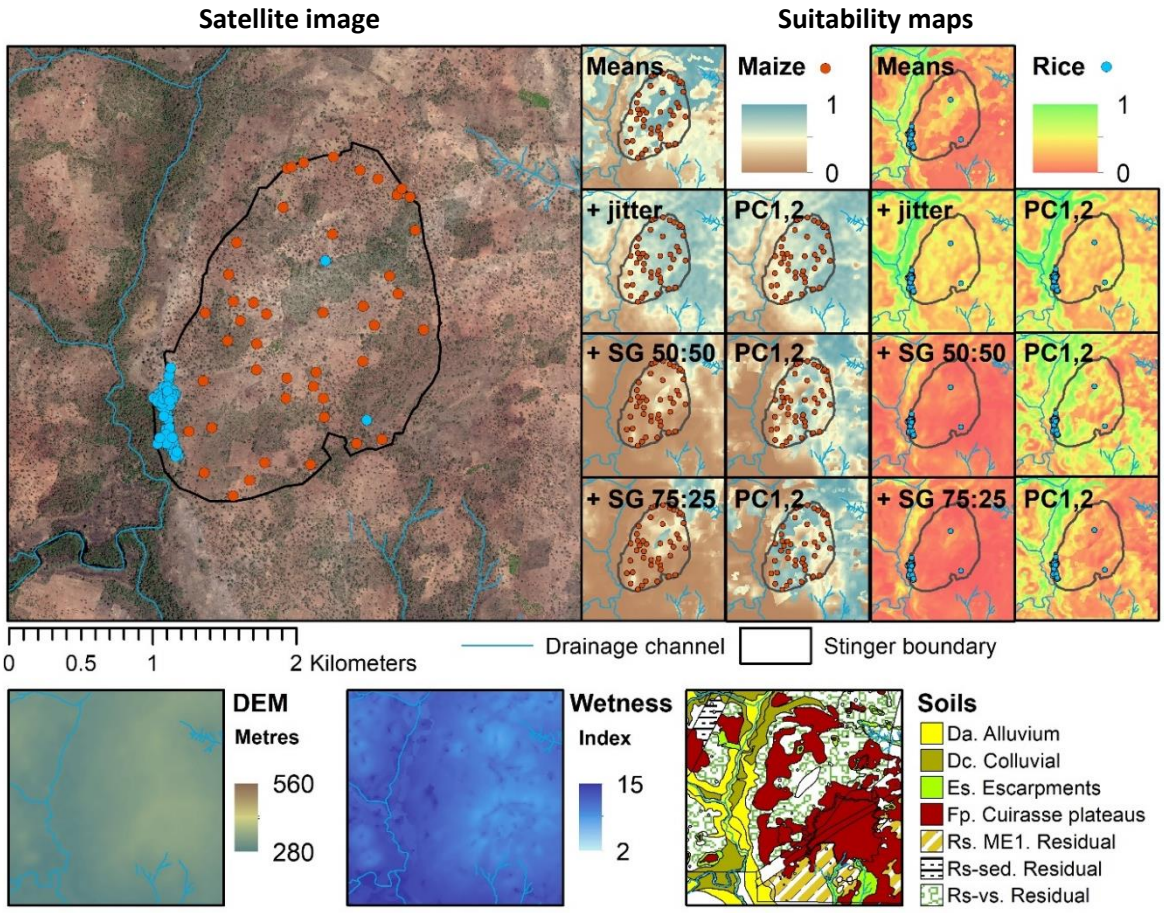

Figure 4-15 Comparison of suitability maps derived using hybrid soil layers for Maize and Rice to corresponding true colour Worldview2 satellite imagery (April 2010) and environmental maps for the Stinger mine site region

The predictions from the *Means* and *Means + jitter* suitability maps in Figure 4-15 are consistent with the known occurrences for the two crops, predicting suitability or marginal suitability at almost all presence point locations. In contrast, the *Means + SG (50:50)* and *Means + SG (75:25)* models demonstrate poor predictive power for the Stinger site – almost all the maize presence locations are rated unsuitable for maize and most rice presence locations have low suitability ratings for rice. The application of PCA to these two sets of hybrid soil layers improved model performance within the Stinger site and their *PC1,2* equivalent models have produced predictions that are far more consistent with the evidence of the presence points. A broad area of visible soil that is indicative of agroforestry cropping can be observed along the left edge of the satellite image. This area is predicted to be unsuitable for both maize and rice by the *Means + SG (50:50)* and *Means + SG (75:25)* models, and their *PC1,2* equivalents have predicted this area to be unsuitable for maize. In contrast, most of this area is predicted to be suitable for either maize or rice by both the *Means* and *Means + jitter* models.

Figure 4-16 and Figure 4-17 compare the suitability maps produced by the models reported in Table 4-12 and Table 4-13 for Maize and Rice with satellite imagery of two 4km by 4km sections of the project region that are distant from the mine sites. Enlargements of the areas marked by black rectangles on the satellite images are shown in Chapter 2 (Figure 2-12 and Figure 2-13) and inspection of these enlargements is helpful for interpretation of the image detail.

In Figure 4-16 dryland agriculture across most of the landscape is observable in the satellite image, with pockets of rice cultivation visible in areas corresponding to the areas of colluvial soil in the soil map. The curaisse plateaus and escarpments marked in the soil map are echoed in the DEM and wetness maps, and are observable in the satellite image as changes in soil colour and lines of uncleared bush.

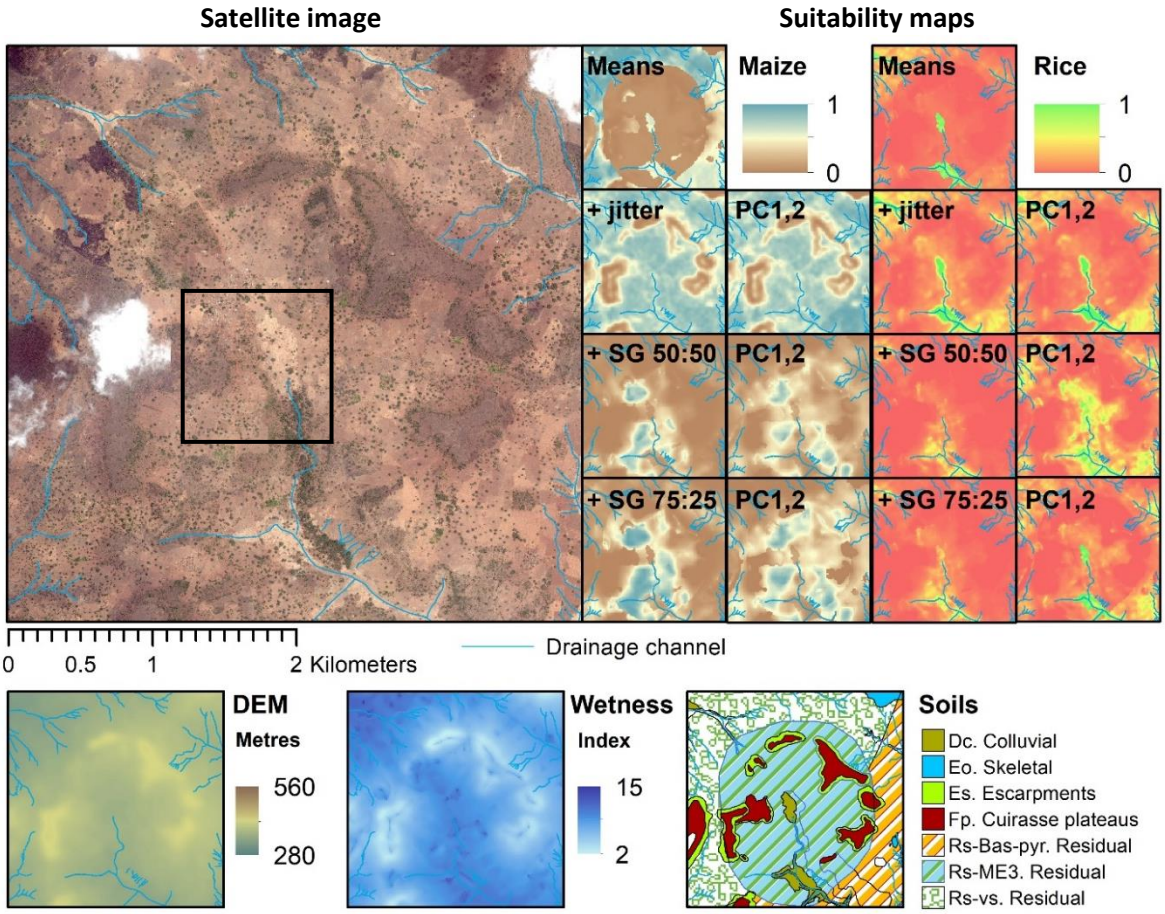

Figure 4-16 Comparison of suitability maps derived using hybrid soil layers for Maize and Rice to corresponding true colour Worldview2 satellite imagery (April 2010) and environmental maps (refer Figure 2-13 on page 36 for enlarged image of marked area); white in the satellite image indicates cloud cover

Inspection of the Maize suitability maps shows great variability in the predictions from the seven models. The *Means + jitter* (smoothed) suitability map appears consistent with the information in the satellite image and environmental maps. Areas of steep slope and lower wetness values are predicted as unsuitable for maize, with all other areas predicted as suitable. The corresponding *PC1,2* suitability map is almost identical and both seem feasible.

In contrast, the Maize suitability map produced using the *Means* hybrid soil layers seems highly infeasible, predicting unsuitability for maize over a large area where the satellite image shows dryland agriculture occurring. It should be noted that the soil type Rs-ME3 does not occur at any of the four mine sites so would not occur in training samples for this model. The maize suitability maps from the *Means + SG (50:50)* and *Means + SG (75:25)* models and their *PC1,2* equivalents also show little correspondence to either the visual evidence in the image or observable features in the environmental maps, and so also seem infeasible. Areas suitable for rice growing are closely linked to soil type. The Rice suitability maps in Figure 4-16 show the areas of colluvial soil sharply

defined in the *Means* map as very suitable for rice. These areas are also rated very suitable in the *Means + jitter* map. Neither the *Means + SG (50:50)* model nor the *Means + SG (75:25)* model predicted high suitability for rice across the area of rice cultivation visible near the centre of the region. However, the *PC1,2* equivalents for both models did make correct predictions for rice suitability for this area.

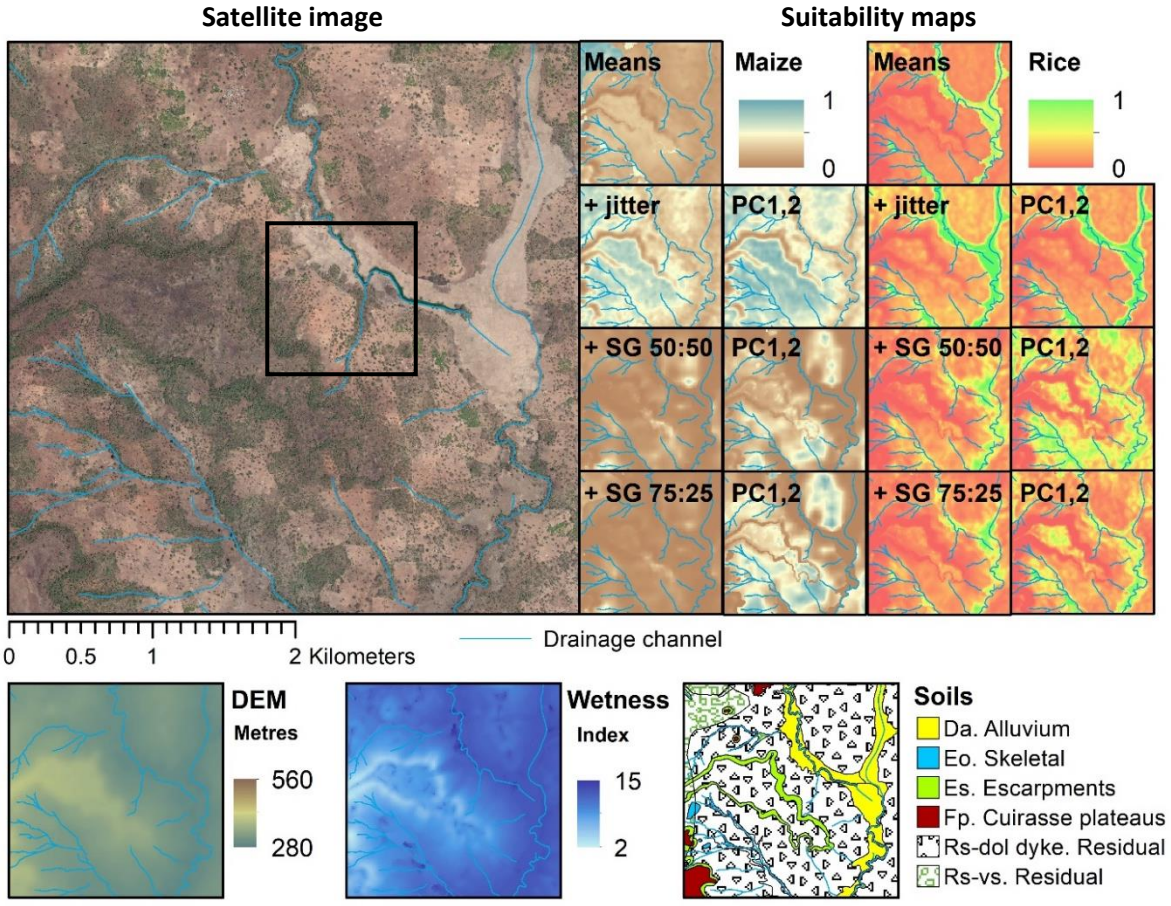

Figure 4-17 Comparison of suitability maps derived using hybrid soil layers for Maize and Rice to corresponding true colour Worldview2 satellite imagery (April 2010) and environmental maps (refer Figure 2-12 on page 35 for enlarged image of marked area)

The region depicted in Figure 4-17 contains a large extent of rice cultivation that is visible in the satellite image and corresponds to the areas of alluvium in the soil map. The Rice suitability maps for the *Means + jitter* model and its *PC1,2* equivalent rate this whole area as highly suitable for rice cultivation, and most other areas as unsuitable. This rice growing area is also clearly visible in the *Means* suitability map, but with slightly lower suitability scores. It is not so clearly distinguished in the *Means + SG (50:50)* and *Means + SG (75:25)* maps or their *PC1,2* equivalents.

With regard to Maize, the *Means + SG (50:50)* and *Means + SG (75:25)* models predict almost no areas of suitability in this region and the *Means* model predicts only a small area of suitability in the top left of the map. This seems at odds with the large areas of dryland agriculture visible in the satellite image. The *Means + jitter* map and its *PC1,2* equivalent both identify areas of unsuitability for maize that correspond to escarpment or to wet alluvial soils and show large areas that are rated as suitable or marginally suitable for maize.

Although qualitative assessment using satellite imagery is illustrated for only three sites and two crops in this thesis, similar examples of predictions inconsistent with observed agriculture occur across the landscape for many of the models. Superficially, the region wide *Means* suitability map in Figure 4-12 appears unrealistic, in spite of the high mean test AUC scores from the *Means* suite of models. Recognition, early in the project, of the insufficiency of categorically valued soil layers as useful predictors had prompted the development of the other sets of hybrid soil layers. So, the high formal accuracy scores for the *Means* models provide a useful reminder of the potential for overfitting to training data and cautions against the use of AUC scores alone as a measure for goodness of fit.

The *Means + SG (50:50)* and *Means + SG (75:25)* models achieved high formal accuracy scores in terms of mean test AUC, and their *PC1,2* equivalents also achieved high mean map correlations scores for similarity of predictions across the model scenarios. However, comparison of the predictions of these models with presence point locations and with observed land was unable to provide qualitative evidence for their feasibility as crop suitability models. From the qualitative assessments illustrated above, only the *Means + jitter* maps and their *PC1,2* equivalents show sufficient consistency with the Earth observation evidence to appear realistic.

4.4 Choosing the best models

This chapter has described the development of multiple crop suitability models to explore the factors that contribute to good models or may lead to the production of poor models. Good models should seem realistic, with suitability ratings that do not contradict Earth observation evidence. Formal evidence for their accuracy are high test AUC scores, and evidence for their robustness are highly correlated suitability maps across a suite of similar models that vary only in the sets of presence points used in their training.

Model parameters

The exploration of Maxent model parameters (refer Section 4.3.3) indicated that the default regularisation values were suitable for this data (beta multiplier = 1), and that the choice of linear, quadratic and product feature classes (LQP) produced the best overall results.

Environmental predictors

Sections 4.3.4 and 4.3.5 explored the effectiveness of environmental predictors. Models were built using combinations of terrain, soil and radiometric layers, and the use of PCA layers was tested. The presence data for this project were spatially clumped in four small areas so cross-validation was used to detect the occurrence of spatial bias in models resulting from the use of certain predictors.

Experiments using terrain layers only showed that models using layers 1=Slope, 2=Wetness and 3=Solar as predictors produced more robust results than models with additional predictors 0=DEM, M=MRVBF or N=MRRTF. They also showed that the use of layers 0=DEM and N=MRRTF, in particular, exacerbated spatial bias in resulting models.

The use of radiometric layers was tested and produced good results in combination with terrain layers. However, the radiometric layers are highly correlated with the soil layers (refer Appendix E) and these data were used in refining the Gryphon Minerals soil map, and so became redundant when used in combination with soil layers.

Early tests using supplied and found soil data (the Gryphon Minerals soil map and SoilGrids layers, respectively) demonstrated the necessity of developing hybrid soil layers in order that soil predictors could be used by the models. Four sets of hybrid soil layers were constructed using different methods, but only the *Means + jitter* hybrid layers produced feasible suitability models. Comparison of the results for models that used the sets of hybrid soil layers (and their PCA equivalents) exposed the weakness of relying on formal performance scores alone as the measures of a good model. The models using hybrid soil layers that achieved the highest test AUC scores and had the most highly correlated suitability maps were demonstrated to be poor models by qualitative assessment using satellite imagery of the region.

Best models

The best set of models from the 47 runs used three terrain layers (123) and the nine jittered hybrid soil layers (ABCDEFGHI) as predictors, and used linear, quadratic and product (LQP) feature classes. The 25 models (for the five crops from the five model scenarios) are compared in the following section, and the best models from this set were chosen as the final output suitability maps for the Burkina Faso project region.

4.4.1 Characteristics of the best set of models

This set of models comprise suitability models for the five main crops created using five model scenarios that differed only in the composition of their training sets of presence points. Four of the model scenarios used spatially independent training and test data and so, in order to provide context for the following results, the numbers and percentages of crop presence points at each of the mine sites are presented in Table 4-14.

Table 4-14 Number and percentage of presence points (centroids of agricultural plots) for each crop at the four mine sites

Mine site	Cotton	%	Maize	%	Peanut	%	Rice	%	Sesame	%	Total	%
Fourkoura	25	*11.2*	50	*14.5*	23	*17.4*	91	*13.5*	20	*9.7*	209	*13.2*
Nogbele	157	*70.1*	194	*56.1*	61	*46.2*	329	*48.8*	91	*44.0*	832	*52.6*
Samavogo	13	*5.8*	57	*16.5*	30	*22.7*	199	*29.5*	45	*21.7*	344	*21.7*
Stinger	29	*12.9*	45	*13.0*	18	*13.6*	55	*8.2*	51	*24.6*	198	*12.5*
Total	**224**	*100*	**346**	*100*	**132**	*100*	**674**	*100*	**207**	*100*	**1,583**	*100*

More than half of all presence points occurred at Nogbele. In the case of Cotton, 70% of presence points were from Nogbele, so the model for Cotton in scenario *P3 Nogbele,* that used training data from the other sites, would have more than twice as many test points as training points.

Test AUC scores

The test AUC scores for the models are presented in Table 4-15. Stinger is wetter and has a different soil profile from the other three sites and the test AUC scores for *P5 Stinger* models are very low for the four dryland crops. This example shows how failing to include representative training samples for a site can reduce the predictive capability of resulting models for that site. This is not always the case as can be seen from the Rice model for scenario *P2 Fourkoura* and the Cotton model for *P4 Samavogo* that both achieved the highest test AUC score for their respective crop despite having no training samples from the test site.

Table 4-15 Test AUC scores for the five main crops in the five model scenarios from models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI)

Model scenario	Cotton	Maize	Peanut	Rice	Sesame	*Mean*
P1 Random 25%	0.71	0.72	0.72	0.89	0.70	*0.75*
P2 Fourkoura	0.66	0.61	0.68	0.94	0.69	*0.71*
P3 Nogbele	0.62	0.69	0.64	0.83	0.65	*0.69*
P4 Samavogo	0.71	0.65	0.61	0.90	0.64	*0.70*
P5 Stinger	0.57	0.47	0.57	0.87	0.40	*0.58*
Mean test AUC	*0.65*	*0.63*	*0.64*	*0.89*	*0.62*	*0.69*

Suitability map correlations

The 25 suitability maps for the five crops in the five model scenarios are plotted in Figure 4-18 for visual comparison to observe similarities and differences.

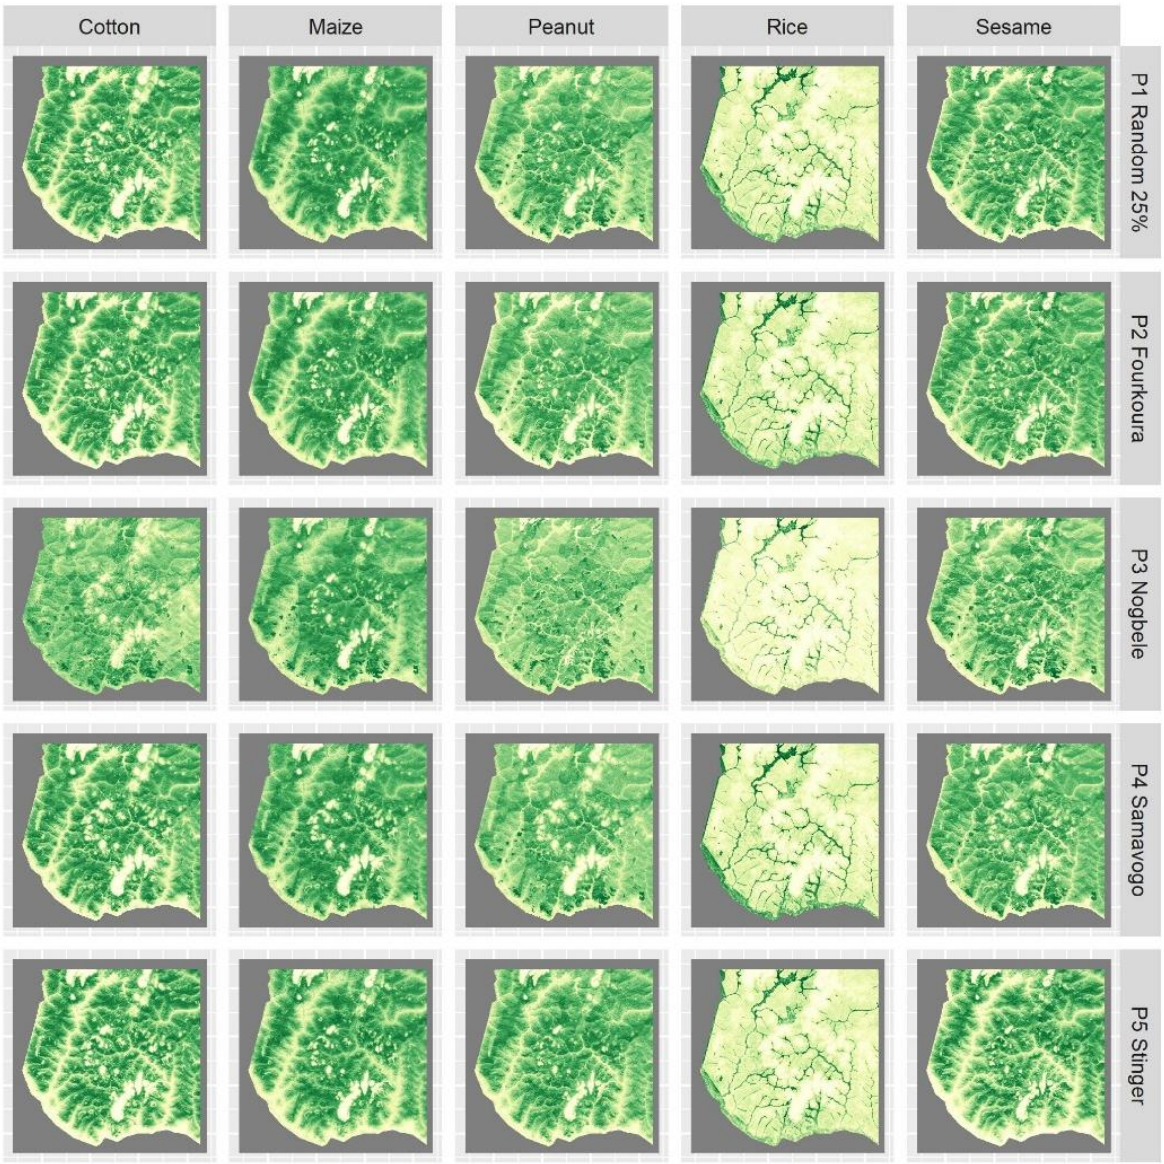

Figure 4-18 Suitability maps from models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI): degree of suitability indicated by greenness

The maps show a high level of consistency regarding the landscape patterns of suitability for each crop across the model scenarios, providing evidence of the robustness for crop models generated by this set of predictors. The qualitative assessments illustrated above in Figure 4-15, Figure 4-16 and Figure 4-17 demonstrated how Maxent models trained using different inputs can produce contradictory suitability predictions for a particular crop in an area, e.g. with one model predicting high suitability for the crop and another rating the area completely unsuitable. Such contradictory results are not evident in Figure 4-18. The maps predict very similar areas of suitability and very similar areas of unsuitability for each crop in all model scenarios, with the main visible difference between the maps being the degree of predicted suitability. Figure 4-19 presents the correlation coefficients for pairwise comparisons of the suitability maps for each crop from the five model scenarios.

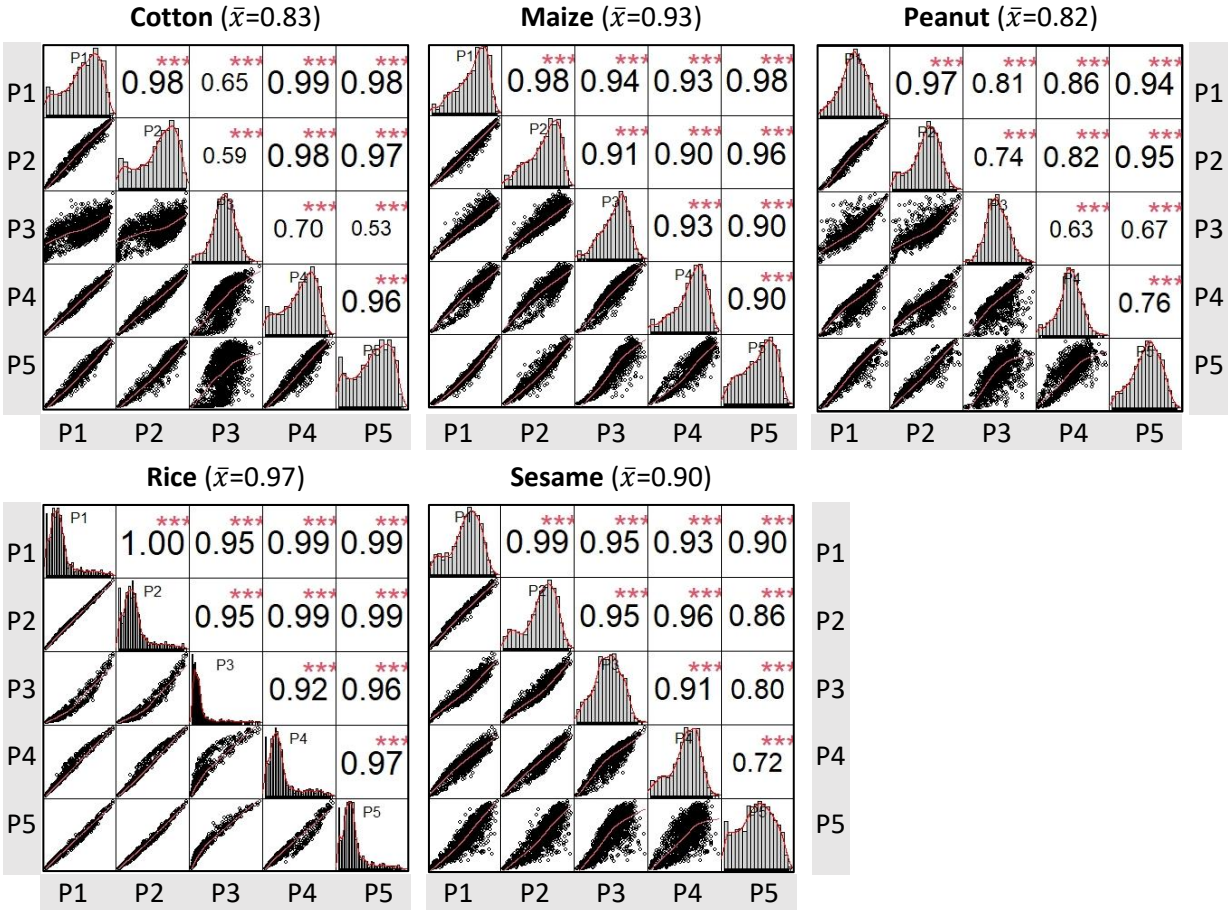

Figure 4-19 Correlations of suitability maps between the five model scenarios from models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI)

The Rice suitability maps from all model scenarios are almost identical and record a mean correlation value from pairwise comparisons of maps of 0.97. The Maize maps are also quite similar and have a mean correlation value of 0.93. The Sesame suitability maps from scenarios *P1-P4* are very similar to each other but show less similarity to the map from scenario *P5 Stinger*.

For both Cotton and Peanut, the lowest correlation coefficients occur for the comparisons with the *P3 Nogbele* suitability maps. The *P3 Nogbele* models for Cotton and Peanut both had less than 80 training samples (67 and 71, respectively) so only linear and quadratic features were used for these models (refer Table 3-1). All other models also used product features for building response curves.

Major divergences in model predictions that would demonstrate substantial spatial bias (such as observed in some of the correlation matrix charts in Figure 4-7) are not apparent in Figure 4-19. The scatterplots reveal some differences in the degrees of predicted suitability for crops between models from different scenarios, but contradictory predictions are outliers. As the suitability maps compared in each of the scatterplots differ only in the samples used for training, these results demonstrate that these models exhibit only minor spatial bias derived from the training samples.

Environmental contributions

The environmental layer contributions for the models are presented in Figure 4-20. The histograms illustrate model reliance on particular layers and show how including or excluding a site from the set of training data may influence the makeup of the crop models.

The environmental profiles of the Rice models are very similar across all model scenarios and have resulted in highly similar Rice suitability maps. This is also the case for the Maize and Sesame models, with similar environmental profiles across their respective sets of models and high mean correlation coefficients for pairwise comparisons of suitability maps across the model scenarios.

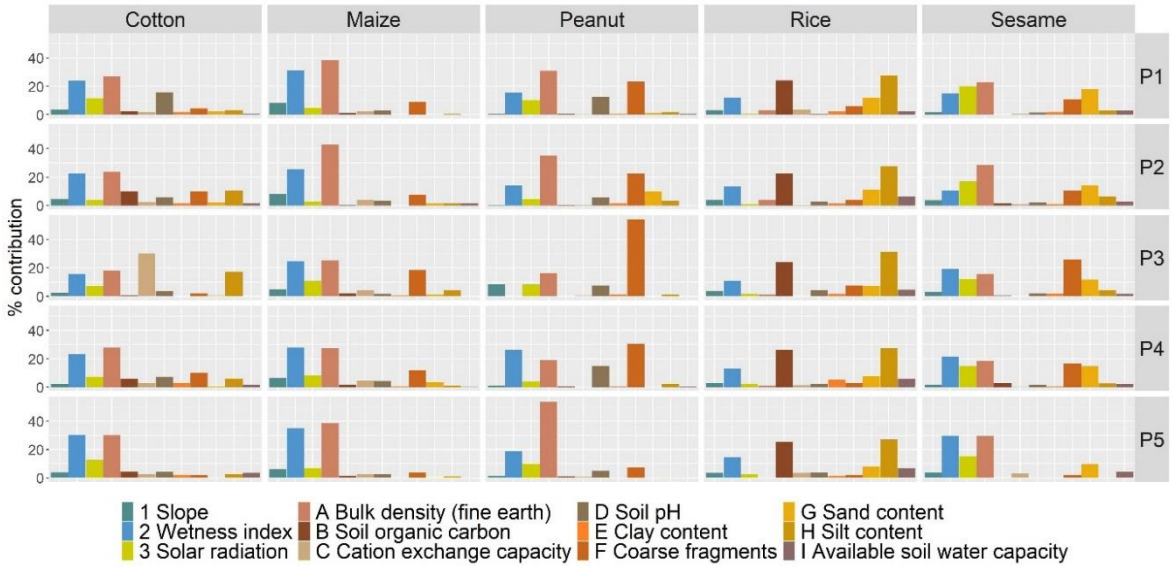

Figure 4-20 Environment contributions (%) to suitability models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise

The environmental profiles of the Cotton and Peanut models from scenario *P3 Nogbele* are least similar to the environmental profiles of the parallel models for these crops. Product features were not used for modelling these two crops in scenario *P3 Nogbele* and training data for the models did not include samples from Nogbele. In the case of Cotton this has resulted in layer *C=Cation exchange capability* emerging as the most significant contributor to this model and the P3 suitability maps being least similar to the other Cotton suitability maps. In the case of Peanut, layer *2=Wetness index* has disappeared as a predictor and the importance of *F=Coarse fragments* has been magnified, with the P3 suitability maps also being least similar to the other Peanut suitability maps.

The five *P5 Stinger* models have no Stinger training samples and all have very low % contributions for layer *F=Coarse fragments*. All other models (that had training samples from Stinger) have higher contributions from this layer, and this is especially obvious in the case of Peanut.

4.5 Final suitability maps

Of the five model scenarios, *P1 Randon 25%* is expected to have produced the best set of crop models for the project region as a whole as these models were trained using a broader geographical range of environmental training samples and so are least likely to exhibit training sampling bias. The models from this model scenario have been chosen as the final outputs for this project.

The final suitability maps, illustrated in Figure 4-21, are the smoothed suitability maps from model scenario *P1 Random 25%* created using three terrain layers (1=Slope, 2=Wetness, 3=Solar) and nine jittered hybrid soil layers (ABCDEFGHI), using feature classes LQP and beta multiplier of 1. The crop models achieved test AUC scores likely to be near the theoretical maximum achievable scores for the target crops, and comparison of the suitability maps with the terrain and soil maps in Chapter 2 shows that they are realistic.

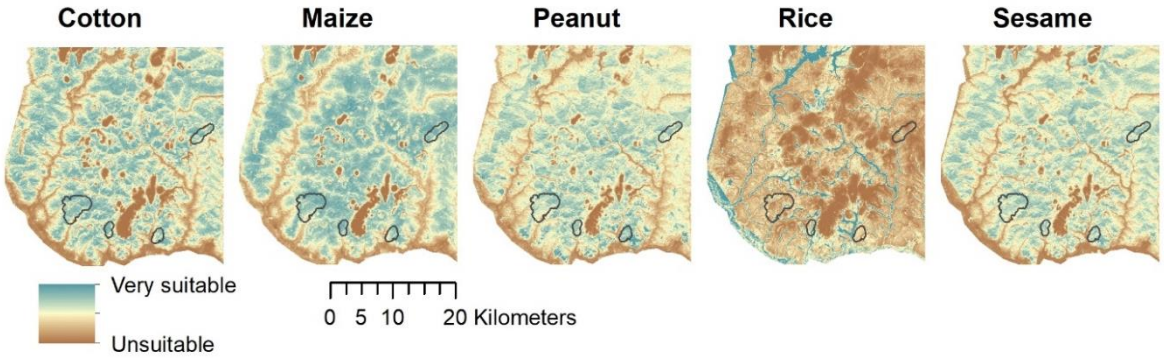

Figure 4-21 Final crop suitability maps for the five main crops from models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise

From the maps it can be seen that some areas are rated as unsuitable for all crops, e.g. outcrops with skeletal soils or very steep terrain. Rice has the most restricted habitat with only the wetter soils near watercourses rated suitable for rice growing. For dryland crops, predicted suitability is higher on better drained soils.

A major challenge with this project was the difficulty in validating results as it is not possible to visit the project area to assess the suitability maps in situ. Satellite imagery proved useful for qualitative assessment, but its usefulness in assessing accuracy is limited. Instead, the detailed planting data for the four mining sites are used to inspect how well the predictions agree with these known planting locations.

Predictions compared with planting locations

Figure 4-22 plots the presence points (using black crosses) for the five main crops over the corresponding suitability maps to allow detailed inspection of the accuracy of predictions for known occurrences. The results are very encouraging, showing that overwhelmingly the known planting locations for each crop occur on land rated by the models as suitable or very suitable for that crop.

The planting locations for rice are densely clustered over land rated very suitable for rice, with very few outliers. For example, near the south-western edge of Stinger the models predict rice as the only crop suitable to grow there, and the planting locations reveal that only rice is actually grown there.

The planting locations for the dryland crops avoid the areas rated highly suitable for rice and other areas unsuitable for agriculture such as skeletal outcrops and escarpment. Where the planting locations occur in areas rated marginally suitable for the crop it is not clear whether these demonstrate limitations of the models or are examples where farmers must make do with the land available.

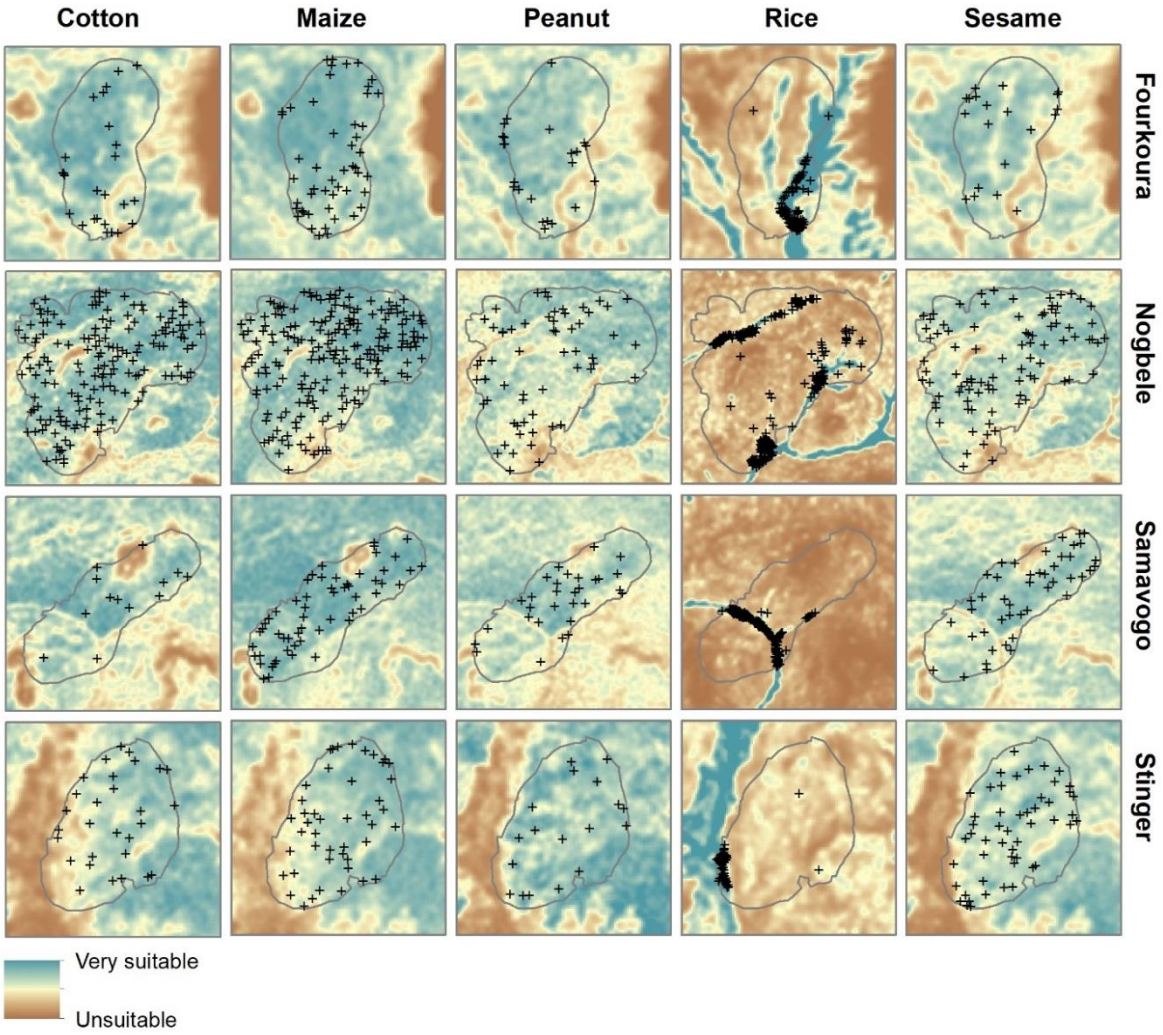

Figure 4-22 Predictions compared to planting locations (+) for the five main crops at the four mine sites - models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise

4.6 Summary

This chapter has demonstrated the application of species distribution modelling techniques to the task of land suitability assessments for agriculture. The primary data sources for the modelling were originally acquired for the purpose of mining exploration and these were supplemented with additional freely available spatial data from public data sources.

The environmental niche modelling algorithm Maxent was used to develop crop suitability models for a region of Burkina Faso in West Africa where subsistence farmers would be relocated to make way for a new gold mine. The compensation maps prepared as part of the application for the mining licence provided known occurrences for locally grown crops and were used to generate presence points for training and testing Maxent. The configuration of the presence data, clustered in four small areas, presented the likelihood of spatial sampling bias affecting models, and the evaluation of modelling results was also complicated by the lack of test data outside of these four small areas.

The methodology that was devised for the task took advantage of the geographical separation of the four presence data sites to develop cross-validation models. Modelling for particular sets of predictors and model parameters was replicated in five model scenarios using different partitions of the presence data for training and testing, four of which had test data that was spatially independent of the training data. The accuracy of the predictions (using test AUC) and similarity of resulting suitability maps (from correlations) then were compared to assess model accuracy, robustness and sensitivity to sampling bias.

Multiple tests of environmental predictors and parameters were performed in order to identify those predictors and parameters contributing to the most accurate and robust models. Summary results were presented for twelve crops, however detailed results were presented only for the five main crops (with more than 100 presence points). Estimates for fractional predicted area and thus theoretical maximum AUC provided useful model accuracy benchmarks for the five crops.

Testing of Maxent parameters showed that the combination of linear, quadratic and product feature classes together with the beta multiplier value of 1 were optimal overall for this task. The tests using different sets of environmental layers revealed that the use of certain layers degraded model results (DEM and index of ridge top flatness) and some layers became redundant when used in combination with other layers (index of valley bottom flatness and radiometric layers). Use of principal component layers was found to mitigate extreme sample selection bias, but did not improve overall results.

Developing effective environmental layers to characterise soil properties was a major challenge in the project. The categorical map of soil type (from Gryphon Minerals) had proven unsuitable as a predictor for Maxent, as had the raster layers of soil properties derived from SoilGrids. Four sets of new raster soil layers were synthesised from the categorical and raster soil data: using the mean soil property value in each soil category for the first set (17 discrete values per layer); adding Gaussian noise to the means for the second set; and adding the soil property raster values to the means in different proportions for the third and fourth sets. Principal component layers were also created from the last three sets. The effectiveness of these seven sets of hybrid soil layers as model predictors was quantitively compared using AUC and correlation results, and later qualitatively assessed by comparing the suitability maps with high-resolution satellite imagery and with the locations of known occurrences.

Using the first set of hybrid soil layers (soil means) produced suitability maps dominated by the shapes of the soil polygons that appeared unrealistic, but still achieved high test AUC scores. The

layers with added Gaussian noise resulted in realistic looking models, but with lower measured accuracy. The final sets of hybrid soil layers (with added SoilGrids values) resulted in models with very high test AUC scores and their PCA versions had even higher scores, far exceeding the estimated theoretical maximum AUCs for all dryland crops.

The suitability maps from models trained using the different sets of hybrid soil layers presented markedly different geographical patterns of predicted suitability for the same crop across the landscape, posing the question: which models (if any) are feasible as predictors of agricultural land suitability for their target crop? Comparing the suitability maps with the satellite imagery revealed numerous examples of predictions inconsistent with observed agriculture occurring across the landscape. The qualitative assessment demonstrated that, for most of these models, the high test AUC scores calculated from the clustered presence data were not indicative of model accuracy elsewhere.

Only the models using the hybrid soil layers with added Gaussian noise had produced results that were consistent with the visual interpretations of land use in the satellite images. The test AUC scores for these crop models was also consistent with the estimated theoretical maximum AUC scores achievable for the crops. All the suitability maps from these models exhibited random noise as a consequence of the noisy soil predictor layers. Smoothing the maps improved their visual appearance and also increased the AUC and correlation scores. These models were used for the final project output.

For this final set of models, the five models for each crop (from the five model scenarios) exhibited similar patterns of predicted suitability across the region, indicating that spatial sampling bias was minor. The environmental contributions to the models for each crop were also similar. The final crop suitability maps chosen as the outputs for this project were the smoothed maps produced in the first model scenario that used training data from all four occurrence sites. These maps will appear again in Chapter 6 where methods for presenting these predictions for dissemination to the intended target users are explored.

Chapter 5 Validation of method: South Australia

5.1 Introduction

The agricultural land suitability assessment described in Chapter 4 predicted suitable land for growing particular crops across a 1,100 square kilometre region in Burkina Faso using inputs derived from the detailed maps of agricultural land use at four small communities in the region. Rich data were available to test the accuracy of these predictions within the four communities. But data were not available to test the predictions outside of the communities and it would not be possible to travel to the region to assess their usefulness in situ.[21] Any doubts as to the robustness of the method or questions as to the accuracy of predictions at other locations distant from the communities could not be resolved using the available data. As observed in Section 3.2, another site that allowed region-wide testing of results was needed to validate the methodology.

In this chapter, the methodology from Chapter 4 is duplicated at two local test locations for which region-wide validation data exists. These locations have landscapes that are familiar to the researcher and that can be visited to verify results or collect more data. A parallel test region of dimensions 40km by 40km was initially defined near Adelaide in South Australia. Parts of this region had missing data for soils and radiometrics, so a second test region of the same size was created north of Adelaide in an area where these data were complete. To duplicate the conditions of the Burkina Faso task, the same environmental layers were created and presence data for locally grown crops was simulated in four small training areas at each site. Published maps showing land use potentials for these crops could be compared to the crop suitability maps generated by the models, allowing region-wide validation of results.

5.2 Parallel test regions

5.2.1 Selection

To mimic the configuration of the modelling from Chapter 4, the parallel test regions are the same size as the Burkina Faso project region and each contains four small local areas for presence data to be generated. The Adelaide test region, in Greater Adelaide, was chosen as it contained agricultural land close to where the research was being undertaken. The Marrabel test region, north-east of Adelaide, was created later in order to test the effectiveness of radiometric data in the modelling. As proxies for the Burkina Faso local communities, four agricultural towns were selected at random from each test region and a two-kilometre buffer was drawn around each town. Maps showing the location of the parallel test regions and satellite imagery of the two regions are given in Figure 5-1. The precise geographic locations of these test regions are provided in Appendix C.

[21] The Burkina Faso project region borders Mali and falls within an Australian Department of Foreign Affairs and Trade defined “do not travel” zone.

Figure 5-1 ERSI base maps showing: Marrabel test region (top right); Adelaide test region (bottom right); locations of the test regions (top left) with Greater Adelaide Capital City Statistical Area (source ABS) outlined in blue; Australia and location of Adelaide (bottom left).

The choice of crops to model for each test region was based on: (1) which crops were typically grown there; and (2) whether validation data existed showing agricultural land suitability for the crops across the region. Very detailed crop planting maps, such as the Burkina Faso community land use maps, were not available for the South Australian sites. However, detailed maps showing the agricultural land potential for growing a wide range of crops across the South Australian agricultural zone were available. These maps are presented in the AgInsight interactive mapping

portal[22] and are discussed in more detail in Section 5.3. The land use potential maps provided the validation data needed for comparison with model generated crop suitability maps, and also provide a sound basis from which to generate pseudo presence data for the small training sites (described in Section 5.4.2)

5.2.2 Marrabel test region

The Marrabel test region is characterised by broad acre farming, with grain production and sheep grazing predominating (refer *Source: I. Ahmer*

Figure 5-2). Six crops were selected for suitability modelling at the Marrabel site following a visit to the region to observe local agriculture: Barley, Canola, Chickpeas, Lucerne (dry land), Olives and Wheat. Buffers of radius two kilometres were defined around the towns of Marrabel, Ngapala, Saddleworth and Tarlee, and these four locations were used for the development of pseudo-presence data for the six crops.

Source: I. Ahmer

Figure 5-2 Landscape at Marrabel with wheat field in the foreground (August 2019)

5.2.3 Adelaide test region

The Greater Adelaide region (outlined in blue in the top left panel of *Source: I. Ahmer*

Figure 5-2) has a rich agricultural sector and produces most of the vegetables and much of the fruit grown in South Australia. The Adelaide test site is situated in the Greater Adelaide region and also intersects two of the state's top wine regions: McLaren Vale in the south and Adelaide Hills to the east.

Six crops were selected for suitability modelling at the Adelaide site: Apples, Brassicas, Faba beans, Grape vines, Olives and Potatoes. Buffers of radius two kilometres were defined around the towns of McLaren Flat, Meadows, Oakbank and Uraidla, and these four locations were used for the development of pseudo-presence data for the six crops.

[22] The AgInsight mapping portal is available at https://www.aginsight.sa.gov.au/. It provides comprehensive agricultural and economic data from South Australia and received awards in 2016 for public sector digital innovation and geospatial excellence.

5.3 Validation data

Land use potential is defined as the potential of soil and land to sustain a specific crop. It describes the capability of land for a particular use and is very relevant in agricultural land suitability assessments (which often consider additional influences such as economics, climate, pest and disease incidence, regulations, social factors, etc) (Rowland, Maschmedt et al. 2016).

The *Land use potential* maps used by the AgInsight mapping portal are based on regional scale soil and land mapping undertaken by the South Australian Government. The spatial datasets are available for download from *Data.SA* (https://data.sa.gov.au/) – the South Australian Government Data Directory. The *Land use potential* assessments for the crops and the mapping methodology employed to generate the maps are summarised in Rowland, Maschmedt et al. (2016).

The *Land use potential* spatial datasets were developed for 45 crops and pasture types using an assessment scheme described in the Maschmedt (2002) document "Assessing Agricultural Land." Five *Land use potential* classes were defined based on land and soil attributes only (and assuming management according to standard industry practices), as listed in Table 5-1, with a sixth class to denote land unsuitable for agriculture.

Table 5-1 Land use potential class definitions, source: Rowland, Maschmedt et al. (2016)

Class	Potential	Definition
Class 1	High	Land with high productive potential and requiring no more than standard management practices to sustain productivity.
Class 2	Moderately high	Land with moderately high productive potential and / or requiring specific, but widely accepted and used, management practices to sustain productivity.
Class 3	Moderate	Land with moderate productive potential and / or requiring specialized management practices to sustain productivity.
Class 4	Moderately low	Land with marginal productive potential and / or requiring very highly specialized management skills to sustain productivity.
Class 5	Low	Land with low productive potential and /or permanent limitations which effectively preclude its use.
Class X	Not applicable	Urban, evaporation pans, quarry, water, rock, saline soil, reservoir, cliff, reef etc.

The *Land use potential* maps were derived by characterising, for each crop, the polygons of a reference soil landscape map in terms of these classes. Many map units contained a mix of classes and so could not be assigned to a single class. To accommodate the variable land use potential within map units, a set of ten mapping categories was devised, and map units were categorised according to the proportions of each class occurring in the map unit.

Table 5-2 Land use potential mapping categories, source: Rowland, Maschmedt et al. (2016)

Mapping category	Proportion of land with moderate to high potential	Most common potential class
Aa	More than 60%	High potential (mostly *Class 1*)
Ab	More than 60%	Moderately high potential (mostly *Class 2*)
Ac	More than 60%	Moderate to high (mixed)
Ad	More than 60%	Moderate potential (mostly *Class 3*)
B	30-60%	Low to high potential (mixed)

C	10-30%	Moderately low to low potential (mixed)
D	1-10%	Moderately low to low potential (mixed)
Ea	Less than 1%	Moderately low potential (mostly *Class 4*)
Eb	Less than 1%	Low potential (mostly *Class 5*)
X	–	–

Table 5-2 lists the ten *Land use potential* mapping categories, and the rules used to derive them from the *Land use potential* classes are provided in Appendix G.

Maps of the *Land use potential* for the crops selected for modelling at the Marrabel and Adelaide test regions are shown in Figure 5-3 and Figure 5-4, respectively.

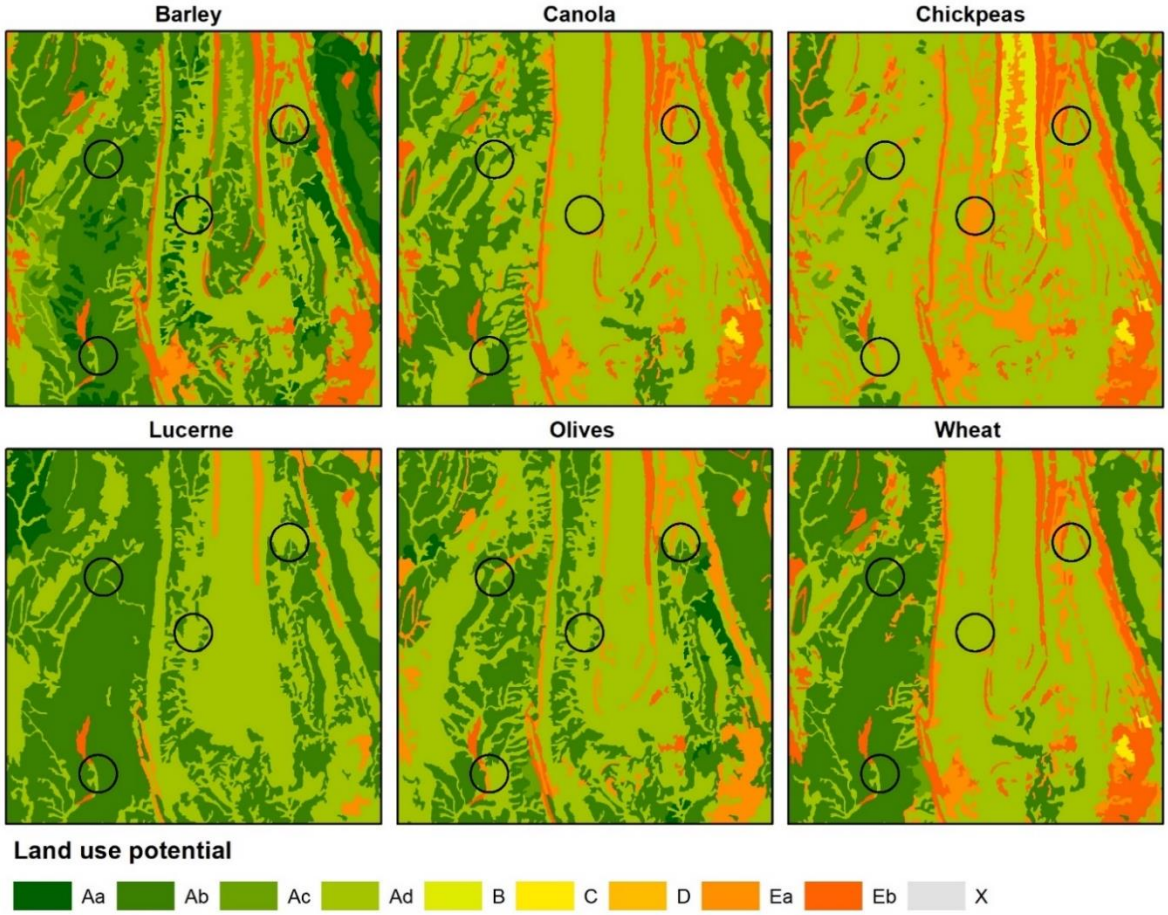

Figure 5-3 Maps showing the Land use potential for the six crops at the Marrabel test region (Source: Data SA)

The *Land use potential* crop maps make an impressive contribution to our understanding of South Australia's agricultural landscape. However, in using these maps it is important to be aware of their limitations. Rowland, Maschmedt et al. (2016) identify a number of key considerations with respect to the use of the *Land use potential* models and maps, including:

- The models are preliminary and have not been subject to field validation.
- The models are based on soil and landscape properties alone and no account has been taken of water quality or availability, climatic factors or existing land use.

- Soil landscape map units are not homogeneous entities – the mapping category is intended to reflect the most common characteristics of the landscape and unspecified variations occur.
- The boundaries between mapping units should be treated as transition zones between legend categories as these boundaries are rarely as sharp in the landscape as the maps imply.
- The maps are intended to provide a regional overview and should not be used to draw conclusions about conditions at particular locations.
- No account was taken of varietal or cultivar differences of individual crops.

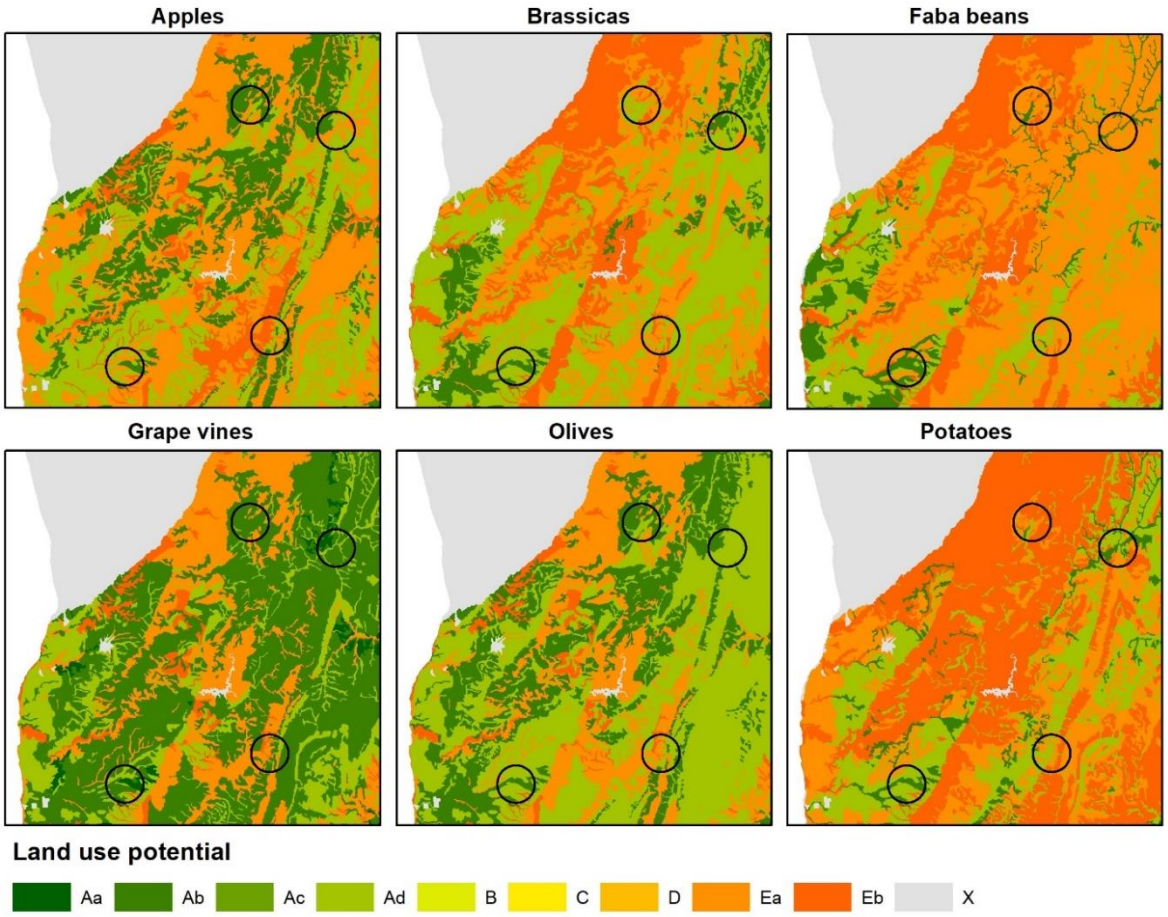

Figure 5-4 Maps showing the Land use potential for the six crops at the Adelaide test region (Source: Data SA)

5.4 Methods and data

The methodology described in Chapter 4 is duplicated for both South Australian locations, with five model scenarios defined for each region and equivalent environmental layers prepared. Presence data was not pre-existing for the SA sites so had to be explicitly generated.

Figure 5-5 provides an overview of the processing performed for the SA test regions and shows the processing inputs and outputs. The *Land use potential* maps have been used as the source for both the presence data and the verification maps. The availability of the region-wide verification maps allows the accuracy of the crop suitability maps to be formally measured – this was not possible for the suitability maps described in Chapter 4 for the Burkina Faso test region.

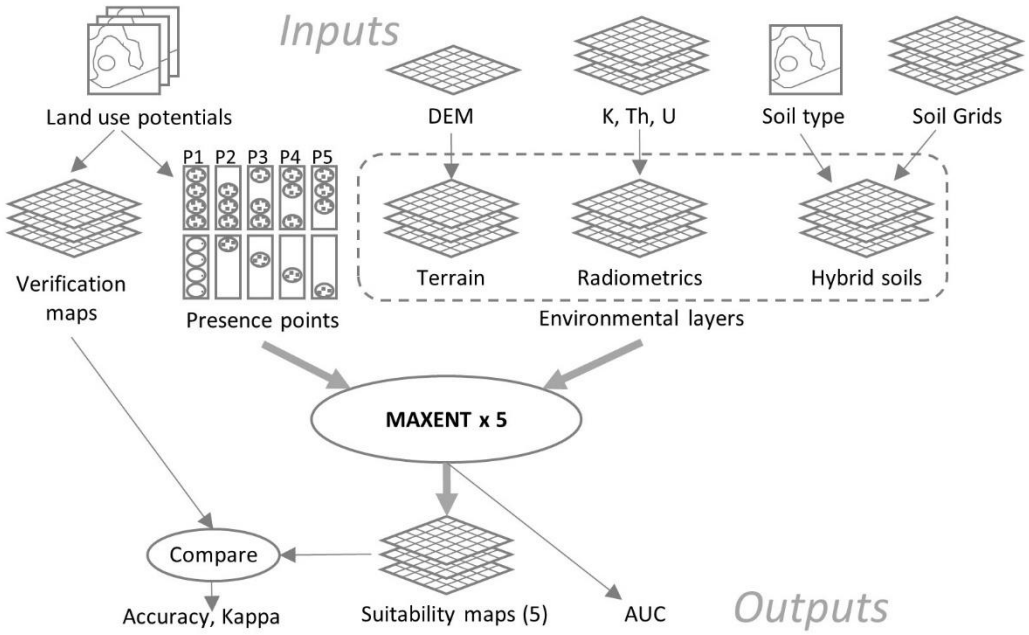

Figure 5-5 Data flow diagram showing inputs and outputs used to model the SA test regions

5.4.1 Model scenarios

The five model scenarios defined for the Marrabel and Adelaide test regions are listed in Table 5-3 and Table 5-4, respectively. Comparison of the results for the five model scenarios is useful in identifying those environmental layers most likely to result in sampling bias.

Table 5-3 Model scenario definitions for the Marrabel test site

Model	Model name	Test data for model	Training data for model
P1	Random 25%	25% of presence points – randomly selected	Remaining 75% of presence points
P2	Marrabel	Presence points at Marrabel	Presence points at the other 3 sites
P3	Ngapala	Presence points at Ngapala	Presence points at the other 3 sites
P4	Saddleworth	Presence points at Saddleworth	Presence points at the other 3 sites
P5	Tarlee	Presence points at Tarlee	Presence points at the other 3 sites

Table 5-4 Model scenario definitions for the Adelaide test site

Model #	Model name	Test data for model	Training data for model
P1	Random 25%	25% of presence points – randomly selected	Remaining 75% of presence points
P2	McLaren Flat	Presence points at McLaren Flat	Presence points at the other 3 sites
P3	Meadows	Presence points at Meadows	Presence points at the other 3 sites
P4	Oakbank	Presence points at Oakbank	Presence points at the other 3 sites
P5	Uraidla	Presence points at Uraidla	Presence points at the other 3 sites

5.4.2 Presence data

Data showing actual crop plantings are not easily available for the South Australian test sites, and the field work that would be required to collect such data is outside the scope of this research project. In the absence of actual presence data for the crops to be modelled, pseudo-presence data has been generated using the crop suitability ratings from *Land use potential* maps. The mapping categories *Aa*, *Ab*, *Ac* and *Ad* (refer Table 5-2) identify areas with moderate to high potential for a target crop and so offer more likely planting locations for that crop.

The pseudo-presence points for each crop have been created in those areas where the *Land use potential* for the crop was rated at *Aa*, *Ab*, *Ac* or *Ad*. The centroids of the *Land use potential* polygons could not be used as presence points: the polygons are very much larger than the field sizes from the Burkina Faso modelling and also are irregular non-convex shapes so that the centroid frequently occurs outside the polygon. Instead, the pseudo-presence data were created by using the polygons rated with moderate to high potential to clip grids of points. The spacing of the pseudo-presence points varied with the *Land use potential* rating to mimic the greater likelihood of more suitable areas being cultivated with that crop, as shown in Table 5-5. For areas rated as having high potential for a crop (*Aa*), each pseudo-presence point corresponds to one hectare spacing, matching the average plot size in the Burkina Faso data. The process for generating the weighted sets of pseudo-presence points for the individual crops is illustrated in Figure 5-6.

Table 5-5 Density of pseudo-presence points for each Land use potential mapping category

Mapping category	Land use potential	Grid spacing	Corresponding area for each pseudo-presence point
Aa	High potential	100 metres	1 hectare
Ab	Moderately high potential	200 metres	4 hectares
Ac	Moderate to high	300 metres	9 hectares
Ad	Moderate potential	400 metres	16 hectares

Note that there is overlap in the suitability areas that have been used to generate pseudo-presence data for the crops at the South Australian sites. This is different from the actual crop plantings data for Burkina Faso, where typically only one crop occurs per plot. This resulted in much larger numbers of pseudo-presence points at the South Australian sites (3,882 for Marrabel and 3,191 for Adelaide for six crops) than true presence points at the Burkina Faso site (1,912 for twelve crops). To adjust for this discrepancy, a random sample of 1,000 pseudo-presence points was used for each of the SA parallel test sites.

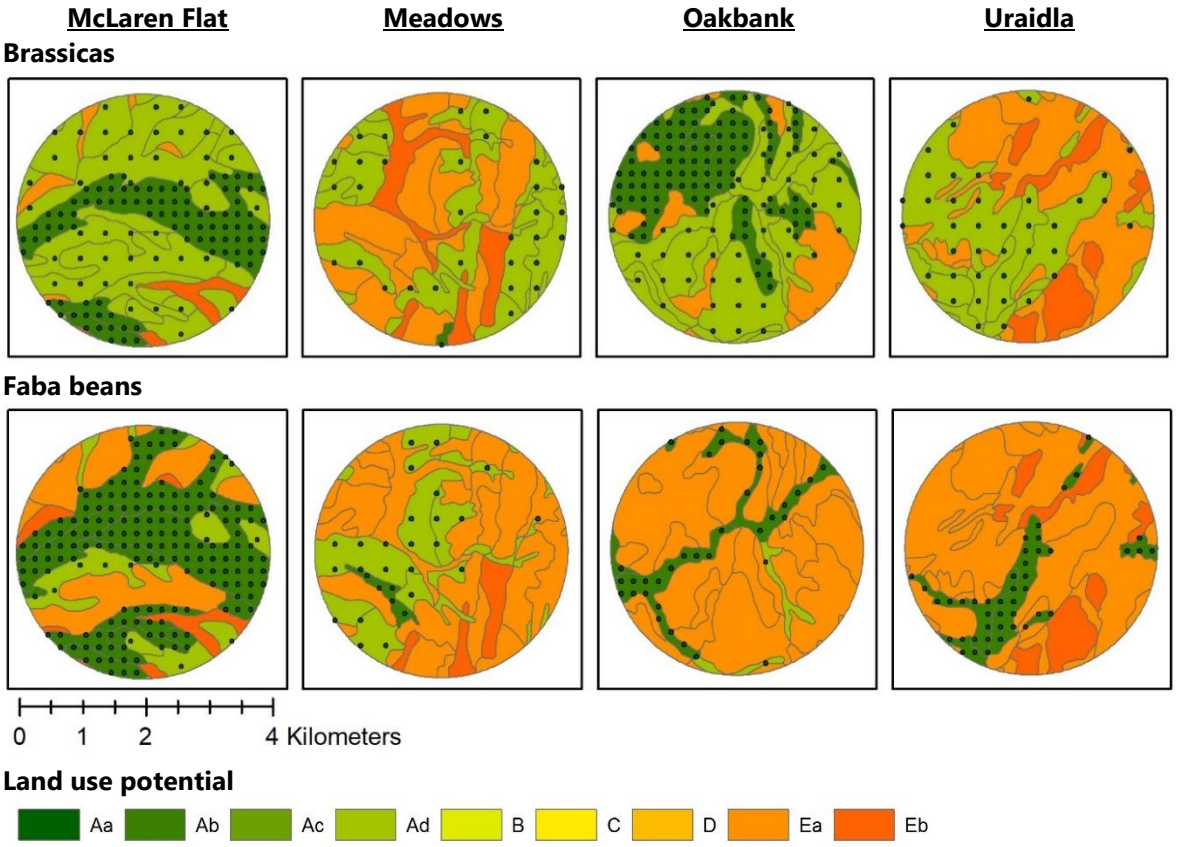

Figure 5-6 Derivation of weighted pseudo-presence data for brassicas and faba beans at the Adelaide test site

The numbers of pseudo-presence points used for each crop are documented in Table 5-6 for the Marrabel test region and Table 5-7 for the Adelaide test region. These tables identify the number of training and test presence points used for each crop in each of the five model scenarios tested at each location.

Table 5-6 Number of training and test presence points used for each crop in each model scenario for the Marrabel test region

	Model scenario	Type	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Total
P1	**Random 25%**	**Training**	207	82	45	156	131	132	753
	Random 25%	Test	68	27	15	51	43	43	247
P2	**Marrabel**	**Training**	187	87	54	172	146	154	800
	Marrabel	Test	88	22	6	35	28	21	200
P3	**Ngapala**	**Training**	236	98	45	175	136	162	852
	Ngapala	Test	39	11	15	32	38	13	148
P4	**Saddleworth**	**Training**	208	75	39	133	108	103	666
	Saddleworth	Test	67	34	21	74	66	72	334
P5	**Tarlee**	**Training**	194	67	42	141	132	106	682
	Tarlee	Test	81	42	18	66	42	69	318
		Total	275	109	60	207	174	175	1000

Table 5-7 Number of training and test presence points used for each crop in each model scenario for the Adelaide test region

	Model scenario	Type	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes	Total
P1	**Random 25%**	**Training**	102	86	65	289	123	87	752
	Random 25%	Test	34	28	21	96	40	29	248
P2	**McLaren Flat**	**Training**	107	69	29	268	122	66	661
	McLaren Flat	Test	29	45	57	117	41	50	339
P3	**Meadows**	**Training**	119	104	79	331	140	93	866
	Meadows	Test	17	10	7	54	23	23	134
P4	**Oakbank**	**Training**	109	65	76	257	126	81	714
	Oakbank	Test	27	49	10	128	37	35	286
P5	**Uraidla**	**Training**	73	104	74	299	101	108	759
	Uraidla	Test	63	10	12	86	62	8	241
		Total	136	114	86	385	163	116	1000

5.4.3 Environmental layers

The environmental layers used in the modelling for the two South Australian test regions were generated using the methods described in Chapter 2 for the Burkina Faso environmental layers. Where possible, the same data sources have been used (SRTM tiles for the DEM, and SoilGrids250m layers).

Terrain layers

A digital elevation model (DEM) was created for the Adelaide region by mosaicking four SRTM tiles (S34E138, S34E139, S35E138 and S35E139). DEMs for the two test sites were created by clipping the mosaic tile, reprojecting to the WGS 1984 UTM Zone 54H coordinate system and applying a simple smoothing filter of radius 3.

From each DEM, five additional terrain layers were generated using identical methods to those described in Chapter 2:

- **Slope** – generated from the DEM using RSAGA tool *slope.asp.curv*.
- **Wetness index** – derived from the DEM by first filling sinks using the RSAGA tool *fill.sinks* using *method=wang.liu.2006*, and then running the RSAGA tool *wetness.index*.
- **Solar radiation** – derived from the DEM using RSAGA tool *pisr2* for latitude=-35.
- **Multi-resolution valley bottom flatness** – created from the DEM using SAGA Morphometry tool *MRVBF*.
- **Multi-resolution ridge top flatness** – created from the DEM using SAGA Morphometry tool *MRVBF*.

Maps of the six terrain layers used for the Marrabel test region are provided in Figure 5-7. The distributions of terrain layer values at the four data presence sites are compared in Figure 5-8 using violin plots to show the shapes of the distributions and black bars to identify mean values. The red line in each plot represents the mean value across the test region for that environmental predictor.

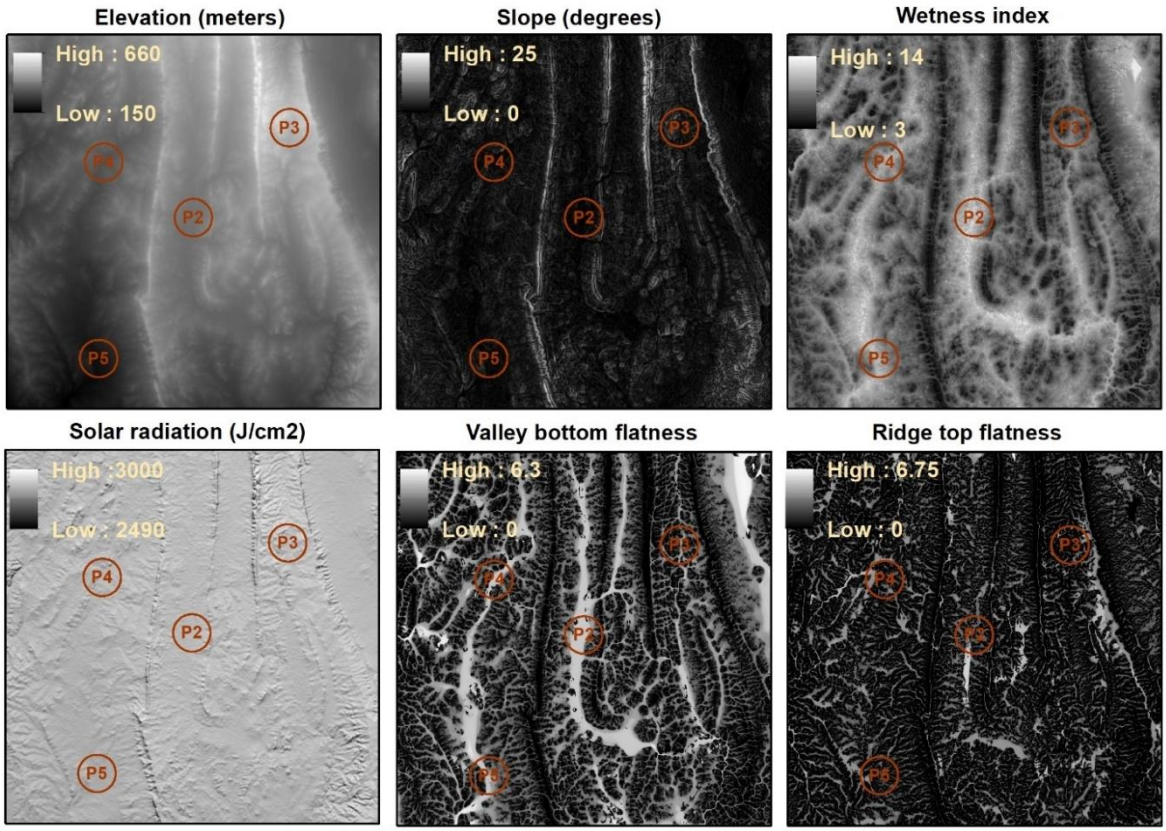

Figure 5-7 Maps of the terrain layers used for the Marrabel test region

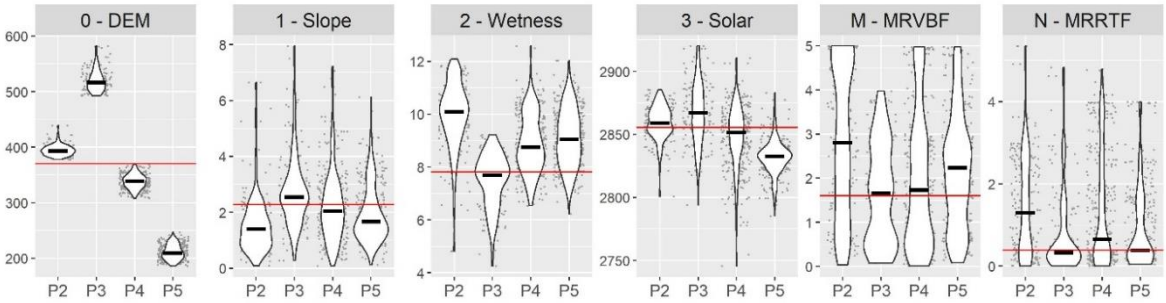

Figure 5-8 Distribution of terrain layer values at presence data sites (P2 Marrabel, P3 Ngapala, P4 Saddleworth and P5 Tarlee) compared to the mean for the Marrabel test region (horizontal red line)

It is observed that site *P3 Ngapala*, at the highest altitude, has steeper slopes and is drier than the other three sites. Site *P2 Marrabel* is the wettest and has more flat land, evidenced by the lowest mean value for slope and highest mean value for valley bottom flatness. Sites *P4 Saddleworth* and *P5 Tarlee* are at lower altitudes – both are wetter and have lower mean slope than the test region as a whole.

Maps of the terrain layers for the Adelaide test region are given in Figure 5-9 and distributions of terrain layer values at the presence data sites are compared in Figure 5-10.

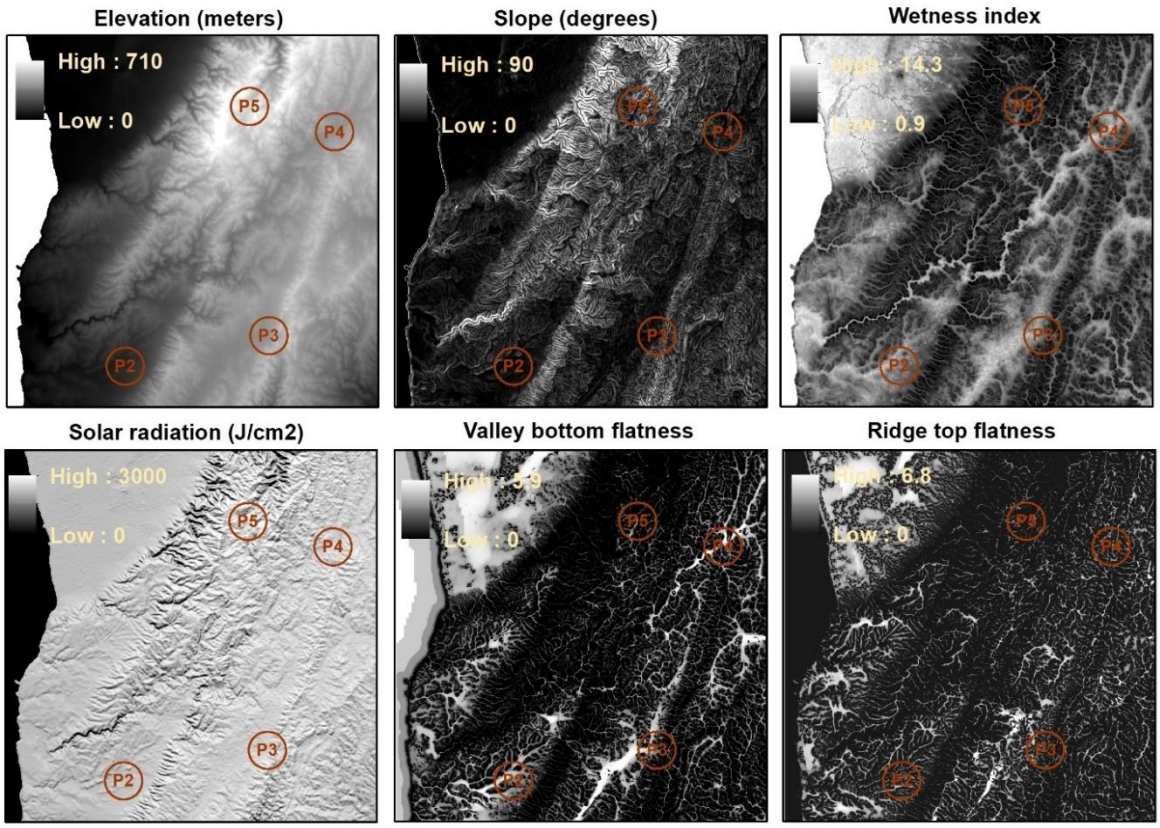

Figure 5-9 Maps of the terrain layers used for the Adelaide test region

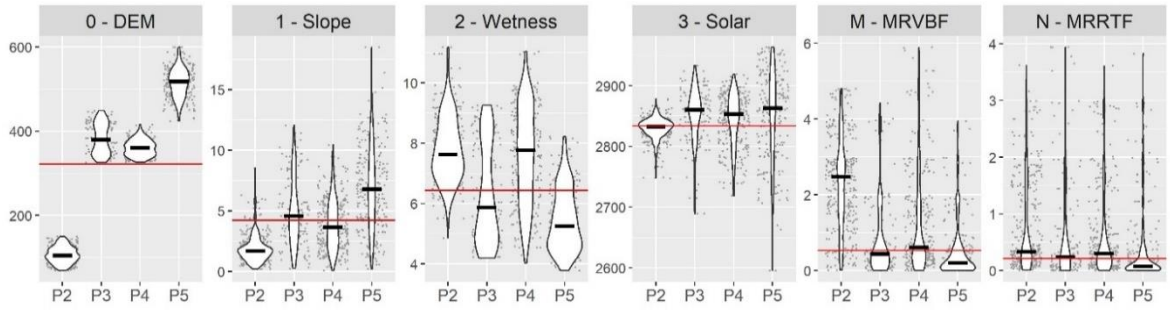

Figure 5-10 Distribution of terrain layer values at presence data sites (P2 McLaren Flat, P3 Meadows, P4 Oakbank and P5 Uraidla) compared to the mean for Adelaide test region (horizontal red line)

Site *P2 McLaren Flat*, near the coast, has very low elevation and is much flatter than the region as a whole, as evidenced by lower mean slope and substantially higher mean valley bottom flatness. In contrast, site *P5 Uraidla*, at the highest altitude, has the most rugged terrain (highest mean slope) and is driest (lowest mean wetness index value). The terrain values for sites *P3 Meadows* and *P4 Oakbank* show the most similarity to the test region averages, with the exception of wetness index which shows that *P3 Meadows* is drier and *P4 Oakbank* is wetter on average than the region overall.

Soil Layers

The preparation of soil layers for the two South Australian test sites duplicates the processing described in Section 2.5. Sets of hybrid soil layers were produced for both test sites from the SoilGrids rasters using a local map of soil type to substitute for the Gryphon Minerals soil map.

South Australian soil map

The polygon map of soil type used in the SA modelling identifies fifteen broad soil groups across South Australia's agricultural zone. The soil groups are based on significant profile features observed and recorded by soil scientists during the course of the State Land and Soil Mapping Program (1986-2001) and are described in the reference book *The Soils of Southern South Australia* by Hall, Maschmedt et al. (2009).[23] The maps of soil type for the Marrabel and Adelaide test regions are shown in Figure 5-11.

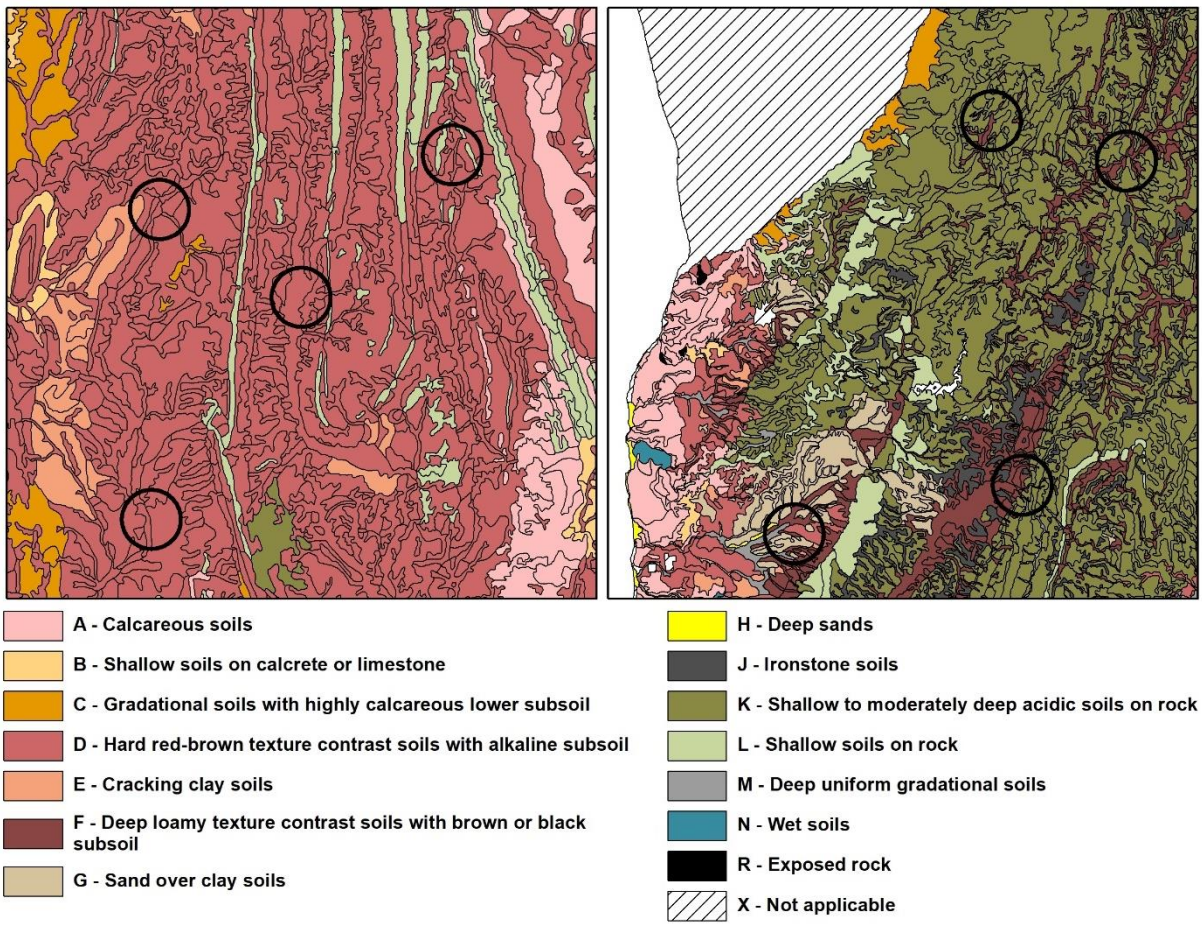

Figure 5-11 Soil maps for the Marrabel test region (left) and Adelaide test region (right)

SoilGrids

SoilGrids250m layers, resampled to 30m pixels and consolidated for the top 30cm of soil, were created for the two South Australian test sites using the process documented in Section 2.5.2. Maps of these consolidated layers are supplied in Appendix D, for reference.

[23] These data are available for free download from https://data.sa.gov.au/data/dataset/soil-groups. View fact sheet at https://data.environment.sa.gov.au/Content/Publications/SoilAttrib_FactSheet02_SoilGroups.pdf.

Hybrid soil layers

New sets of hybrid soil layers were created for each test site using the soil polygon boundaries from the South Australian soil map and the SoilGrids values for the respective soil types using the process described in Section 2.5.3 and illustrated in Figure 2-14.

Soil and Landscape Grid of Australia layers

In contrast to the Burkina Faso site, detailed information on soil attributes is available for Australian locations. The Soil and Landscape Grid of Australia (SLGA) provides raster maps for a range of soil attributes that were created by combining Australia-wide digital soil data and regional maps. The raster soil attribute data are provided at six defined depth levels and 3 arc second spatial resolution (approximately 90m x 90m pixels) and are accompanied by estimates of reliability for each pixel. [24]

The existence of these high-quality data allowed a formal comparison to be performed between the use of these data and the use of the hybrid soil layers in the Maxent modelling. Nine soil attributes that corresponded to the SoilGrids attributes used for the hybrid soil layers were selected, as shown in Table 5-8. The data were processed for use in the modelling by: (1) consolidating the data for the top 30cm (weighted average of values for depth levels 0 – 5cm, 5 – 15cm and 15 – 30cm); (2) clipping the consolidated rasters to the test regions; and (3) reprojecting each resulting raster to a UTM projected co-ordinate system with cell size of 30m x 30m using bilinear interpolation.

Table 5-8 Soil and Landscape Grid of Australia attributes used to create environmental layers

Layer	Attribute	Code	Description	Units
A - BLDFIE	Bulk Density (whole earth)	BDw	Bulk Density of the whole soil (including coarse fragments) in mass per unit volume	g/cm3
B - OCDENS	Organic Carbon	SOC	Mass fraction of carbon by weight in the less than 2mm soil material	%
C - CECSOL	Effective Cation Exchange Capacity	CEC	Cations extracted using barium chloride plus exchangeable H + Al	meq/100g
D - PHIHOX	pH (CaCl2)	pHc	pH of 1:5 soil/0.01M calcium chloride extract	none
E - CLYPPT	Clay	CLY	Less than 2um mass fraction of the less than 2mm soil material	%
F - CRFVOL	Coarse Fragments	CFG	Mass fraction of the soil material greater than 2mm	%
G - SNDPPT	Sand	SND	20um – 2 mm mass fraction of the less than 2mm soil material	%
H - SLTPPT	Silt	SLT	2-20 um mass fraction of the less than 2mm soil material	%
I - WWP	Available Water Capacity	AWC	Available water capacity	%

[24] Refer https://www.clw.csiro.au/aclep/soilandlandscapegrid/ for data download and product descriptions.

Radiometric layers

The radiometric map for South Australia used in this project was produced by the Department of Primary Industries and Resources, South Australia in 2011 to aid geological exploration.[25] The data were formatted as a three-band raster using the GDA94 coordinate system and cell size of approximately 100m. The raster band values were unsigned integers scaled to the range 0-255.

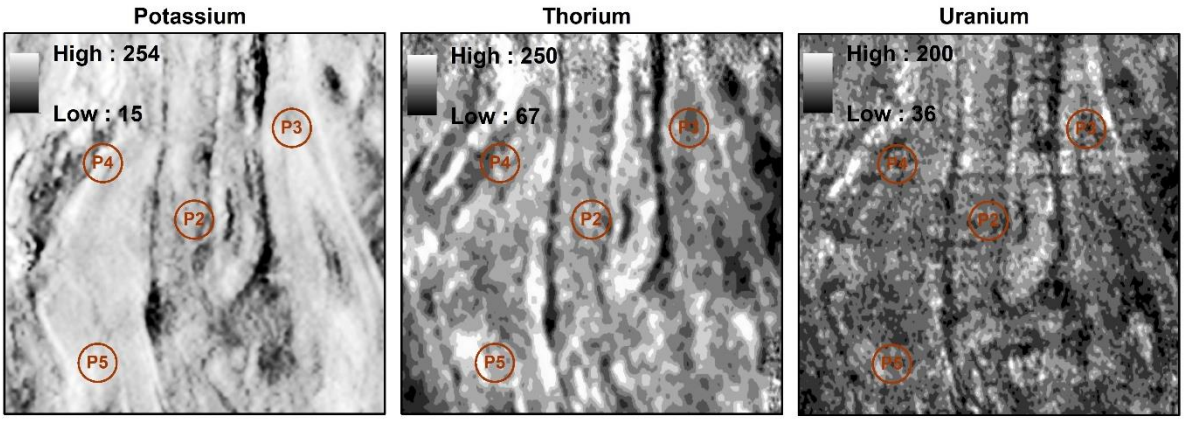

Figure 5-12 Maps showing the digital values for radiometric layers at the Marrabel test region

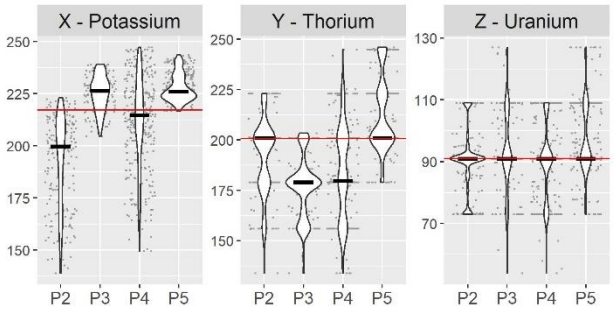

Figure 5-13 Distribution of radiometric layer values at presence data sites (P2 Marrabel, P3 Ngapala, P4 Saddleworth and P5 Tarlee) compared to the mean for Marrabel test region (horizontal red line)

The raster bands were each reprojected to 30m cell size using bilinear interpolation and clipped for the two SA test regions. The resulting radiometric layers for the Marrabel test region are shown in Figure 5-12 and the distributions of the layer values are plotted in Figure 5-13. Note that radiometric environmental layers were not produced for the Adelaide test region as radiometric data were not available for the western part of the region.

[25] The radiometric grid of Australia is available for free download from Geoscience Australia, refer https://ecat.ga.gov.au/geonetwork/srv/eng/catalog.search#/metadata/131974.

5.4.4 Verification Maps

The verification maps for the crops in each test region are raster maps derived from the *Land use potential* shapefiles described in Section 5.3. Suitability ratings were assigned to each of the *Land use potential* polygons based on the mapping category code as per Table 5-9. The shapefiles were converted to rasters using the suitability rating for the raster value to produce sets of crop suitability rasters with values in the range 0 to 1. Note that the suitability rasters so created, with equivalent symbology, are visually identical to the crop maps in Figure 5-3 and Figure 5-4.

Table 5-9 Suitability ratings assigned to Land use potential codes

	Suitable				Unsuitable					
Land use potential code	Aa	Ab	Ac	Ad	B	C	D	Ea	Eb	X
Suitability rating	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0	NA

The *Land use potential* codes Ad, Ac, Ab and Aa denote areas of moderate to high potential for a crop, and so indicate areas suitable for growing that crop. Classifying these four codes as *suitable* and all other codes as *unsuitable* for a crop allows the fractional predicted area (FPA) for the crop to be calculated from its verification map.

Table 5-10 shows the fractional predicted areas for the six crops modelled at the Marrabel test region. Most areas of the test region are suitable for most crops – Chickpeas have the smallest FPA at 71% and all other crops have FPAs of over 80%. The fractional predicted area for a crop affects the achievable AUC score for models (refer Section 3.4.2 and Table 3-4). In the case of Lucerne, with FPA of 96%, a test AUC score of 0.52 could be expected for a very good classifier. The maximum theoretical AUC is less than 0.6 for all crops except Chickpeas (0.65).

Table 5-10 Fractional predicted area and corresponding theoretical maximum achievable AUC for crops in the Marrabel test region

	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat
Fractional predicted area	0.87	0.82	0.71	0.96	0.88	0.82
Theoretical maximum AUC	0.57	0.59	0.65	0.52	0.56	0.59

Table 5-11 shows the fractional predicted areas for the six crops modelled at the Adelaide test region. There is much greater variation in FPA for the crops in this test region, ranging from 25% for Faba beans up to 72% for Olives and Grape vines. The smaller FPAs result in higher achievable AUC scores for good models for these crops in this test region.

Table 5-11 Fractional predicted area and corresponding theoretical maximum achievable AUC for crops in the Adelaide test region

	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes
Fractional predicted area	0.49	0.48	0.25	0.72	0.72	0.28
Theoretical maximum AUC	0.76	0.76	0.88	0.64	0.64	0.86

Binary suitability maps

The *Land use potential* raster verification maps described above can be converted to binary verification maps by mapping the raster cell values to 0 (unsuitable) if they are less than 0.5 or to 1 (suitable) otherwise (see Figure 5-14).

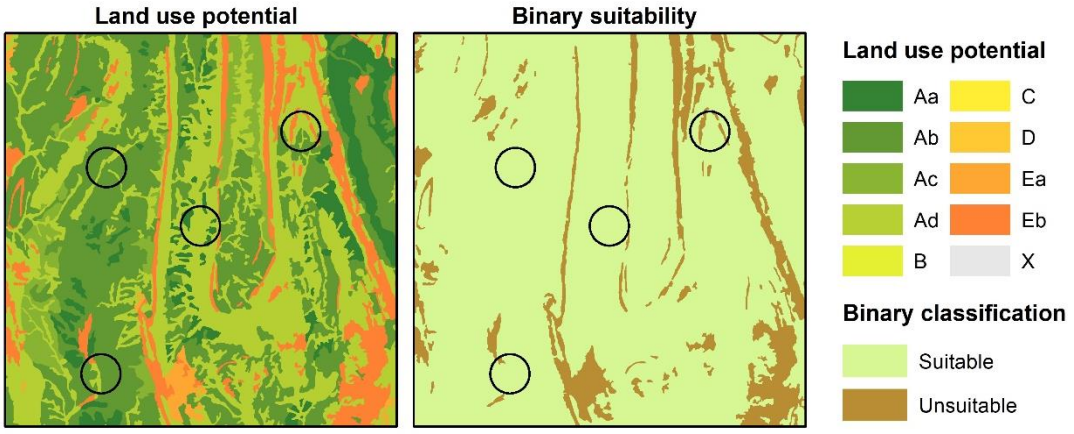

Figure 5-14 Corresponding maps of Land use potential and binary suitability for Barley at the Marrabel test region; Land use potential codes Aa, Ab, Ac and Ad are classified as suitable

Maxent suitability maps can be converted to binary classification maps with known FPA by using the fixed cumulative value of 1-FPA to threshold predictions. The known FPAs for the six crops in each region (Table 5-10 and Table 5-11) are used to specify the thresholds for the Maxent predictions so that each binary classification map has the same FPA as its corresponding binary verification map, see Figure 5-15. The two binary maps may then be formally compared to measure the amount of agreement between the binary classifications (accuracy) and the degree to which this agreement is due to chance (kappa statistic).

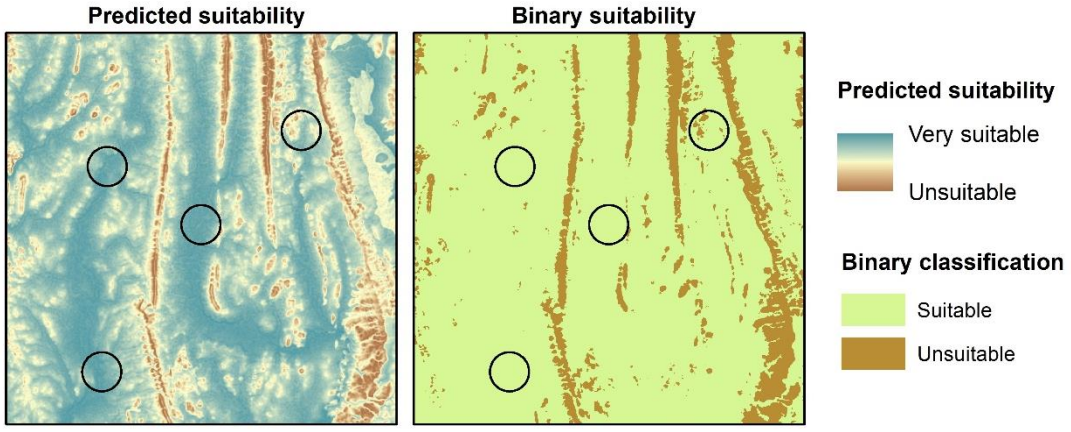

Figure 5-15 Corresponding maps of continuous and binary predicted suitability for Barley at the Marrabel test region; the area classified suitable is equal to fractional predicted area for Barley

Interrater comparison of binary suitability maps

The binary verification maps derived from the *Land use potential* maps provide the region-wide validation data used in this chapter for the Maxent predictions. These six binary suitability maps for each region provide benchmarks against which the crop suitability predictions generated by Maxent can be formally compared. Each Maxent suitability map is first transformed to a binary classification map with the same FPA as its corresponding binary verification map. Interrater

comparison of the two binary maps will calculate the accuracy and kappa values measuring the agreement between the prediction and benchmark maps.

It is not possible to interpret these numbers as measurements for the correctness of predictions. The *Land use potential* maps have not been subject to field validation and have some limitations in relation to spatial accuracy. Their soil landscape map units are not homogeneous and, as can be seen from Table 5-2, soil polygons rated as suitable for a crop may include up to 40% unsuitable land. In addition, the boundaries between the mapping units represent transitions between suitability categories and are rarely as sharp in the landscape as the maps imply (Rowland, Maschmedt et al. 2016). This is illustrated in Figure 5-16 which shows corresponding map detail for *Land use potential* and modelled suitability predictions for Wheat, overlaid with soil polygon boundaries and elevation contours.

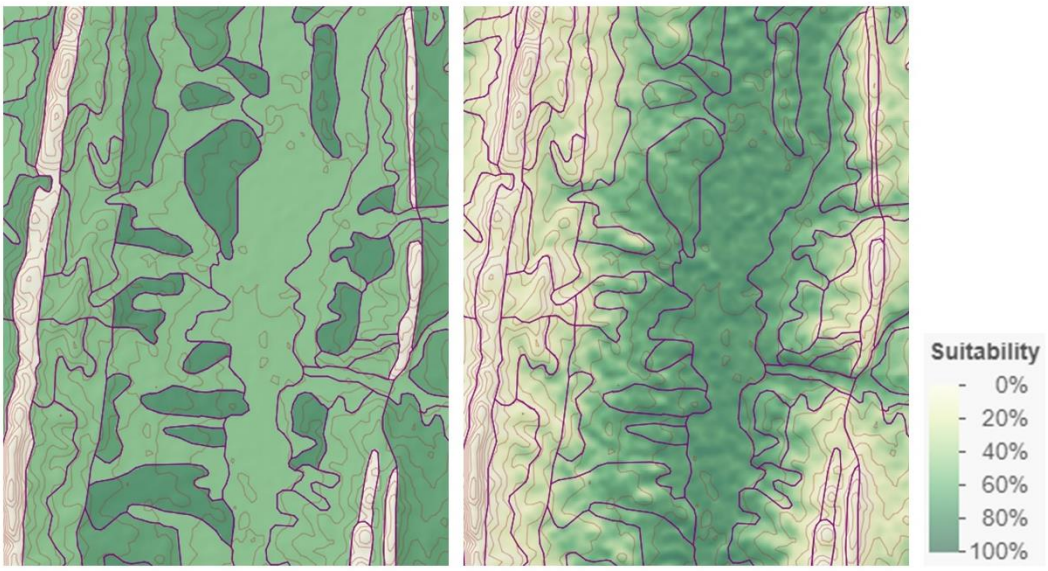

Figure 5-16 Comparison of validation map (left) and suitability predictions (right) – soil polygons outlined in purple, 10m contour lines shown in brown

Never-the-less, the binary verification maps derived from *Land use potential* maps provide a consistent and soundly based standard against which to compare the predictions from alternative suitability models. In the results presented below the accuracy and kappa values from interrater comparisons with the binary verification maps are used as performance measures to compare and rate the effectiveness of the crop suitability models.

5.4.5 Outputs

The same outputs as those listed for the Chapter 4 processing were produced for all runs for the South Australian test regions. Additional outputs were generated to facilitate region-wide validation of suitability models. These were:

- Binary suitability maps created by thresholding the continuously valued suitability maps using the threshold value that results in the same fractional predicted area as the corresponding binary verification map;
- The accuracy and kappa values resulting from the interrater comparison of the binary suitability maps and their corresponding binary verification maps.

5.5 Results and discussion

The results from the crop suitability modelling for the Marrabel and Adelaide test regions are presented below. Model robustness and sensitivity to sampling bias is assessed by comparing the accuracy of predictions on test data (AUC) and the similarity of suitability maps (correlations) across the five model scenarios, as was done in Chapter 4. However, model accuracy is now formally measured using the accuracy and kappa values resulting from the interrater comparison of binary suitability maps and their corresponding binary verification maps.

The crop suitability models that correspond to the "*final suitability maps*" from Section 4.5 are examined first to assess the effectiveness of the set of predictors and parameters used for the West African models for modelling at dissimilar sites. Results for models built using other combinations of predictors and parameters are then compared to assess whether the conclusions drawn in Chapter 4 also apply to these results.

5.5.1 Marrabel test region

The Marrabel test region has a similar terrain profile to the Burkina Faso project region, being relatively flat (mean slope < 2.5 degrees) but with some areas of steep slope and low wetness values that would seem unsuitable for agriculture. The maps of *Land use potential* (Figure 5-3) show that most of the Marrabel test region is cultivatable and rated suitable for the crops that have been modelled.

Corresponding "final suitability maps"

Figure 5-17 presents the crop suitability maps for the Marrabel test region (model scenario *P1 Random 25%*) that were generated using the same environmental layers and parameters as the final models presented in Chapter 4. The predictors for these models were three terrain layers (slope, wetness index and solar radiation) and the nine hybrid soil layers created by adding Gaussian noise to the SoilGrids mean values by soil type. Linear, quadratic and product feature classes and the default beta multiplier of one were used for the models. The (smoothed) continuous predictions generated by Maxent are shown in the top row for comparison with raster verification maps created using the suitability values from Table 5-9. Below these are the binary prediction and verification maps created by thresholding each crop suitability map based on the known fractional predicted area for the crop across the test region.

The maps of *continuous* predictions show similar overall patterns of suitability and unsuitability to the verification maps, but vary in places in the degree of suitability predicted. Predictions of high suitability are observable in a number of areas rated as moderately suitable in the verification maps, and vice versa. However, it should be noted that the continuously valued suitability maps are not directly comparable with the discrete-valued verification maps that have just nine possible suitability scores.

The *binary* verification maps are directly comparable to binary suitability maps for the same crop thresholded to the same fractional predicted area. Very similar patterns of suitability are observable between the binary prediction maps and their corresponding binary verification maps.

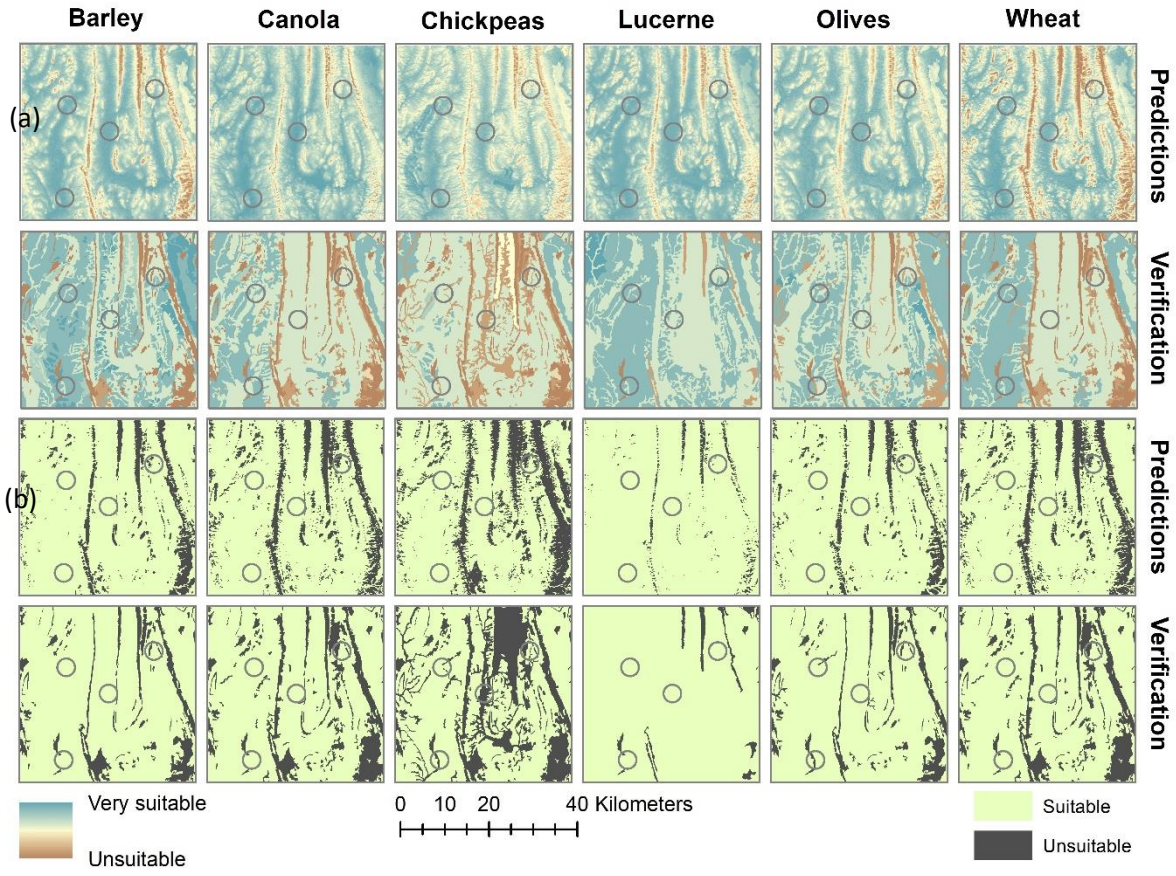

Figure 5-17 Visual comparison of Maxent predictions with verification data for Marrabel test region: (a) continuous predictions and (b) binary predictions based on fractional predicted areas for crops; models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI)

Formal measures of accuracy for the predictions are provided in Table 5-12, both for the original Maxent predictions and after smoothing of the suitability maps. These measures are: the threshold-independent test AUC scores from the continuously valued suitability maps; and the accuracy and kappa values from the interrater comparison of the thresholded binary suitability maps with their corresponding binary verification maps.

The test AUC scores for the original Maxent predictions exceed the theoretical maximum AUC scores for all crops, and the maps have good measured accuracy with good kappa values indicating non-random predictive performance. The mean measured accuracy of the six binary maps from the Maxent predictions is 0.86 with mean kappa value of 0.44. Smoothing the maps improved mean measured accuracy slightly to 0.87 with mean kappa of 0.47. Smoothing produced the greatest improvement in binary accuracy and kappa values for Chickpeas (from accuracy of 0.70, κ=0.29 to accuracy of 0.75, κ=0.36), but also slightly reduced the Chickpeas test AUC score.

The thresholded smoothed Maxent suitability maps for all crops except Chickpeas had accuracy greater than 85% and kappa values above 0.4. For Lucerne, with very high FPA of 96%, the predictions were 96% accurate when compared to the verification map, but the kappa value of 0.4 demonstrated that this was not random agreement.

Table 5-12 Accuracy of crop suitability models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI) – Marrabel test region, model scenario = P1 Random 25%

	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Mean
Fractional predicted area (FPA)	*87%*	*82%*	*71%*	*96%*	*88%*	*82%*	-
Theoretical maximum AUC	*0.57*	*0.59*	*0.65*	*0.52*	*0.56*	*0.59*	*0.58*
	Maxent suitability maps						
Test AUC	0.75	0.78	0.69	0.72	0.76	0.71	0.73
Binary accuracy (same FPA)	0.90	0.85	0.70	0.96	0.88	0.85	0.86
Kappa for binary accuracy	0.55	0.49	0.29	0.38	0.43	0.52	0.44
	Smoothed Maxent suitability maps						
Test AUC	0.76	0.80	0.67	0.75	0.75	0.73	0.74
Binary accuracy (same FPA)	0.91	0.86	0.75	0.96	0.89	0.86	0.87
Kappa for binary accuracy	0.56	0.51	0.36	0.40	0.44	0.52	0.47

Comparing model scenarios – sample selection bias

The maps presented in Figure 5-17 show the predictions of the crop suitability models from scenario *P1 Random 25%.* These models were trained using presence data from all four locations and so are likely to be the best set of models from the five model scenarios. The four other models for each crop in this suite of similar models were developed using the same set of environmental predictors and Maxent parameters, but used spatially independent training and test data. The robustness of the final models is assessed by comparing results from all models in the suite. Note: the tables and figures below compare the results for the *smoothed* Maxent suitability maps for the six crops in the five model scenarios.

Table 5-13 shows the test AUC scores for the 30 crop models in the suite. Most crop models achieved test AUC scores higher than the theoretical maximum AUC for the crop. The *P1 Random 25%* test AUC scores were highest for most crops; however, the models for Chickpeas and Olives in scenario *P4 Saddleworth* both achieved the highest test AUC score for their respective crop, despite having no training samples from the test site.

Table 5-13 Test AUC scores for crops in the five model scenarios from models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI) – Marrabel test region

Model scenario	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Mean
Theoretical max AUC	*0.57*	*0.59*	*0.65*	*0.52*	*0.56*	*0.59*	*0.58*
P1 Random 25%	0.76	0.80	0.67	0.75	0.75	0.73	0.74
P2 Marrabel	0.77	0.78	0.51	0.73	0.65	0.70	0.69
P3 Ngapala	0.45	0.35	0.42	0.42	0.48	0.42	0.42
P4 Saddleworth	0.71	0.74	0.74	0.69	0.78	0.71	0.73
P5 Tarlee	0.73	0.74	0.64	0.69	0.76	0.71	0.71
Mean test AUC	**0.68**	**0.68**	**0.60**	**0.65**	**0.68**	**0.65**	**0.66**

All models from model scenario *P3 Ngapala* had very poor test AUC scores. These models were trained using presence data from the other three sites but tested using presence data from Ngapala alone. The test site at Ngapala is the driest and most rugged of the four sites. Failing to include representative training samples from Ngapala has reduced the predictive capability of these models on test data from this site.

Table 5-14 presents the accuracy and kappa values for interrater comparisons of the binary suitability maps with corresponding binary verification maps for the six crops in the five model scenarios. The mean accuracy across all crop models is 0.87 with mean kappa value of 0.46.

Comparing the results for the individual crops it is observed that the Lucerne models have highest accuracy (0.96) and the models for Barley have the highest kappa values (mean 0.55) in all model scenarios. The models for Chickpeas have the lowest accuracy of all the crop models in every model scenario.

Table 5-14 Accuracy of thresholded maps from the five model scenarios – models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI) – Marrabel test region

Model scenario	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Mean
FPA	*87%*	*82%*	*71%*	*96%*	*88%*	*82%*	-
P1 - Random 25%	0.91	0.86	0.75	0.96	0.89	0.86	0.87
Kappa	*0.56*	*0.51*	*0.36*	*0.40*	*0.44*	*0.52*	*0.47*
P2 - Marrabel	0.90	0.85	0.73	0.96	0.90	0.85	0.87
Kappa	*0.55*	*0.48*	*0.34*	*0.41*	*0.48*	*0.50*	*0.46*
P3 - Ngapala	0.89	0.86	0.72	0.96	0.88	0.85	0.86
Kappa	*0.51*	*0.50*	*0.30*	*0.36*	*0.44*	*0.50*	*0.43*
P4 - Saddleworth	0.91	0.85	0.74	0.96	0.89	0.85	0.87
Kappa	*0.57*	*0.45*	*0.33*	*0.39*	*0.43*	*0.48*	*0.44*
P5 - Tarlee	0.91	0.87	0.77	0.96	0.89	0.87	0.88
Kappa	*0.56*	*0.54*	*0.40*	*0.37*	*0.44*	*0.55*	*0.48*
Mean accuracy	**0.90**	**0.86**	**0.74**	**0.96**	**0.89**	**0.86**	**0.87**
Mean kappa	***0.55***	***0.50***	***0.35***	***0.39***	***0.45***	***0.51***	***0.46***

Comparing models for the same crop across all model scenarios reveals similar calculated accuracy and kappa values, with no model scenario resulting in strikingly different results for any of the crops. Overall, the binary suitability maps from scenario *P5 Tarlee* showed greatest agreement with the verification maps, but the mean accuracy and kappa values for this model scenario were only very slightly higher than those from the other four model scenarios.

Figure 5-18 presents the scatter plots and correlation coefficients for pairwise comparisons of suitability maps from the five model scenarios for each crop to illustrate the degrees of similarity between the crop models. The suitability maps show high levels of similarity: the mean correlation coefficient overall is 0.92, and all results had high significance level (p-value < 0.001).

The crop suitability maps generated in the four model scenarios P1, P2, P3 and P4 show very high similarity, having correlation coefficients between 0.9 and 1 for all pairwise comparisons other than with the *P2 Marrabel* map for Olives (0.88). Despite very poor test AUC scores from the spatially independent test data in scenario *P3 Ngapala*, the suitability maps produced by these models are extremely similar to the *P1 Random 25%* maps, with correlation coefficients greater than 0.95 for all crops and perfect correlation for Wheat (1.00).

The continuous-valued suitability maps from the *P5 Tarlee* models (that had no training points from Tarlee) show the least similarity to the maps from other models (although their binary thresholded versions have already been shown to have very similar accuracy). The Tarlee site differs in that it receives less solar radiation on average than the other three test sites and the region as a whole (refer Figure 5-8).

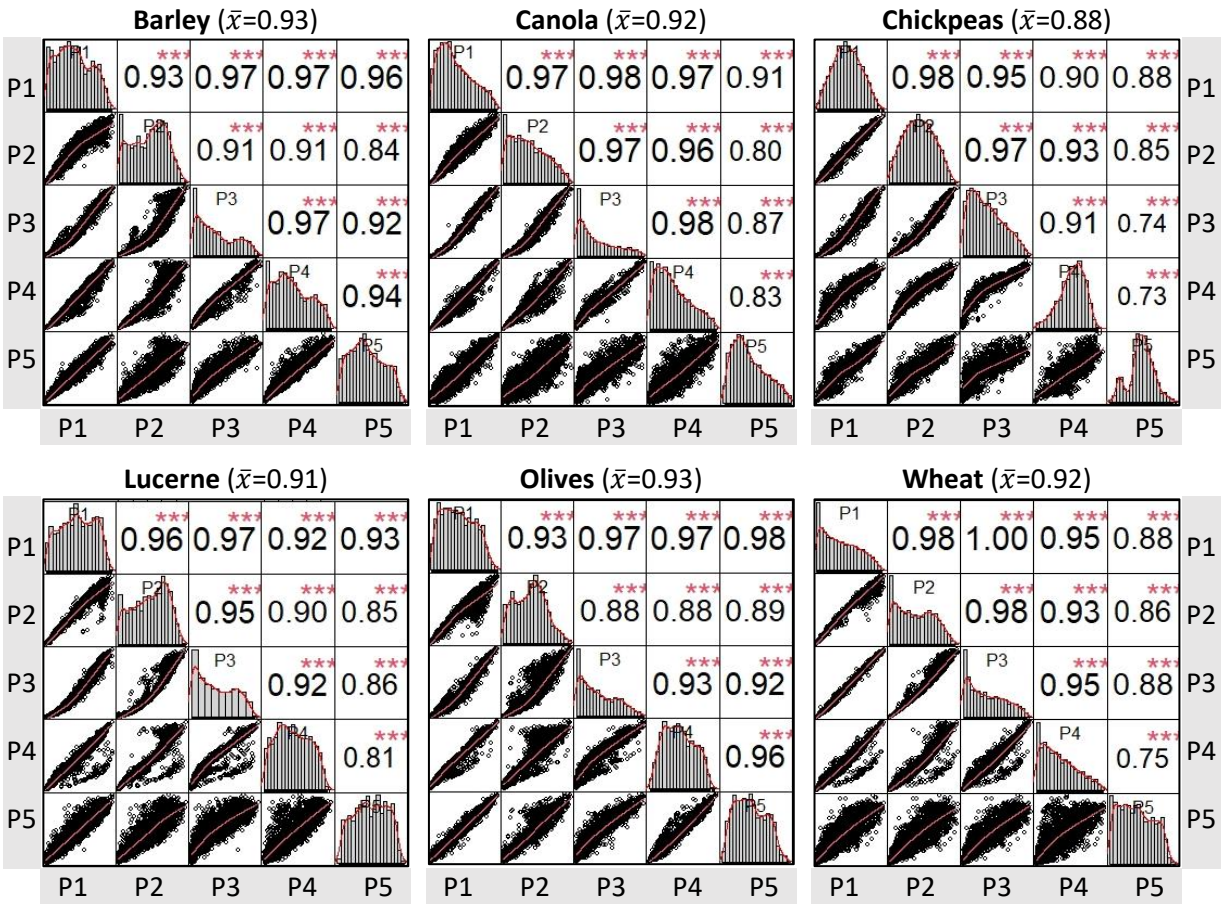

Figure 5-18 Correlations of suitability maps between the five model scenarios from models built using terrain layers (123) and hybrid soil layers with added Gaussian noise (ABCDEFGHI) – Marrabel test region (mean = 0.92)

The environmental layer contributions for the models are presented in Figure 5-19. The histograms illustrate model reliance on particular layers and show how including or excluding a site from the set of training data may influence the makeup of the crop models. For example, layer 3=Solar radiation makes a much smaller contribution to the *P5 Tarlee* models than to the crop models in the other model scenarios.

Layer 2=Wetness index is the most significant environmental predictor for most models. The patterns of predicted suitability in the maps of continuously-valued predictions in Figure 5-17(a) match closely the patterns in the Wetness index map in Figure 5-7, with areas of higher wetness corresponding to areas of higher predicted suitability. When resolved into maps of binary predictions, these patterns of suitability, for most crops, match closely the patterns of suitability in the binary verification maps and result in high accuracy and kappa values from the interrater comparisons.

The exception is Chickpeas. Compared to the other five crops, the Chickpeas models had the lowest test AUC scores and the least agreement with the verification maps. Close inspection of the maps for Chickpeas in Figure 5-17 reveals contradictory predictions of suitability for Chickpeas down centre of test region in an area that has high values for wetness index and valley bottom flatness (compare Figure 5-7). The suitability maps predicted high suitability for Chickpeas in this area in all model scenarios (only the maps for *P1 Random 25%* are shown) whereas the

verification maps derived from *Land use potentials* classify this area as unsuitable. Chickpeas are sensitive to waterlogging and prefer well-drained, non-acidic soils with medium to heavy clay texture (Pulse Australia 2016). Field validation would be required to establish whether waterlogging is a risk in this area.

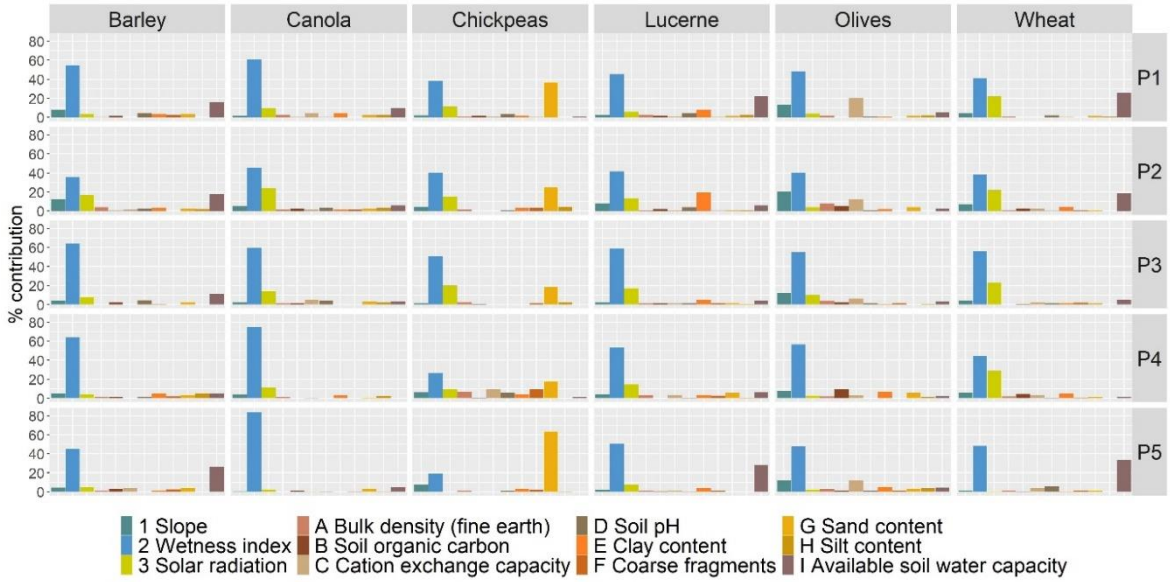

Figure 5-19 Environment contributions (%) to suitability models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise – Marrabel test region

Comparing alternative models

All the models discussed above were generated using the same environmental layers and parameters as the final models presented in Chapter 4. The results for these models have been examined in detail to assess how well they predict suitability for the crops across the test region as a whole. Other suites of suitability models using different sets of environmental predictors and Maxent parameters were generated for the Marrabel test region as part of this research to test conclusions reached in Chapter 4. Summary results from these models are presented below.

Table 5-15 presents summary performance scores for suites of models built to test different combinations of environmental layers. In Chapter 4, altitude (layer 0=DEM) was shown to be a poor predictor for suitability for the Burkina Faso project region. This is also the case for the Marrabel test region. The suite of models built using all 18 environmental layers performed very poorly, having mean test AUC of 0.4 (worse than random prediction) and the lowest scores for map similarity and accuracy from all the tests reported in the table. Removing layer 0=DEM as a predictor substantially improved mean test AUC to 0.64 (better than the mean of the theoretical maximum AUCs for the crops) and improved all other summary performance scores.

The models built using only the three terrain layers 1=Slope, 2=Wetness index and 3=Solar radiation had the same mean test AUC and were almost as accurate as the more complex models built using 17 layers (without 0=DEM). Including additional terrain predictors M=MRVBF and N=MRRTF increased test AUC slightly but did not improve accuracy. The further addition of radiometric layers as predictors degraded all performance scores. This result is consistent with the findings from Chapter 4 that the use of radiometric layers and the terrain layers 0=DEM, M=MRVBF and N=MRRTF did not lead to increased predictive capability. The models

corresponding to the final suitability maps from Chapter 4 (123ABCDEFGHI) produced the best results in Table 5-15, achieving the highest scores for all performance measures.

Table 5-15 Summary performance scores for models built using different combinations of environmental layers; means calculated from model scores for all crops in all model scenarios

Test description	Environmental layers	Feature classes	Mean test AUC	Mean correlation coefficient*	Mean accuracy	Mean kappa
All layers	0123ABCDEFGHIMNXYZ#	LQPH	0.40	0.63	0.81	0.29
Without DEM	123ABCDEFGHIMNXYZ#	LQPH	0.64	0.82	0.86	0.43
Terrain + soil	*123ABCDEFGHI#*	*LQP*	*0.66*	*0.92*	*0.87*	*0.46*
Terrain only	123	LQPH	0.64	0.91	0.84	0.39
Terrain only	123MN	LQP	0.66	0.90	0.84	0.39
+ Radiometric	123MNXYZ	LQP	0.61	0.81	0.82	0.31

* Mean of correlation coefficients from pairwise comparisons of suitability maps for the same crop in all model scenarios

Soil layers are the mean SoilGrids values for soil types with added Gaussian noise (Means + jitter hybrid soils)

Hybrid soils layers

Table 5-16 compares results for suites of models built using the sets of hybrid soil layers created using each of the methods described in Section 4.3.5.

Table 5-16 Summary performance scores for models built from terrain and soil layers (123ABCDEFGHI) using sets of hybrid soil layers constructed using different methods; means calculated from model scores for all crops in all model scenarios

Soil layers = ABCDEFGHI PCA soil layers = PQ Method	Feature classes	Mean test AUC	Mean correlation coefficient*	Mean accuracy	Mean kappa
Means + jitter	LQP	0.64	0.88	0.85	0.43
Means + jitter: smoothed maps	*LQP*	*0.66*	*0.92*	*0.87*	*0.46*
PCA of Means + jitter	LQP	0.66	0.90	0.86	0.32
Means (50%) + SG (50%)	LQP	0.59	0.69	0.83	0.37
PCA of Means (50%) + SG (50%)	LQP	0.67	0.84	0.82	0.36
Means (75%) + SG (25%)	LQP	0.60	0.73	0.86	0.45
PCA of Means (75%) + SG (25%)	LQP	0.65	0.81	0.84	0.37

* Mean of correlation coefficients from pairwise comparisons of suitability maps for the same crop in all model scenarios

Similar to the results reported in Chapter 4, the smoothed suitability maps from the models using the hybrid soil layers with added Gaussian noise (*Means + jitter: smoothed maps*) achieved the highest mean map correlation score measuring similarity of predictions across the five model scenarios for each crop. These suitability maps also achieved the highest mean accuracy and kappa scores of all the models using hybrid soil layers.

It is noted that the very high test AUC values observed for some suites of models using hybrid soil layers in Chapter 4 are not apparent in the Marrabel results reported above. The West African presence points are true occurrences, whereas the South Australian presence points are samples

from locations predicted likely to be suitable for their crops. This difference may have mitigated the potential for overfitting to training data in the South Australian models.

Comparison with Soil and Landscape Grid of Australia soil layers

The lack of suitable available soil data for modelling agricultural suitability in the Burkina Faso project region necessitated the development of the hybrid soil layers that combined SoilGrids raster layers with a local polygon soil map. As the purpose of this chapter was to validate the methods used for the West African site, similar hybrid soil layers were produced for the two South Australian test regions and the results for Maxent crop suitability models using the hybrid soil layers as environmental predictors have been reported above.

High-quality soil attribute data are available for Australian sites from the Soil and Landscape Grid of Australia (SLGA). The predictive performance of crop suitability models built using soil layers derived from these data is compared to models build using hybrid soil layers for the 'best' models (scenario *P1 Random 25%*) in Table 5-17, and across all models scenarios in Table 5-18.

Table 5-17 Comparative performance of models using hybrid soil layers and SLGA soil layers as predictors for the 'best' models (models from scenario P1 built from terrain and soil layers 123ABCDEFGHI)

	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Mean
Fractional predicted area (FPA)	*87%*	*82%*	*71%*	*96%*	*88%*	*82%*	*-*
Theoretical maximum AUC	*0.57*	*0.59*	*0.65*	*0.52*	*0.56*	*0.59*	*0.58*
	Hybrid soil layers with added Gaussian noise						
Test AUC	0.76	0.80	0.67	0.75	0.75	0.73	0.74
Binary accuracy	0.91	0.86	0.75	0.96	0.89	0.86	0.87
Kappa for binary accuracy	0.56	0.51	0.36	0.40	0.44	0.52	0.47
	Soil and Landscape Grid of Australia soil layers						
Test AUC	0.81	0.83	0.81	0.82	0.86	0.80	0.82
Binary accuracy	0.89	0.88	0.70	0.94	0.89	0.87	0.86
Kappa for binary accuracy	0.55	0.59	0.29	0.21	0.49	0.57	0.45

The measured accuracy against the validation data was similar for both sets of models for most crops, although higher accuracy scores for Chickpeas and higher kappa score for Lucerne are observed for the models using the hybrid soil layers.

More similarity of predictions across the five model scenarios was observed for the suitability maps from the models using the hybrid soil layers than for the models using SLGA soil layers (mean correlation coefficients of 0.92 and 0.81, respectively). It is observed that the models using SLGA soil layers had much higher test AUC values for all crops in the 'best' models, but lower *mean* test AUC values across the five model scenarios. This indicates greater spatial bias occurring in the models using SLGA soil layers and that the models using hybrid soil layers are more robust against sample selection bias.

Table 5-18 Summary performance scores for models using hybrid soil layers and SLGA soil layers as predictors - models built from terrain and soil layers 123ABCDEFGHI, means calculated from model scores in all model scenarios

	Barley	Canola	Chickpeas	Lucerne	Olives	Wheat	Mean
Fractional predicted area (FPA)	*87%*	*82%*	*71%*	*96%*	*88%*	*82%*	-
Theoretical maximum AUC	*0.57*	*0.59*	*0.65*	*0.52*	*0.56*	*0.59*	*0.58*
	Hybrid soil layers with added Gaussian noise						
Mean test AUC	0.68	0.68	0.60	0.65	0.68	0.65	0.66
Mean map correlations (*r*)	0.93	0.92	0.88	0.91	0.93	0.92	0.92
Mean binary accuracy	0.90	0.86	0.74	0.96	0.89	0.86	0.87
Mean kappa	0.55	0.50	0.35	0.39	0.45	0.51	0.46
	Soil and Landscape Grid of Australia soil layers						
Mean test AUC	0.61	0.67	0.58	0.58	0.68	0.63	0.63
Mean map correlations (*r*)	0.79	0.81	0.74	0.80	0.83	0.85	0.81
Mean binary accuracy	0.89	0.87	0.70	0.94	0.89	0.87	0.86
Mean kappa	0.54	0.58	0.27	0.20	0.49	0.57	0.44

Conclusions

The results presented above demonstrate that the methods of predicting agricultural land suitability for crops presented in Chapter 4 can be used effectively at dissimilar locations. The final maps of Maxent predictions for the six crops in the Marrabel test region (produced identically to the final maps from Chapter 4) were realistic and had high measured accuracy when compared to region-wide validation maps. The cross-validation results also showed the models to be robust against sample selection bias from the training data.

The hybrid soil layers with added Gaussian noise (synthesised from a detailed categorical map and coarse resolution raster layers) were effective environmental predictors for the Marrabel crop models. Their use resulted in model predictions that were stable across cross-validation scenarios with accuracy that was similar to that of models that used soil layers from high-quality soil property rasters (SLGA data). The other methods for creating hybrid soil layers that were tested all produced unreliable results that were both less accurate and more unstable.

Similarly to the models in Chapter 4, the environmental predictors contributing to the best Marrabel crop models were slope, wetness index, solar radiation and the nine hybrid soil property layers with added Gaussian noise; and use of further terrain and radiometric layers was shown not to improve predictive capability.

5.5.2 Adelaide test region

The terrain maps for the Adelaide test region show that it contains many areas with steep slope and low wetness values that would seem unsuitable for agriculture (refer Figure 5-9). The maps of *Land use potential* for the six crops modelled (Figure 5-4) show much variation in areas of predicted suitability for the different crops. Grape vines and Olives have the largest ecological niches with almost three quarters of the region rated suitable for each crop; for Apples and Brassicas about half the region is predicted suitable; and for Faba beans and Potatoes only about a quarter of the region is predicted suitable.

Corresponding "final suitability maps"

Figure 5-20 presents the crop suitability maps for the Adelaide test region (model scenario *P1 Random 25%*) that were generated using the same environmental layers and parameters as the final models presented in Chapter 4, and compares these to the verification maps of land use potential.

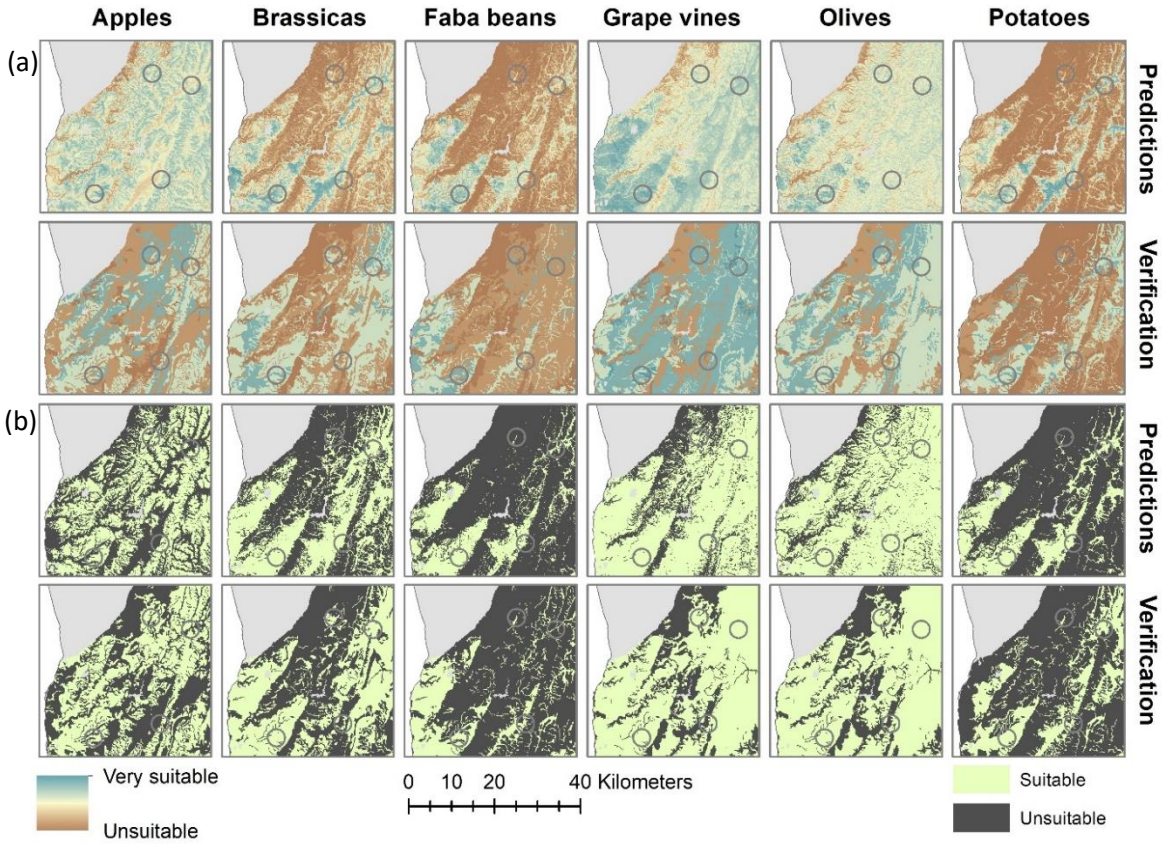

Figure 5-20 Visual comparison of Maxent predictions with verification data for Adelaide test region: (a) continuous predictions; and (b) binary predictions based on fractional predicted areas for crops. Models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise.

The thresholded maps of predictions for each crop show similar patterns of predicted suitability to those in the corresponding binary verification maps, although the degrees of predicted suitability in the continuously valued maps vary in places. Formal measures of accuracy for the binary predictions are provided in Table 5-19.

The highest accuracy scores occurred for the two crops with the smallest ecological niches – Faba beans and Potatoes. The theoretical maximum AUC scores for these crops were the highest (0.88 and 0.86, respectively) and the models for both crops achieved test AUC scores very close to these values (0.88 and 0.84, respectively). Both crop models had binary accuracy scores greater than 80% with corresponding kappa scores greater than 0.5. In contrast, the model for Olives had the lowest test AUC score (0.57), and the model for Apples had the lowest binary accuracy score (68%, κ=0.36).

Table 5-19 Accuracy of crop suitability models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise – Adelaide test region, model scenario P1

	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes	Mean
Fractional predicted area (FPA)	*49%*	*48%*	*25%*	*72%*	*72%*	*28%*	-
Theoretical maximum AUC	*0.76*	*0.76*	*0.88*	*0.64*	*0.64*	*0.86*	*0.76*
	Smoothed Maxent suitability maps						
Test AUC	0.59	0.70	0.88	0.61	0.57	0.84	0.70
Binary accuracy (FPA)	0.68	0.74	0.84	0.79	0.76	0.81	0.77
Kappa for binary accuracy	0.36	0.49	0.56	0.46	0.32	0.54	0.45

Comparing model scenarios – sample selection bias

The different model scenarios were designed to test sensitivity to sample selection bias. The AUC results for both the Burkina Faso project region (Table 4-15) and the Marrabel test region (Table 5-13) produced test AUC values less than 0.5 (worse than random) in only one model scenario (respectively, *P5 Stinger* and *P3 Ngapala*). For the Adelaide test region, test AUC values less than 0.5 occurred for some crops in all model scenarios having spatially independent training and test data (scenarios *P2*, *P3*, *P4* and *P5*). Only scenario *P1 Random 25%* had test AUC scores indicating better than random performance for all crops. And only the crops Faba beans and Potatoes had test AUC scores indicating better than random performance in all model scenarios.

Table 5-20 Test AUC scores for crops in the five model scenarios from models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise – Adelaide test region

Model scenario	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes	Mean
Fractional predicted area (FPA)	*49%*	*48%*	*25%*	*72%*	*72%*	*28%*	-
Theoretical max AUC	*0.76*	*0.76*	*0.88*	*0.64*	*0.64*	*0.86*	*0.76*
P1 Random 25%	0.59	0.70	0.88	0.61	0.57	0.84	0.70
P2 McLaren Flat	0.10	0.38	0.62	0.13	0.13	0.64	0.33
P3 Meadows	0.62	0.53	0.80	0.47	0.56	0.80	0.63
P4 Oakbank	0.28	0.56	0.78	0.35	0.38	0.77	0.52
P5 Uraidla	0.30	0.41	0.57	0.33	0.35	0.60	0.43
Mean	0.38	0.52	0.73	0.38	0.40	0.73	0.52

Figure 5-21 presents the scatterplots and correlation coefficients for pairwise comparisons of the suitability maps for each crop from the five model scenarios. Major divergences in model predictions are apparent in many of the scatter plots, indicating substantial spatial bias in the predictions for all crops in some model scenarios.

Extreme sample selection bias is observable in all crop models from scenario *P2 McLaren Flat*. The mean test AUC from all crops models in this model scenario is just 0.33, and all correlation coefficient values that are negative or close to zero (other than for Apples) occur for pairwise comparisons with the maps from this model scenario. The McLaren Flat test site is distinctive in being much flatter and having a different soil profile to the other three test sites. Not including training samples from this site has resulted in very different patterns of predicted suitability across the region as a whole.

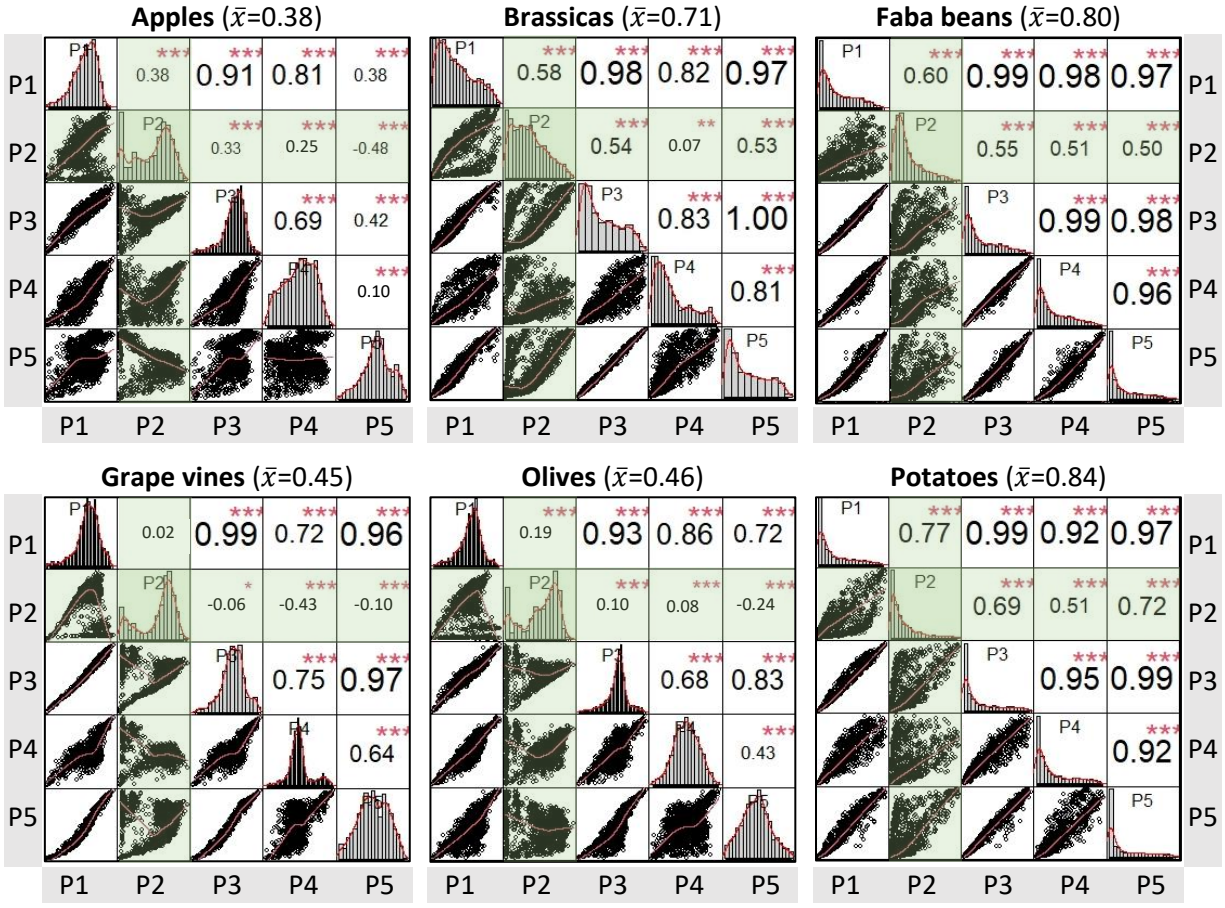

Figure 5-21 Correlations of suitability maps between the five model scenarios from models built using terrain layers (123) and hybrid soil layers (ABCDEFGHI) with added Gaussian noise – Adelaide test region; correlations with maps from scenario P2 McLaren Flat are shaded green

The thresholded crop suitability maps from scenario *P2 McLaren Flat*, that had no training samples from McLaren Flat, are compared in Figure 5-22 to those from scenario *P1 Random 25%* that used training samples from all test sites. Contradictory predictions of suitability are observable for large sections of the western side of the test region – these areas are rated suitable for all crops in the *P1* suitability maps, but rated unsuitable for all crops in the *P2* suitability maps. When training data from McLaren Flat is used (scenarios *P1*, *P3*, *P4* and *P5*), all models produce very similar suitability predictions for Faba beans and Potatoes (correlation coefficients > 0.9) and for Brassicas (correlation coefficients > 0.8).

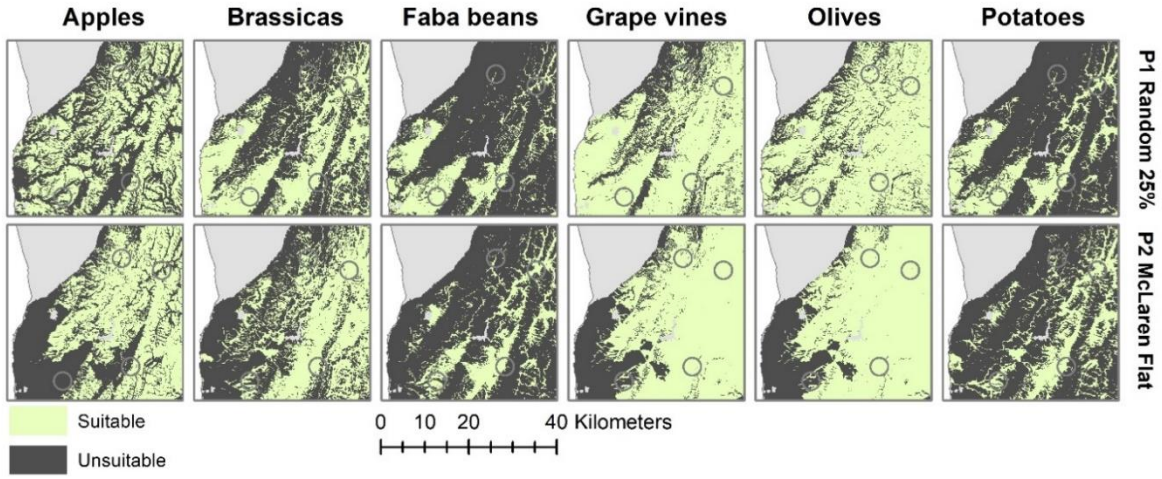

Figure 5-22 Thresholded crop suitability maps from P1 Random 25% models with training samples from all test sites and P2 McLaren Flat models with no training samples from McLaren Flat

In the case of Apples, the correlation results reveal highly inconsistent suitability predictions between the model scenarios. As they grow, apple trees develop wide spreading root systems that have a few deep vertical anchoring roots and a network of fibrous roots in the top 1 m of soil system (Westwood 1988). The soil layers used as predictors for the Maxent models characterise soil qualities in the top 30cm only. The results for Apples, Olive and Grape vines suggest these layers are inadequate for modelling agricultural plants with deeper root systems.

Comparison with Soil and Landscape Grid of Australia soil layers

The predictive performance of crop suitability models built using soil layers derived from the Soil and Landscape Grid of Australia (SLGA) data is compared to models build using hybrid soil layers for the 'best' models (scenario *P1 Random 25%*) in Table 5-21, and across all models scenarios in Table 5-22.

Table 5-21 Comparative performance of models using hybrid soil layers and SLGA soil layers as predictors for the 'best' models (models from scenario P1 built from terrain and soil layers 123ABCDEFGHI)

	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes	Mean
Fractional predicted area (FPA)	*49%*	*48%*	*25%*	*72%*	*72%*	*28%*	-
Theoretical maximum AUC	*0.76*	*0.76*	*0.88*	*0.64*	*0.64*	*0.86*	*0.76*
	Hybrid soil layers with added Gaussian noise						
Test AUC	0.59	0.70	0.88	0.61	0.57	0.84	0.70
Binary accuracy	0.68	0.74	0.84	0.79	0.76	0.81	0.77
Kappa for binary accuracy	0.36	0.49	0.56	0.46	0.32	0.54	0.45
	Soil and Landscape Grid of Australia soil layers						
Test AUC	0.82	0.84	0.91	0.79	0.76	0.91	0.84
Binary accuracy	0.67	0.75	0.82	0.72	0.68	0.86	0.75
Kappa for binary accuracy	0.34	0.49	0.52	0.30	0.21	0.65	0.42

Table 5-22 Summary performance scores for models using hybrid soil layers and SLGA soil layers as predictors; models built from terrain and soil layers (123ABCDEFGHI); means calculated from model scores in all model scenarios

	Apples	Brassicas	Faba beans	Grape vines	Olives	Potatoes	Mean
Fractional predicted area (FPA)	*49%*	*48%*	*25%*	*72%*	*72%*	*28%*	*-*
Theoretical maximum AUC	*0.76*	*0.76*	*0.88*	*0.64*	*0.64*	*0.86*	*0.76*
	Hybrid soil layers with added Gaussian noise						
Mean test AUC	0.38	0.52	0.73	0.38	0.40	0.73	0.52
Mean correlation coefficient	0.38	0.71	0.80	0.45	0.46	0.84	0.61
Mean binary accuracy	0.63	0.72	0.82	0.76	0.73	0.80	0.74
Mean kappa	0.25	0.43	0.51	0.35	0.26	0.50	0.38
	Soil and Landscape Grid of Australia soil layers						
Mean test AUC	0.41	0.48	0.64	0.40	0.42	0.67	0.50
Mean map correlations	0.59	0.66	0.66	0.58	0.59	0.74	0.63
Mean binary accuracy	0.62	0.70	0.80	0.70	0.68	0.83	0.72
Mean kappa	0.25	0.41	0.45	0.24	0.20	0.58	0.36

As with the Marrabel results, both sets of models had similar measured accuracy. The models using SLGA soil layers also had much higher test AUC scores for all crops in the 'best' models, but lower *mean* test AUC scores across the five model scenarios.

For both sets of models, the least similarity of suitability maps between model scenarios (measured by mean correlation coefficient) occurred for the crops with deeper root systems (Apples, Grape vines and Olives) and the greatest map similarity and highest measured accuracy occurred for the crops with the most restricted ecological niches (Faba beans and Potatoes).

Conclusions

The results for the Adelaide test region demonstrate limitations that are relevant to the application of the methods used in this thesis. The research design for the South Australian modelling mimicked the configuration of clustered West African occurrence data and used cross-validation to detect and assess the influence of sample selection bias on generated models.

The Adelaide test region contains a great deal of productive farmland on its coastal plains and hilly inland, and its diverse topography also includes many areas too rugged for agriculture. Unrepresentative training data makes spatial bias in models highly likely, and this is clearly observable in the suitability maps produced in model scenario *P2 McLaren Flat*. Having only training examples from hilly areas, these models were unable to detect suitability for growing crops on flatter land. This deficiency was overcome, in this case, by inclusion of representative training points from McLaren Flat in the other model scenarios. However, if a strong spatial bias applies to the full set of occurrence data available for training models then the problem cannot be corrected in this way. The risk of spatially biased training data is high in situations where presence data are incidental occurrence data collected for another purpose and may necessitate the collection of further training samples from other locations to address the problem.

The selection of the environmental variables to use in modelling should be based on biological reasoning and, where possible, include specific variables known or suspected to affect a species distribution. The predictor layers used for the West African modelling were selected to model

suitability for field crops that grow in top 30cm soil. These predictors were effective for similar crops in the Adelaide test region, but proved inadequate as predictors for modelling tree crops that have deeper root systems.

5.6 Summary

The purpose of Chapter 5 has been twofold. Firstly, to validate the methodology used in Chapter 4 by applying it to sites for which region-wide validation data are available. And, secondly, to demonstrate the transferability of the method to other sites with different terrain, climate and styles of agriculture. The methods used in Chapter 4 were duplicated for two local sites in South Australia for which region wide validation data are available and that can be visited to assess results and evaluate prototype presentation methods.

The validation data for the South Australian sites was in the form of categorical maps of land use potential for a number of commonly grown crops. Occurrence data were not available for these sites but could be artificially generated from the validation maps. Pseudo-presence points were created for six locally grown crops at each site and were clustered into four small areas to mirror the configuration of the Burkina Faso occurrence data. Equivalent raster environmental layers were generated using the methods described Chapter 2. The Maxent software was then used to train crop suitability models from the pseudo-presence points and environmental predictors in five cross-validation model scenarios, as was done in Chapter 4.

Evaluation of the crop models used test AUC and correlation coefficients, as in Chapter 4, together with region-wide accuracy scores obtained from comparison with the validation data. The validation maps of categorically valued land use potential polygons were not directly comparable to the continuously valued suitability rasters. To facilitate formal interrater comparison, the maps were converted to binary rasters with values indicating suitability or unsuitability for their target crop. The fractional predicted area (FPA) and theoretical maximum AUC were calculated for each crop from its binary verification map and each suitability raster was thresholded to produce a binary suitability map with the same FPA. Interrater comparison of the pairs of maps produced scores for accuracy and kappa coefficient.

In the case of the Marrabel test region, the methods and environmental predictors used to produce the final suitability maps in Chapter 4 resulted in robust crop suitability predictions that closely matched the region wide validation data - evidenced by highly correlated suitability maps between model scenarios and high measured accuracy from the interrater comparisons of binary suitability and verification maps.

In the Adelaide test region, the same environmental predictors were shown to be effective predictors for the field crops with root systems in the top 30 cm of soil, but not for modelling suitability for other crops with deeper root systems. Alternative sets of hybrid soil layers using different formulations of the source soil property rasters may have provided effective predictors for these crops (e.g. based on depth to bedrock, consolidated to greater depths, or unconsolidated at all depths) but were these not investigated during this project.

A core theme in this thesis has been the use of cross-validation model scenarios that have spatially independent training and test data as a technique for detecting and assessing the influence of sample selection bias in resulting models. The Adelaide test region provided a powerful example of how biased training data can spatially distort model predictions, and this issue can never be ignored when data-driven modelling algorithms are used. In this chapter the feasibility of the South Australian crop predictions was readily assessable by comparison with

region-wide validation data (that also rendered the modelling task redundant for these sites). When a common bias affects *all* occurrence data the techniques of comparing test AUC scores and correlating suitability maps from cross-validation model scenarios are unlikely to detect it; further qualitative assessments (such as those used in Chapter 4) are required to establish the feasibility and goodness of models.

Chapter 6 Disseminating results

6.1 Introduction

This chapter investigates ways of presenting the crop suitability maps developed in Chapter 4 so that they can be effectively delivered into affected communities to inform farmers and land managers making decisions about land use. For the maps to be useful to their target audience they need to be designed to be usable by this target population and, in this case, illiteracy and poor access to internet resources bring special challenges to the task of producing informative maps.

In designing a map for a particular use it is necessary to consider the context for the map which Griffin, White et al. (2017) define to include four components. These are: (1) the map user; (2) the environment where the map is being used; (3) the activity performed using the map; and (4) the graphical object that is the map. Distinct map contexts will emerge from individual map use situations. Here, the primary intended map users are subsistence farmers who will be displaced by a new gold mine in rural south-west Burkina Faso, with the primary map activity to be identifying the land use potential of particular land for growing particular crops. Different map production technologies (e.g. paper versus electronic maps) are available for developing these land use potential maps, with the different technologies producing different requirements for effective map design.

This chapter first considers the map user profile to identify issues of significance for effective map design for this audience. The opportunities and constraints that apply to different styles of map production are then considered, leading to the development of several prototype maps designed specifically for the target audience. These map solutions are described in detail and include both printable formats suitable for users without access to electronic devices, and electronic versions suitable for browsing on a computer, tablet or mobile phone, with and without internet access.

6.2 Map user profile

The social impact assessment (SIA) prepared for the Banfora Gold Project (Intersocial Consulting 2014) provides socio-economic baseline data for the local region and reports on very poor communities with low levels of education and literacy.

Education

The SIA reports very low levels of education for members of the affected population, with most having had no education and only a few having completed primary school. In the adult population over the age of 20, approximately 55% of males and 78% of females have had no formal education, and only 24% of males and 8% of females have completed primary school. Although education is meant to be compulsory and free for children aged 6-16 years, inadequate government resources means fees are often levied for the cost of school supplies and infrastructure, impacting school attendance in poor communities. The SIA reports that in Léraba Province almost 40% of children aged 6-11 years of age were not attending primary school, and just 29% of boys and 16% of girls aged 12-18 years were attending secondary school.

Languages

Burkina Faso is a linguistically diverse country with 59 spoken languages (Kone 2010). French is the official language but is not widely spoken and less than 15% of the population speak French on a daily basis (Kone 2010). The languages Dioula, Fulfulde and Mòoré are also recognised as national languages: Mòoré is the most widely spoken language with almost half the population being speakers; Fulfulde is a local lingua franca and widely spoken as a first language in the north and east of the country; and Dioula developed as a trading language and is spoken in western Burkina Faso (NALRC 2010). Many of the local languages are only spoken and do not have a written form.

Dioula is the principal language spoken in the south-western Cascades Region (including the Banfora area) followed by Mòoré and Senoufo. Within the study region in Léraba Province the Senoufo ethnicity and language predominate, although Dioula is also spoken by most of the population and used to communicate with members of local authorities who are not native to the region (Intersocial Consulting 2014). The Dioula language has a written form that uses a phonetic alphabet but Senoufo has no written form.

Literacy

The literacy rate for Burkina Faso is low, with less than half the population able to read and write. In 2018 the literacy rate for the population aged 15 years and older was recorded as 39.3%, but this was a substantial improvement on the literacy rate of 22.5% in 2006 and of just 12.8% in 1996 (UNESCO Institute of Statistics 2022). Kone (2010) suggests that the use of French as the language of formal instruction in schools has contributed to low literacy rates and high dropout rates. However, the author also observes that attempts to include first-acquired language as the medium of instruction in schools have generally proven unsuccessful due to the perception that spoken and written competence in French offers greater opportunities for social and physical mobility than literacy in local languages, and to the challenges of providing text books and teaching resources in multiple languages.

Although French remains the primary language of instruction in schools, the government has created adult literacy classes in local languages to combat illiteracy across the broader population. The SIA reported that 8% of males and 7% of females in the survey population had undergone training in “alphabetisation” centres and many others in the focus groups had expressed interest in attending such courses (Intersocial Consulting 2014, page 49).

The literacy levels for the population affected by displacement are quantified in the SIA. Overall, only one out of four community members can read or write. Knowledge of French is very limited: three quarters of the population reported not to speak, read or write any French, and only 3% had gained a good level of spoken and written French. The most common languages spoken by the affected communities are Senoufo and Dioula. Senoufo does not have a written form, but in all affected villages at least one person can read Dioula (Intersocial Consulting 2014, page 80).

Access to technology

The SIA reports that cell phones are the one of the most commonly owned household assets in the survey population, with 90% of households having a cell phone. More than 70% of households also reported having a solar panel for electricity generation (Intersocial Consulting 2014, page 77).

Source: I. Ahmer

Source: I. Ahmer

Figure 6-1 Solar panels on Banfora rooftop; solar panels for sale in Banfora (2018)

The rapid adoption of mobile phones across Burkina Faso was examined by Hahn and Kibora (2008) who observed that the growth of mobile phone usage in Burkina Faso and many other African countries has been faster than in Western countries, in spite of the difficult economic situation of many of the users. The use of mobile phones in Burkina Faso commenced in 1996 and within ten years had penetrated the most remote areas of the country. Their use is now embedded in society, although SMS and MMS services are little used in a country with low literacy rates. Low priced new phones are imported, typically from Dubai, and second-hand and even damaged phones come from France and other European countries. This has given rise to a new business of mobile phone maintenance offering a variety of services such as decoding to remove prior network restrictions, replacement of broken parts, and recharging of batteries. In rural areas people creatively look for sources to charge their phone batteries, e.g. through solar panels and car batteries. In such a poor country many mobile phone owners try to adopt a 'zero budget strategy' to avoid ongoing usage costs and various local practices have emerged that use phones to communicate without incurring call costs (Hahn and Kibora 2008).

In 2015 the Pew Research Centre conducted a survey of 40 countries to determine smart phone ownership and internet usage. The results of the survey for Burkina Faso showed that 79% of adults owned a cell phone, 14% owned a smart phone and 18% used the internet at least occasionally or reported owning a smart phone. The survey identified a strong relationship between per capita income and internet access, and demonstrated that younger, more educated and higher-income people everywhere had greater access to the internet (Poushter 2016).

Although it is not anticipated that households in the project region own computers or smart phones, it is expected that some access to computers may be available. The SIA reports a large number of community committees and organisations across villages in the project area - if access to computers were to prove a significant issue, it would seem within the scope of the mining company's social licence to also provide computers to some of these community groups.

Implications for maps

The profile of the map users raises several major issues that must be considered if the final maps are to be useful to them. Literacy and education levels are very low in the target community so the maps must be intuitive to understand, language independent and with sufficient embedded context for the users to correctly interpret the spatial dimensions of the maps. Further, to successfully deliver the information into such poor communities, the maps must have no costs associated with them.

6.3 Styles of map presentation

Maps can be designed to be printed on paper or be browsable on electronic devices such as computers, tablets or smart phones. The medium for presenting a map is the major technical context factor that constrains map design. Griffin, White et al. (2017) observe:

> *"… many design decisions relate to the choice of which representations to use, which information to display on a device of a given size, and how to coordinate representations so that map users can (cognitively) fit all of the information displayed together."*

Paper maps present a single static view of their content, whereas maps displayed using electronic devices can be dynamic applications that allow users to adjust the scale or alter the view in other ways. Both styles of map presentation present constraints to the map designer, but each has advantages that can be effectively leveraged for particular map use situations.

6.3.1 Paper maps

Traditional paper maps are static maps that present a single view of their subject. Their content does not change, so all map elements and symbology used in the original design are always visible. In contrast to electronic maps, paper maps have almost zero reliance on technology. They are easily reproduced and have high persistence (as paper is an enduring material).

A paper map implementation imposes major design constraints in relation to the amount of information that can be fitted onto a page or that can be displayed in a single map. All map elements must be presented with fixed extent and scale, often severely limiting the amount of available detail that can be displayed. Complementary maps showing different content for an area often need to be presented as separate maps, with the user having to mentally combine the information on the maps (rather than turning features on or off as can be done with electronic maps).

Although paper maps can be challenging to design, they provide a robust means for delivering content. The map user always views the entire map content as it was designed to be viewed, and it is not possible for the user to accidently hide content or otherwise change the view of maps presented.

6.3.2 Electronic maps

Electronic maps offer many advantages for the map user. As neither the scale nor the extent of the map is fixed, the user is able to pan to areas of interest and zoom in to view very fine detail at those locations. Overlaying the suitability maps on informative base maps allows the user to make additional intuitive assessments on issues such as proximity to infrastructure and terrain features, and an up-to-date base map of satellite imagery allows a user to closely inspect what currently exists at a location.

The most serious limitation with electronic map implementations is the high dependence on technology. Without access to an electronic device that can display maps and a reliable power supply the maps are inaccessible. Even with these basic requirements, access to the maps cannot be guaranteed. It may be beyond the skills of the user to install the necessary software to run the application, and the software may not run on a particular device for which it had not been tested. Even after successful installation on a target device the application may have weak persistence as hardware and software changes can easily impact the capacity to run it on the device.

Base maps are commonly used to provide added context for subject maps in electronic map applications, and map applications designed to run with base maps often lose much of their functionality without the base maps. Many of the popularly used base maps are supplied as map tiles by map servers on the internet in real time and are unavailable without an internet connection.

Whether target users have access to suitable equipment and sufficient computing skills to run the applications are critical issues to consider when designing electronic map applications. Awareness of internet availability and its cost is also an important consideration when using base maps from the internet.

6.3.3 Maps for illiterate and semi-literate users

For a map to be useful it needs to be functionally correct, pleasing and efficient to use, and easy to learn and remember. Illiteracy places strong limits on map design; firstly, because it requires a map to be understandable without relying on text components and, secondly, because map reading and map use skills of illiterate people may operate differently to those of literate people (Griffin, White et al. 2017).

Medhi, Sagar et al. (2006) explored the development of map applications with text-free user interfaces for illiterate and semi-literate household workers in India for whom maps were not common artefacts. Their extensive field studies showed that graphics and photorealism were effective substitutes for text, numbers were also easily recognisable, and landmarks were important for geographic navigation. They also noted the importance of paying attention to subtle graphics cues as user responses may be affected by psychological, cultural or religious bias.

The design and evaluation of a text-free online map interface for illiterate people and non-local-language speakers is described by Bao (2016). Symbols, audio, photographs and video were all effective for replacing the traditional text. Although audio was the favoured medium for illiterate people using interactive electronic maps, it was not useful for non-native speakers who instead required text-free interfaces with the support of visual representations.

6.3.4 Design and implementation issues for the crop suitability maps

The profile of the map users raises several major issues that must be considered if the final maps are to be useful to them. Literacy and education levels are very low in the target community so the maps must be intuitive to understand, language independent, and with sufficient embedded context for the users to correctly interpret the spatial dimensions of the maps. To successfully deliver the information into such poor communities, the maps must also have no costs associated with them.

The maps of crop suitability predictions for individual crops presented in Figure 4-21 are the subject maps for dissemination. These suitability maps are georeferenced raster images that are not suitable for overlay on each other and so must be presented as a set of complementary maps

that are each viewed individually. Additional reference maps and/or map data are required to supply the necessary geographic context for the map users to mentally locate themselves on the maps and identify locations of interest. There are many types of reference maps that can be used: political maps show administrative boundaries; physical maps show terrain features (rivers, mountains, valleys, etc.) with topographic maps also using contour lines to show elevation; road maps show transport networks (roads, railways, etc.) and the locations of towns and villages. Satellite imagery, at appropriate resolution, can also be useful for viewing terrain, observing landcover and identifying landmarks.

Readability is a critical consideration when designing maps: how much detail will fit into the available image space before the design becomes too cluttered and confusing for the map user to make sense of it? With paper maps, all detail must be readable at the scale at which the map is reproduced. With electronic maps, certain detail is often relevant only at certain scales (zoom levels) and can usually be suppressed if it clutters the display at other scales.

Prototype maps that address these project specific issues have been developed for both paper and electronic presentation formats and are described in the following sections.

6.3.5 Map prototypes

The map prototypes from this thesis are available for inspection on the Box cloud storage system at the following links:

https://universityofadelaide.box.com/s/q746bzprxljdw6wkgmqw1xgtb4b730j1

https://tinyurl.com/448x9233

The 'Paper maps' subdirectory contains each paper map prototype in both JPEG and PDF formats.

The 'Leaflet maps' subdirectory contains the HTML files and supporting Leaflet files for each of the electronic map prototypes. In order to use the Leaflet maps it is necessary to download the entire folder 'Leaflet maps' (242MB compressed). Please decompress the generated 'zip' archive before using the online maps.

Opening the HTML file for the map will display the map in a web browser. The Leaflet maps for the Burkina Faso project region that are described in this chapter are as follows:

Map file	**Language**	**Base maps**
B_Leaflet_maps.html	English	Internet
B_Leaflet_maps_French.html	French	Internet
B_Leaflet_maps_Dioula.html	Dioula	Internet
B_Leaflet_maps_Dioula_standalone.html	Dioula	Within the maps

The 'Leaflet maps' subdirectory also contains Leaflet maps to display the suitability maps and verification maps for the Marrabel and Adelaide test regions. These maps are not discussed in this thesis.

6.4 Paper map prototypes

The choice of paper size constrains the amount of detail that can be included on a paper map. A map that can be printed using a normal desktop printer or photocopied at a standard paper size[26] usually has negligible reproduction costs and so can be an extremely low-cost artefact. Maps designed for A4 and A3 paper are easily reproduced and can be printed at larger sizes if required, but a map designed for poster display is rarely easily readable when printed at the smaller paper sizes and is usually much costlier or more difficult to print at actual size. To ensure ease of reproduction for the paper maps from this project, it is desirable that the map pages be no larger than A3 in size, but still be readable if printed on A4 paper.

For a map to be useful it must offer enough information about geographic location and terrain for the users to readily locate themselves on the map and interpret results correctly. Paper maps are printed at a fixed scale and must be easily readable at this scale. Where the necessary map detail is greater than can be shown in the available image space, maps are typically partitioned into multiple smaller extents that are shown as separate maps using the map book approach of atlases and street directories. The map book approach increases the complexity of the map artefact as it requires an additional index map to look up the relevant map page, and often splits geographical areas of interest over multiple map pages, increasing the cognitive load on the map user. Whether a map book approach is needed is determined by the amount of detail that must be included in the maps.

6.4.1 Map content

The primary purpose of the map artefact is to disseminate mapped predictions of crop suitability across the project region.

Subject maps

The five crop suitability maps are raster images, each containing 1333 x 1333 pixels, georeferenced to cover the same physical extent. Each pixel represents a distinct 30 m x 30 m location in the project region, and the relative suitability for growing a crop at that location is indicated by the colour used to print that pixel in the relevant crop suitability map. Each of the five maps must be presented separately as overprinting of the rasters would distort their content. Each suitability map needs a title to identify the target crop and a legend for interpreting the pixel colours. However, if all maps are rendered using the same symbology, a single legend can be used for the five maps.

Reference maps

Reference map data are necessary to allow the map user to identify locations of interest in the subject maps. In some circumstances reference data can be overlaid as additional detail on the subject maps. But when the detail in the reference map would obscure the detail of the subject map it must be presented as a separate map with sufficient contextual cues for the map user to cognitively combine the content.

The most useful context maps for the map users of this application are: (1) a navigational map showing roads and locations of towns and villages in the project region; and (2) a physical terrain

[26] Measurements in mm for standard international paper sizes: A0 = 1188 x 841; A1 = 841 x 594; A2 = 594 x 420; A3 = 420 x 297; A4 = 297 x 210.

map revealing landscape features. A locator map with an extent indicator for the project region is also necessary to pinpoint its geographical location at the global scale.

Satellite imagery is often offered as a reference base map in electronic map applications, allowing a user to zoom in to inspect observable features in finer detail. Satellite imagery was tested as potential reference data for the paper maps but it proved ineffective and uninformative at all scales at which the maps could feasibly be printed, and so was not used.

Navigation map

The majority of the content for the navigation map comes from OpenStreetMap, a free geographic database of the world released with an open-content licence.[27] When all relevant layers (roads, waterways, natural features, land use, and places with place names) were plotted for the project region the resulting map was uncluttered and was easily readable as a single printed map. Additional map detail to enhance the relevance of the map included the international border and local administrative boundaries.[28]

Terrain map

A digital elevation model (DEM), suitably symbolised using intuitive colours for elevation, provides an informative image of the physical landscape for a region. A light hill-shade effect can be applied to assist three-dimensional interpretation of the data and contour lines can be added to quantify the changes in elevation. Although relatively few permanent waterways are marked on maps of the project region, the paths of water flow in the landscape are particularly informative for farmers. These can be added to the terrain map using the drainage channel layer that was supplied by Gryphon Minerals Ltd.

Locator map

A simple country outline for Burkina Faso with a square to mark the extent of the project region proved sufficient as a locator map due to the easily recognisable border location of the project region.

Combining the map content

Exploration of the map content showed that the detail for the whole project region could be readily mapped at the scale of a printed page, and so it was unnecessary to partition the region to make an atlas of sub-regions. The crop suitability maps cannot be overlaid on each other, although reference map data can be overlaid on the crop maps or displayed adjacent to them. Thus, a map artefact to disseminate the suitability predictions will necessarily contain multiple map frames, and the size of each map frame will constrain the amount of map detail that can be included in it.

Complementary design opportunities have been explored to create two prototype paper map solutions. The first is a single page map design that shows five small crop suitability maps, with navigation and terrain information presented in separate reference maps. The second prototype is a multi-page set of large suitability maps that have navigational and terrain information

[27] The OpenStreetMap layers used in this project were downloaded from https://extract.bbbike.org.

[28] International border sourced from the ESRI World Countries map package. Administrative boundaries downloaded from https://data.humdata.org/dataset/cod-ab-bfa.

overlaid. Both map prototypes are designed to be easily interpretable regardless of the language of the map user or personal level of literacy.

6.4.2 Language independence

The most serious map design challenge for this project is the diversity of languages used and the low level of literacy among the target map users. Three quarters of the community were reported to be illiterate, and those who can read are more likely to be literate in a local language such as Dioula than in the official language French.

In order for the maps to be interpretable by illiterate map users, all the map legends are text-free (refer Figure 6-3 below). The terrain legend shows a landscape profile that has been coloured using symbology corresponding to the altitude values in the map. The single suitability map legend uses degrees of greenness to indicate degrees of suitability and indicates poor suitability by a graphical image of a brown dying plant and high suitability by an image of a healthy green plant. The symbology used for the suitability maps is visually distinct from the symbology used for the terrain map to avoid confusion of interpretation.

The suitability map titles use a combination of photo-realism and text to identify the relevant crops (see Figure 6-2). A characteristic crop photograph is supplied for illiterate map users and the crop name is written in Dioula (red), French (black) and English (blue). Icons for the flags of Burkina Faso, France and United Kingdom are shown beside the crop names to provide explanation of the languages used.[29]

Figure 6-2 Suitability map titles: crop photograph and text in Dioula (red), French (black) and English (blue)

The title for each map, "*prévisions d'aptitude des cultures*", is the French translation of "*crop suitability predictions*" and the word "*terrain*" has the same meaning in both French and English. These titles and the text in the scale bar are written only in French and shown in black type.

6.4.3 Single page map prototype

The paper map prototype combining all necessary content on a single page is shown in Figure 6-4. It was designed using the ArcGIS Pro software for printing on A3 paper. The layout presents a large navigation map referencing a small locator map in the top half of the page, and six smaller maps showing the terrain and suitability predictions for the five crops in the lower half of the page.

[29] Diola translations sourced from https://www.webonary.org/dioula-bf; French translations sourced from https://www.translate.com/english-french.

The purpose of the navigation map is to allow map users to orient themselves in geographic space. Towns, roads and rivers are powerful landmarks – they are symbolised in the navigation map using intuitive colours and styles (e.g. brown lines of different weight for roads, blue rivers, dashed grey lines for administrative boundaries), thus making a legend for these features redundant. A subtle indicator of terrain features has also been added using a pale grey hill-shade effect base layer.

The reference maps and the suitability maps contain complementary content that must be combined in the mind of the map user in order for the maps to deliver their informational content. To make this synthesis possible, and to simplify the cognitive process, the maps need to contain compelling geographic reference cues to assist the map user. Grids with spacings of twenty kilometres (red), five kilometres (dark grey) and one kilometre (light grey) are printed over every map to facilitate identification of the same location on each individual map and allow the map user to link the content from the different maps.

A scale bar is provided for the navigation map and the grids make scale bars redundant on the smaller maps. To provide further geographic reference cues, the locations of the towns and villages are also marked on each of the smaller maps (but are only named on the navigational map). Figure 6-3 shows corresponding detail from the navigation map, terrain map and two crop suitability maps at the size it would appear if printed on A3 paper.

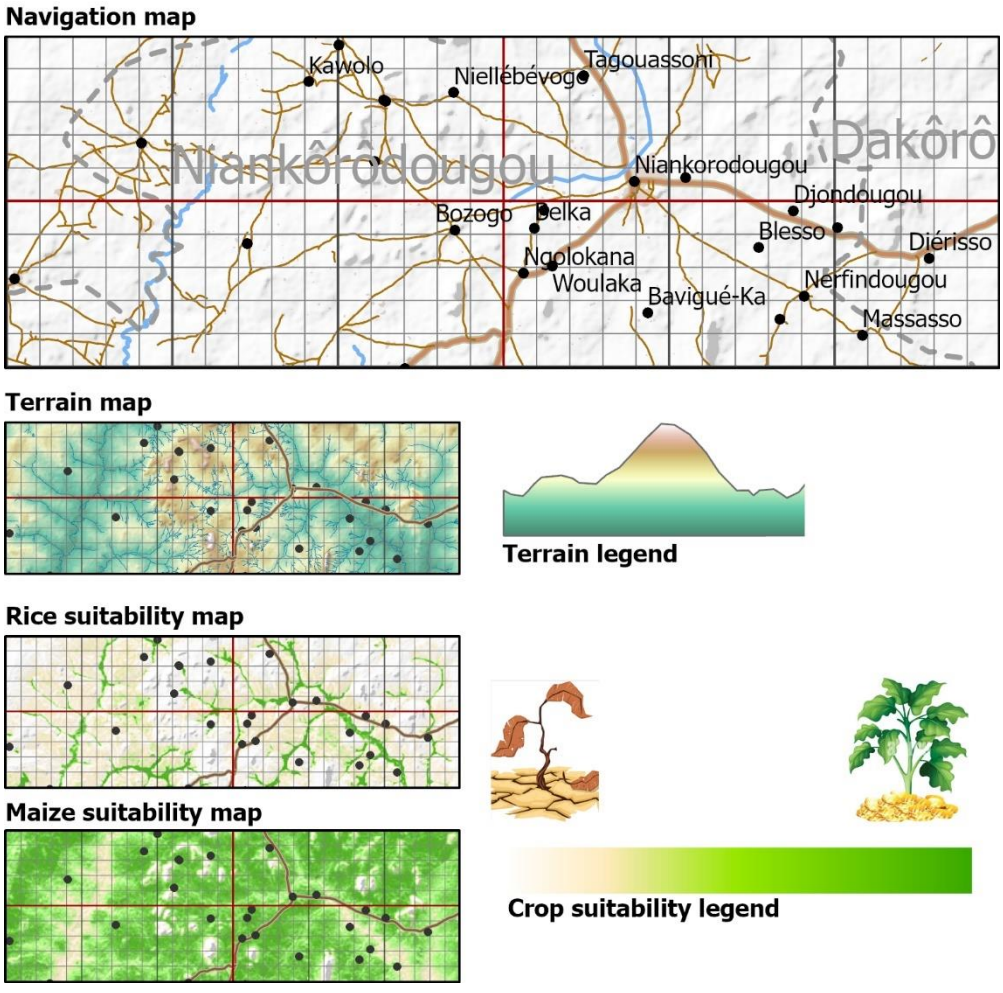

Figure 6-3 Detail from the centre of the navigation map with corresponding detail from the terrain map and two crop suitability maps (actual size when maps printed on A3 paper)

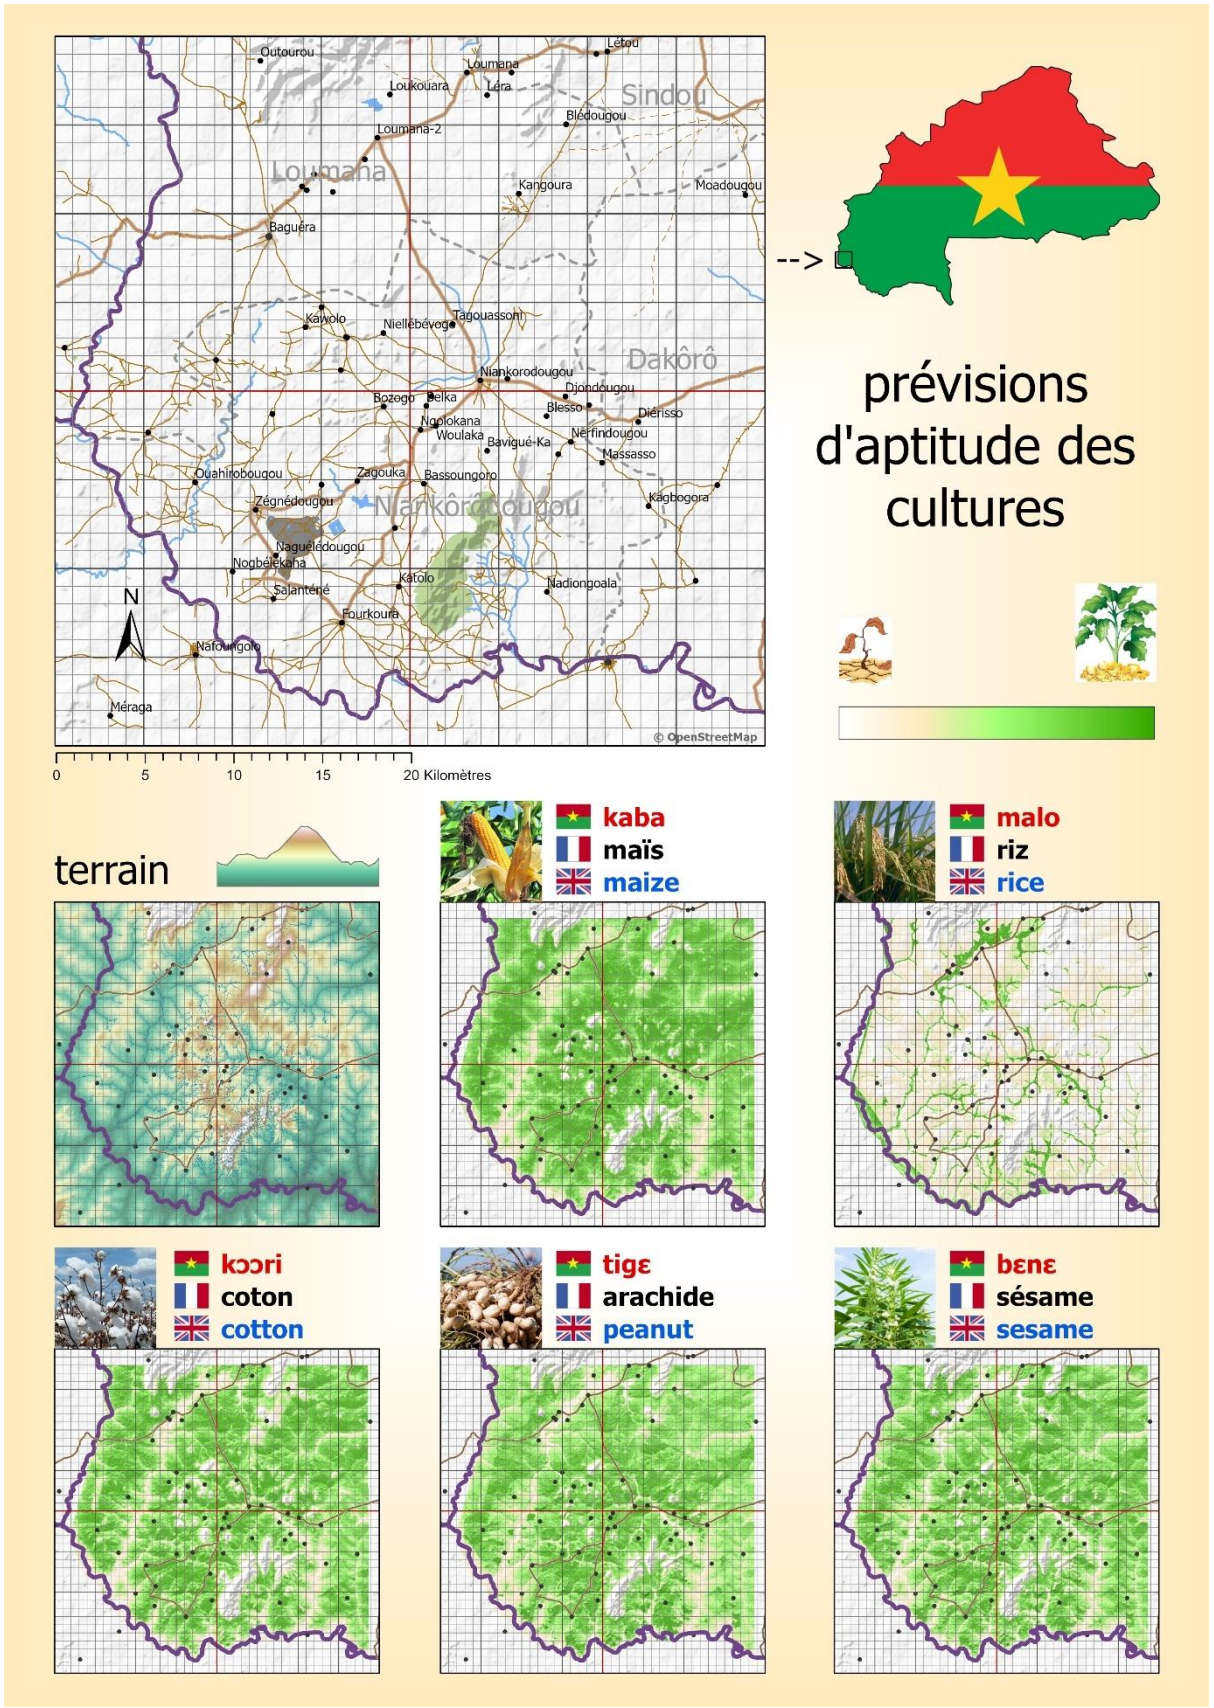

Figure 6-4 Paper map showing crop suitability predictions for five crops (designed for printing on A3 paper)

6.4.4 Multi-page map set prototype

The multi-page paper map prototype is shown in miniature in Figure 6-5 and at larger scale in Figure 6-6 and Appendix H. The set of maps consists of a terrain map and five crop suitability maps, each designed using the ArcGIS Pro software for printing on A3 paper.

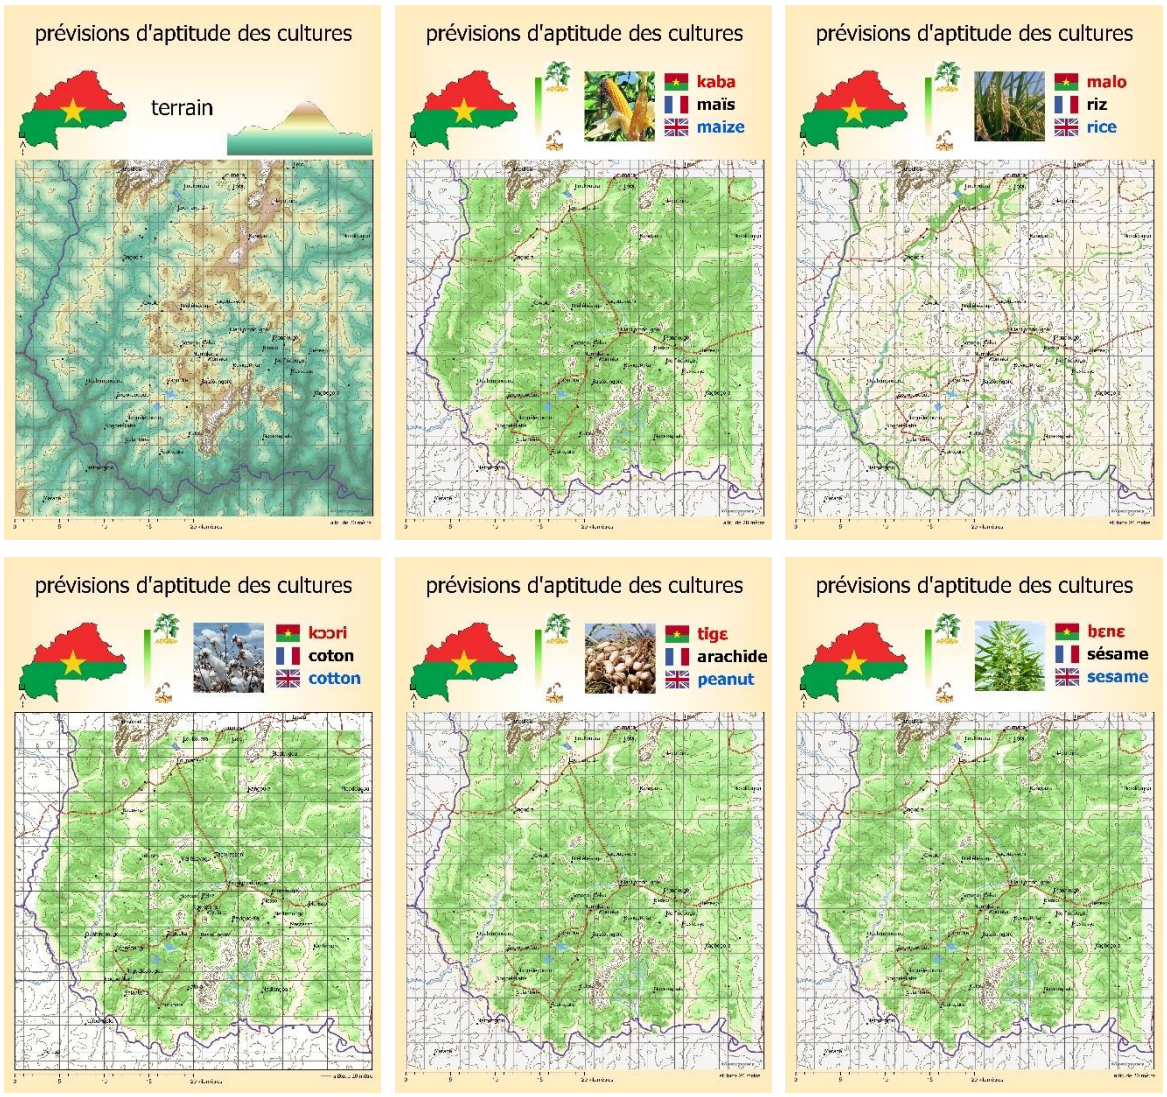

Figure 6-5 Set of paper maps showing terrain and crop suitability predictions for five crops (each designed for printing on A3 paper)

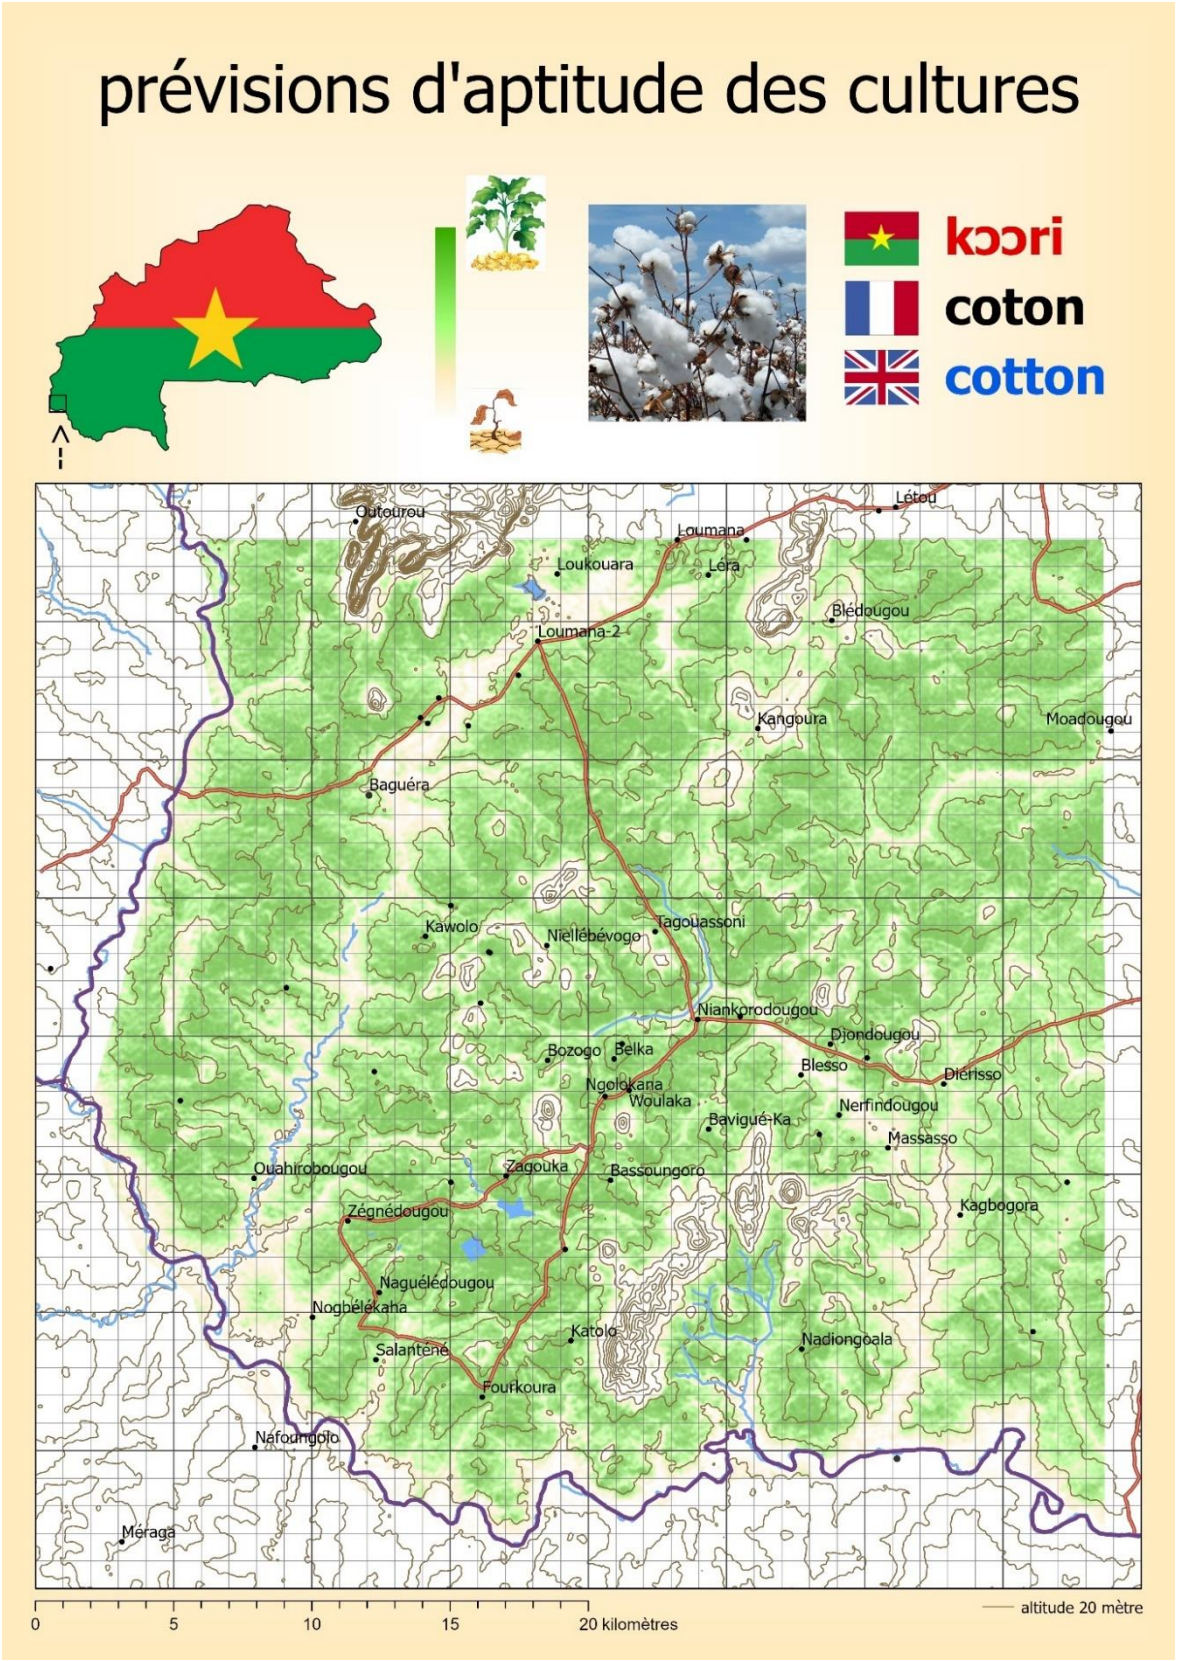

Figure 6-6 Paper map showing crop suitability predictions for cotton (designed for printing on A3 paper)

The maps use the same titles and legends as the single page map prototype, but the larger map frames allow greater amounts of reference data to be overlaid on the raster images than was

possible in the previous prototype. Contour lines showing 20 metre changes of elevation are shown on all the maps to allow greater insight into the terrain features influencing suitability for particular crops. The locations of towns and villages are also identified on all maps. The rivers and major roads have been overlaid on the crop suitability maps, thus making a separate navigational map redundant. Gridlines of 1 km and 5 km spacing have again been provided to allow the map user to link map content for particular geographical locations.

Figure 6-7 illustrates the complementary information presented in the terrain and crop suitability maps. Drainage channels and areas of steep slope are discernible in the terrain map. Although labelling the contour lines with their altitude was not practical on maps of this scale, the symbology of the terrain map quickly resolves any confusion as to whether one contour line marks an area of lower or higher altitude than another.

Terrain map

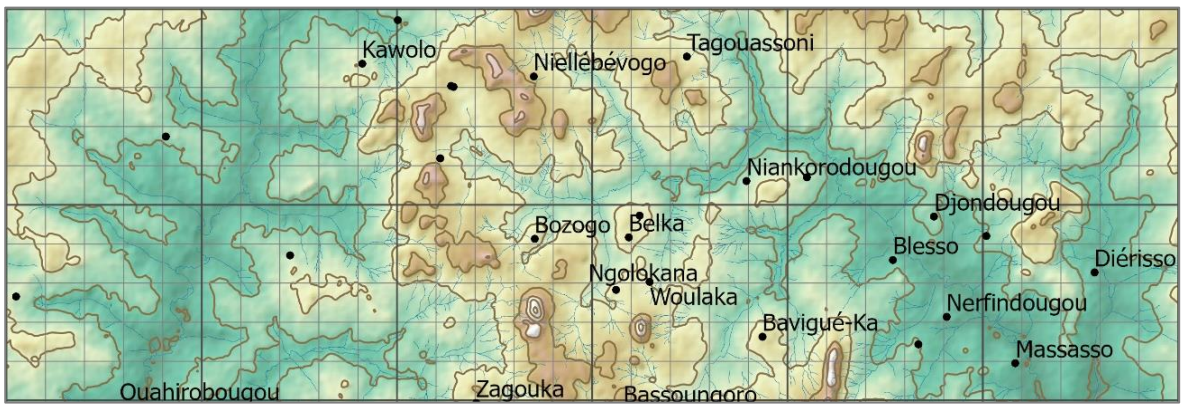

Maize suitability map

Figure 6-7 Detail from the centre of the terrain and maize suitability maps - 20 metre contour lines shown in brown and major roads shown in red

6.5 Electronic map prototypes

An effective electronic map application for this project needs to be cost-free, robust and simple to operate. The application should be able to run on any electronic device capable of displaying maps and require neither proprietary software nor technical support for its operation. The maps need relevant and descriptive context maps plus GPS positioning for mobile devices to answer the question: "I am here, what does this app have to say about this location?"

6.5.1 Leaflet software

"Leaflet" is free and open-source software for creating interactive maps that work efficiently across all major desktop and mobile platforms. The only software required to display a Leaflet map is a web browser, and this software now comes preinstalled on most computers, tablets and smart phones. Because Leaflet was designed with simplicity, performance and usability in mind, it offers intuitive map controls and allows GPS positioning.

Leaflet is implemented as a JavaScript library but can also be accessed from within the R programming environment using the R package 'leaflet'.[30] A Leaflet map application developed using R is distributed as an HTML file with an associated subdirectory of Leaflet map files. Simply opening the HTML file will display the interactive map in the chosen web browser, regardless of internet connectivity. Access to the internet is relevant only when it is needed for map content. A map will display correctly without an internet connection if all map content has been provided by the map developer. However, if the geographic context for the application is provided using base maps served as map tiles from the internet, this context will be unavailable without a connection.

6.5.2 Leaflet map with tiled base maps

The use of base maps served from the internet using map tiles was popularised by Google Maps and is now used by nearly all interactive web maps. The amount of detail supplied by the base map tiles varies according to the zoom level chosen by the user, with more detail becoming visible as a user zooms in at a location. Base maps that show road networks, satellite imagery or terrain are easily interpreted so that there is no need for legends or even titles to explain them.

Subject maps are typically overlaid on a base map using some transparency so that the geographic context from the underlying base map is still visible. A map user can navigate to an area of interest and compare the subject maps with geographic detail from the base maps. The Leaflet layer controls allow the user to switch between base maps and overlay a chosen subject map.

Leaflet maps that overlay original map content on publicly available base maps can be developed very easily using R. Figure 6-8 shows examples of three language specific Leaflets maps created to disseminate the results of this research project. Three alternative base maps are offered: (1) OpenStreetMap - for navigation; (2) ESRI World Imagery - providing satellite photographs of the Earth's surface; and (3) ESRI Shaded Relief - showing landform. Each of the five crop suitability maps can be overlaid on any of the base maps. The same symbology is used for each of the suitability maps so that only one legend is required. GPS positioning will identify the map user's current location on the map if their display device is GPS enabled.

[30] Refer https://leafletjs.com/ for the JavaScript library and https://rstudio.github.io/leaflet/ for the R package.

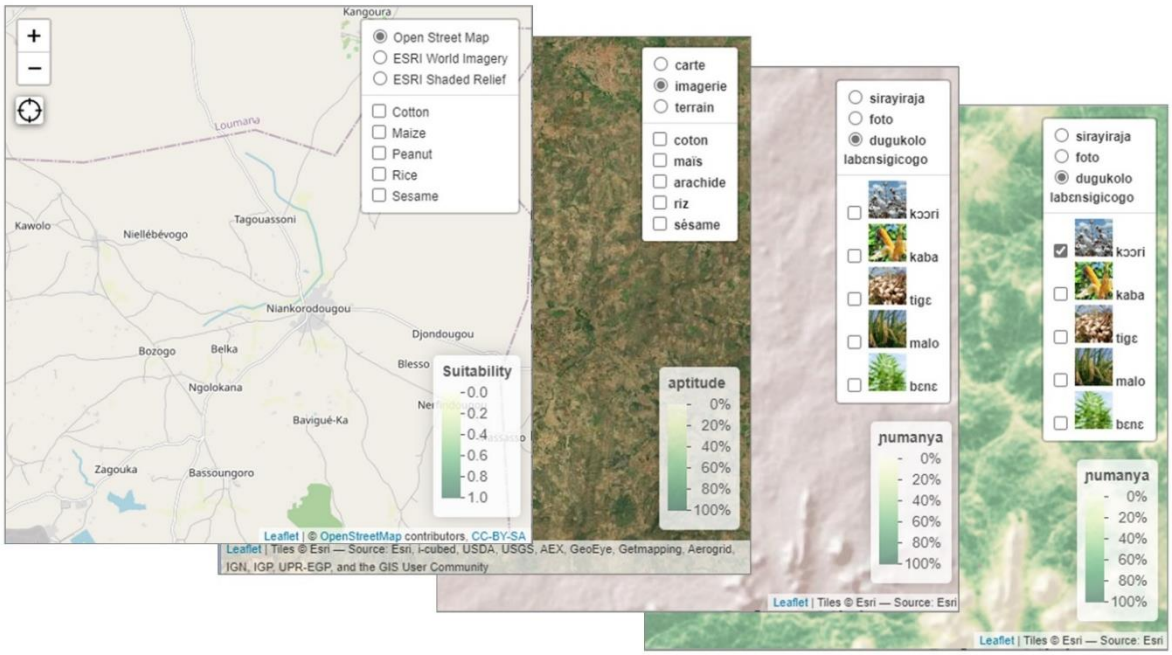

Figure 6-8 Examples of Leaflet maps with language specific controls showing alternative base maps: (1) English + OpenStreetMap; (2) French + satellite imagery; (3) Dioula + terrain; (4) Dioula + terrain + cotton suitability map

A completely text-free or a multi-lingual interface would be difficult to achieve in a simple map as too many images or too much text would clutter the map controls. Instead, language specific versions of the maps have been created for English, French and Dioula. Text descriptions for the base maps are used in all versions for simplicity – their content is readily understandable, and appropriate map icons are difficult to design. The English and French maps assume literacy in that language to identify the target crops for the suitability maps, but images of the crops have been added to the Dioula version to assist illiterate map users. In addition, the legend in the French and Dioula versions expresses the suitability rating using the range 0% to 100% as this scale is more easily understandable than the 0 to 1 values of the suitability rasters used in the English version.

Leaflet's intuitive map controls make a user guide unnecessary. The legend and map controls are always kept visible in these maps (and so can't accidently be hidden) to simplify the interface for users with little experience with using maps or electronic devices. Explorative use of the interface quickly reveals its functionality: checking a radio button selects the base map; the tick box for a crop turns on or off the overlay of its suitability map; and the pan and zoom functions are easily learned from more experienced users, as is the concept of GPS positioning.

Using the maps involves selecting a base map and optionally choosing a crop suitability map to overlay. A user would typically use the navigation map or GPS positioning to identify a location of interest and inspect the suitability for growing the different crops at this location. Changing the base map to landform is informative as to the terrain most suitable for particular crops; the user may then change the base map to world imagery and turn off the display of the crop suitability map to inspect the current land use at the location of interest.

The Leaflet controls do not prevent a user from displaying more than one crop suitability map at the same time, allowing for the possibility of user error and the potential for misinterpretation of results. However, the most serious limitation of these maps relates to internet connectivity. Access to the internet is critical for their usability as all geographical context is supplied by the

base maps. Without internet access the base maps do not display, and without this context it is not possible to meaningfully interpret the crop suitability maps.

6.5.3 Standalone Leaflet map

Although the electronic maps described above provide a simple and easy-to-use solution for disseminating the crop suitability maps, their dependence on the internet makes them impractical in situations where map users, due to issues of cost or location, do not have access to the internet. In this section an alternative Leaflet map solution is described that has all map content embedded within the application and does not require internet access for functionality.

The standalone Leaflet map makes use of the geographical content developed for the paper maps. Its design reverses the role of the base and overlay maps used in the web-based Leaflet maps - rather than displaying the crop suitability maps over base maps with geographic context, the crop maps are used as the base maps and geographic content is overlaid, as required, to provide context and insight.

The standalone Leaflet map offers seven base maps and three overlay maps. The first base map, for navigation, displays OpenStreetMap content on a hill-shade base layer. The second is the digital elevation model (DEM) showing landform, and the remaining five are the crop suitability maps. Only one base map can display at any one time, so the navigational and landform cues from the first two maps are lost when a crop map is displayed. To supply the necessary geographic context for the crop maps, three optional overlays are provided: (1) a simplified road map; (2) contour lines derived from the DEM showing 10 metre changes in altitude; and (3) drainage channels in the landscape (supplied by Gryphon Minerals Ltd).

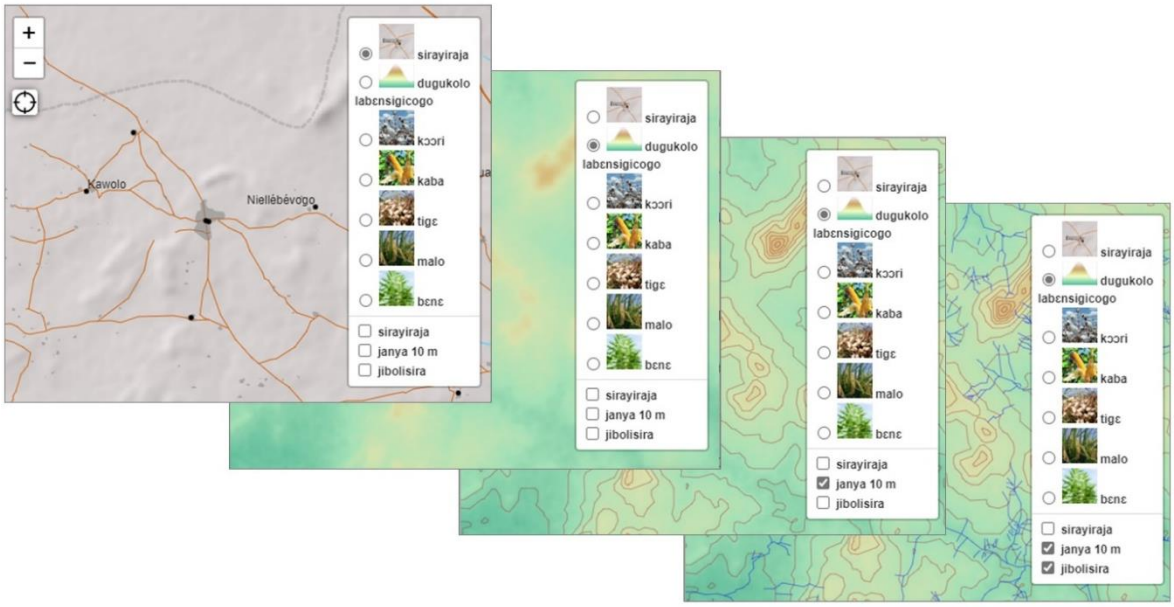

Figure 6-9 Detail from reference maps: (1) navigation map; (2) terrain map; (3) terrain overlaid with 10 metre contour lines; (4) terrain overlaid with contour lines and drainage channels

The two reference base maps are illustrated in Figure 6-9. The first image shows the navigation map (which contains the same map detail as the navigation map in Figure 6-4). The following three images show corresponding detail from the terrain map and demonstrate how the overlays

enhance the image by quantifying the changes in elevation and mapping the flow of water in the landscape.

Three crop suitability base maps are illustrated in Figure 6-10, with all three overlay maps visible. The closeness of contour lines indicates terrain steepness and the drainage channels mark areas of lower elevation; together they provide insight into the environmental niches of the individual crops. As only one crop map will display at a time, the possibility of user error that can occur in the web-based Leaflet maps is avoided.

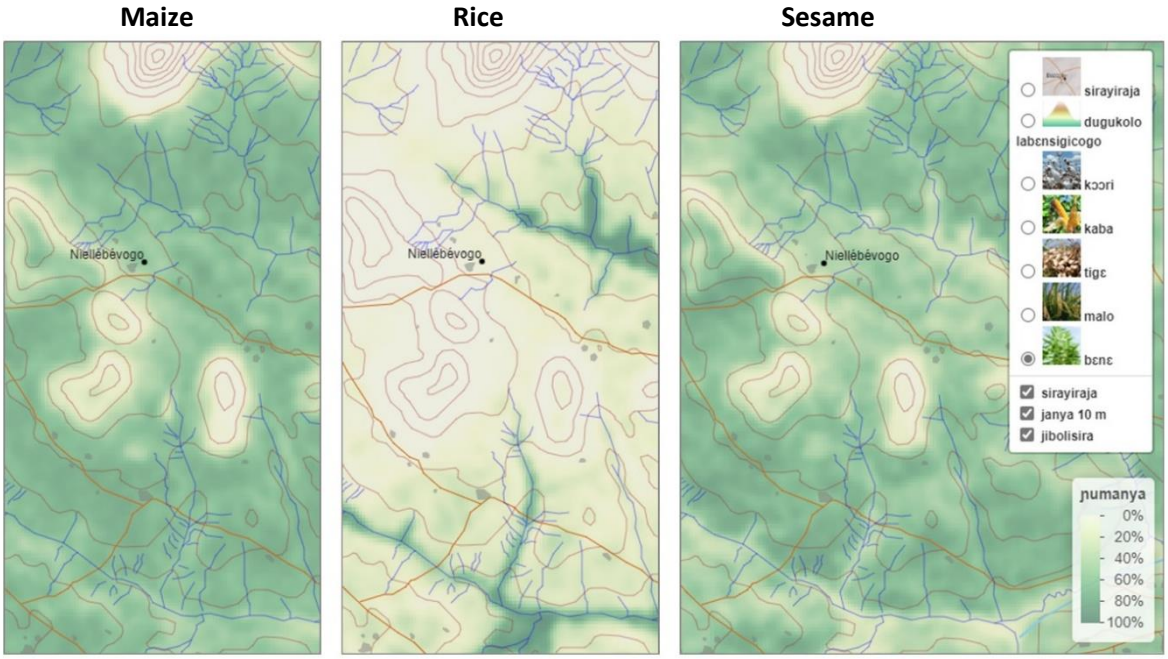

Figure 6-10 Crop suitability maps for Maize, Rice and Sesame with navigation detail, 10 metre contour lines and drainage channels displayed

A Dioula language version only of the standalone map was developed (for the translation of map text see Table 6-1). Images are used for each of the base maps in the map controls to assist illiterate map users, and the image for the terrain base map provides a legend for its symbology. Images were deemed unnecessary for the overlay maps as their content becomes readily understandable when used in combination with the reference base maps. A scale bar is provided on the interactive maps (not shown in the figures above); however, the altitude of contour lines is not marked as this detail would clutter most map views.

Table 6-1 Translation of Dioula text used in the standalone Leaflet map

Dioula	**English**	**Dioula**	**English**
kɔɔri	cotton	sirayiraja	road map
kaba	maize	dugukolo labɛnsigicogo	land-form/relief/topography
tigɛ	peanut	janya 10 m	altitude 10 m
malo	rice	jibolisira	drainage channel
bɛnɛ	sesame		

The ability to view high-resolution satellite imagery showing current land cover and land use is a valuable feature of the web-based Leaflet maps but was not implemented in the standalone version. Similar to the paper maps, satellite imagery proved ineffective and uninformative at all

resolutions that could feasibly have been used, and so was not included in the standalone Leaflet map.

6.6 Summary

This chapter explored mapping solutions appropriate for delivering the crop suitability maps from this project into the target communities. Socio-economic baseline data for the local region revealed very poor communities with low levels of education and literacy. Access to technology was not guaranteed and reliable internet access would be highly unlikely. For maps to be useful in these communities they must be easy to use and have no associated costs.

The main map content consists of the raster images for the project area showing predicted suitability for growing five main local crops. Additional map detail is necessary to provide geographical context and the most useful for this task is navigational features (roads and locations of towns and villages) and a physical terrain map to reveal landscape features.

Map text would be problematic for many target users who may be illiterate or have literacy in a different language. Making the maps language independent involved choosing intuitive map symbology wherever possible so that legends were unnecessary for these features, and designing clearly understandable text-free symbolic map legends when legends were necessary. Images were used in the suitability map titles to identify the relevant crops.

Traditional paper maps and electronic maps are complementary styles of map presentation that each have their own advantages and limitations for map design and map use. Both styles of map presentation have been explored for disseminating the crop suitability maps from this thesis.

Two paper map prototypes were designed for printing on A3 paper, but these maps could be printed at poster size for display and were still readable at A4 size if multiple copies were made as handouts. The first prototype fitted all map content on a single page, presenting multiple maps with common navigational cues to facilitate cognitive combination of the map content by the map users. The second prototype consisted of six full-page maps (terrain plus five crops) with reference detail overlaid on the raster images.

The two electronic map prototypes have pan and zoom functions that allow map users to inspect areas of interest in great detail. Both display in web browsers without the need to install any further software. One version displays semi-transparent suitability maps over base maps (navigation map, shaded relief, satellite imagery) and uses image tiles served from the internet. The second version is complementary in its design, allowing the map user to overlay navigational and terrain information on the suitability maps – this version is fully functional without internet access. Images are used in the map controls and any complementary map text is presented using the local language Dioula, thus demonstrating the language independence of the maps to any non-Dioula speaking map users.

Effective and free mapping solutions for this research task that are language independent and suitable for use by their target audience have been demonstrated in this chapter. Paper maps were designed for map users without access to other technology. Interactive maps were produced for map users with access to electronic devices, including a standalone version suitable for users without access to the internet and another web-based version that also allowed viewing of Earth observation imagery of the region, if the user had internet access.

Chapter 7 Conclusion

7.1 Context for research

Large scale surface mining projects generate extreme environmental and social impacts and create ongoing disruption to the livelihoods and way of life of their local communities. The establishment of a new mine often involves expropriation of land, relocation of communities and compensation payments. The FAO released its *Voluntary Guidelines on the Responsible Governance of Tenure* in 2012, with Section 16 specifically applying to expropriation and compensation (FAO 2012). The compensation guidelines were formulated to protect affected landholders from impoverishment and other risks. In examining the compliance of national-level legal provisions to the FAO guidelines, Tagliarino (2017) observed "*robust compensation procedures established by law, coupled with respect for the rule of law, can help ensure that expropriations promote sustainable development outcomes that balance property rights with the public interest*."

Compensation for compulsory land acquisition can take different forms. In countries with private land ownership (such as Australia) compensation is typically based on the market value of the property with additional payments to cover disturbance costs. In countries with customary land tenure (such as Burkina Faso) market value calculations do not apply. Instead, structural compensation is provided by building new villages to resettle displaced communities and crop compensation payments are made to farmers in recompense for economic displacement. Maps documenting community land use are critical for the crop compensation process. Tagoe, Mantey et al. (2012) describe the role of the land surveyor in the context of land acquisition for mining as being to accurately survey the farm, plot the farm polygons and generate an undisputable farm map known as the crop identification map from which compensation payments are calculated.

This research project was proposed in the context of the establishment of a large new industrial gold mine in south-west Burkina Faso that necessitated displacement and compensation for four farming communities. Some of the displaced people might be offered paid jobs in the mine, for others there would be new business opportunities from the access to capital and new markets, but most would need to re-establish existing livelihoods in new locations. The purpose of this research project was to establish whether the detailed data acquired during exploration could be repurposed for agricultural land suitability mapping with the goal of identifying potential farming land close to the planned mining complex to assist relocated farmers.

7.2 Summary of research

The thesis answered the question *"can land suitability assessments for agriculture be done effectively using the data by-products of mining exploration" (Research Question 1)* in the affirmative by demonstrating how this could be done, using a combination of the supplied and publicly available data. The data supplied by Gryphon Minerals Ltd was location specific, up-to-date (at the time) and finely detailed: the radiometric survey, crop compensation maps and updated soil map were created during the exploration phase and were augmented with purchased satellite images of the region. Some of the purchased imagery has since become freely available - the SRTM digital elevation model has been available for free download since 2014, and

there are now multiple sources of freely available high-resolution Earth observation imagery to use for qualitative assessment of model predictions.

The maps documenting community land use that were prepared to support the application for the mining licence for the Wahgnion Gold Project both facilitated and proved crucial to the agricultural suitability modelling undertaken in this project. Important local knowledge was derivable from the crop maps in that the recorded crops had existing markets and were proven suitable for the local climate and farming methods, and the environmental conditions supporting their cultivation were implicit in the recorded planting locations. These maps were a highly reliable source of known occurrences for crop species that could be used to derive presence points to train data-driven environmental niche modelling algorithms and so allow existing cropping patterns to inform suitability maps showing potential expansion areas for the crops *(Research Question 2)*.

The chosen methodology (maximum entropy modelling) employed a correlative species distribution modelling approach that linked the locations of known occurrences of a species with geospatial environmental data to derive the species-environment relationships characterising the fundamental ecological niche for that species. However, the tightly clustered presence samples used to train and test the models presented an inherent risk of sampling bias causing unrealistic predictions. The most commonly used metric for assessing species distribution models (area under the curve (AUC) for the receiver operator curve) was demonstrated to be misleading for this application and so other validation methods were devised to assess the quality of predictions. Multiple cross-validation models, that used spatially independent sets of training and test data, were compared to detect sampling bias, and Earth observation imagery was used for heuristic assessment of the feasibility of model predictions.

The thesis demonstrated that the soil categories used for mining exploration were ultimately useful as soil categories for the agricultural land suitability models, but the supplied categorical map of soil type was not effective as a model predictor in that format. Publicly available raster maps of predictions for multiple soil properties (SoilGrids) also proved ineffective due to their coarse spatial resolution. A new method was devised to process both sources of soil data into a set of raster layers that combined the soil categories and fine spatial detail of the supplied map with the multi-dimensionality of the soil property maps. These new soil rasters were effective as model predictors *(Research Question 3)*.

Region-wide validation of results was not possible with the available data and travel to the Burkina Faso project region was restricted, so the methodology was duplicated at two local sites in South Australia for which region-wide verification data were available. The South Australian results validated the methods used and demonstrated their transferability to other sites with different terrain, climate and styles of agriculture.

Finally, the thesis showed how land suitability assessment results could be delivered to subsistence farmers in remote locations by producing cost-free and language-independent mapping solutions that would overcome the local challenges of poverty, illiteracy and poor access to technology *(Research Question 4)*.

7.3 Research contributions

The thesis has demonstrated that species distribution modelling techniques, that are usually used to predict the geographic distribution of species for ecological purposes, can be applied to the task of land suitability assessment for agriculture. In this project the known planting locations

from crop maps produced for mining compensation were used to train the Maxent algorithm to predict the fundamental environmental niche for particular crops and so identify other potential planting locations for these crops.

The thesis demonstrated the use of alternative methods for detecting sampling bias and assessing the quality of model outputs. The occurrence data available for this project were not sampled across the target region but were spatially clumped in four small areas. The issue of likely sampling bias that was inherent in the configuration of the training data for the models was explicitly addressed throughout the project by the use of complementary techniques to assess the quality of the model outputs in terms of feasibility, accuracy and robustness. In particular, the commonly used AUC metric was shown to be an unreliable and misleading measure of model goodness for models trained on these data.

Importantly, the thesis presented a new method for converting categorically valued maps into continuously valued raster layers for use in environmental niche modelling (ENM). The default pixel values of a categorical raster are an unordered set of integers for the number of categories. ENM algorithms handle categorical predictors by decomposing their maps into a series of binary maps, creating one binary map for each category value, so exaggerating the differences between the classes and making no allowance for gradual transitions between classes or for mixed landscapes. In this project the values from SoilGrids rasters were used to assign meaningfully ordered values to soil type categories to create new soil property rasters for use as continuous model predictors – so overcoming the limitations of the binary data models imposed by the ENM algorithms for categorical data. Many available maps of potentially useful environmental predictors exist in the form of categorically valued polygon maps that partition the landscape into sets of mutually exclusive environmental classes – the techniques presented in this thesis have the potential to greatly improve the usefulness of such maps for ENM.

7.4 Limitations

The predictions for the fundamental environmental niches for the crops that were modelled in this project *cannot* be interpreted as recommended planting locations. Landscapes have potential for many different types of land use that offer differing and often conflicting ecosystem services. For example, from an agricultural perspective the wetlands in the inland valleys of African river systems are assumed to form the basis of robust production systems suitable for growing rice. However, from an ecological perspective, the wetlands are fragile ecosystems performing important ecological functions such as water purification, carbon sequestration, protection against flooding and erosion, and providing habitat for wildlife. The crop suitability maps produced by this project contribute to an understanding of the potential of the landscape for provisioning services (food production) only. The broader ecosystems involved and *all* services provided by them must be considered during land evaluation and planning to properly balance the agricultural and environmental benefits and impacts of any land use change.

7.5 Significance

This research project was undertaken in the context of gold mining under the Burkina Faso mining code; however, the methodologies developed during the project have the potential for far wider application. The methods in this thesis for predicting local agricultural land suitability are transferable to other sites with different terrain, climate, crops and styles of agriculture and are well suited to mining applications in developing countries where accurate records of local

agricultural activity (in the form of crop compensation maps) are produced as part of the environmental and social impact assessment.

The extractive industries must engage positively with local communities to ensure sustainable coexistence and mutually beneficial outcomes. Community engagement involves working collaboratively to address issues affecting the social well-being of people affected by a mining project. It requires identifying the most critical challenges facing communities and then asking how the resources and competencies of the mining project can be used to help overcome these challenges (Baba, Mohammad et al. 2021). In relation to expropriation of farming land, affected farmers have an urgent and critical need to find alternative suitable land to continue their livelihoods. The thesis has demonstrated how existing mining exploration data and spatial technologies can be used to perform agricultural land suitability modelling to support local agricultural planning and so assist relocated farmers to obtain better agricultural outcomes.

The reciprocity framework for community engagement in the extractive industries focuses on how mining projects can contribute to the capacities of *individual* community members who are affected by projects (Baba, Mohammad et al. 2021). In this thesis, the research results are presented using language-independent maps that would be accessible by the local population to enable individuals to use the results for their own decision making.

In summary, this process of reusing of mining exploration data for land management purposes and presenting the results in a manner accessible by local populations could become a model for future mining projects and contribute to more successful collaborations between the mining sector and local communities globally.

References

Ahmad, F., L. Goparaju and A. Qayum (2017). "Agroforestry suitability analysis based upon nutrient availability mapping: a GIS based suitability mapping." AIMS Agriculture and Food **2**(2): 201-220.

Al Jazeera English. (2012). "Burkina Faso children toil in gold mines." Retrieved 14 November 2016, from https://www.youtube.com/watch?v=sEtTtVzYDKs.

Al Jazeera English. (2014). "Miners risk death in Burkina Faso gold rush." Retrieved 14 November 2016, from https://www.youtube.com/watch?v=kJfU-2Twraw.

Attua, E. M. and J. B. Fisher (2010). "Land Suitability Assessment for Pineapple Production in the Akwapim South District, Ghana: .A GIS-MultiCriteria Approach." Ghana Journal of Geography **2**: 37.

Baba, S., S. Mohammad and C. Young (2021). "Managing project sustainability in the extractive industries: Towards a reciprocity framework for community engagement." International Journal of Project Management **39**(8): 887-901.

Bao, B. (2016). Design and evaluation of text-free map interfaces : a thesis presented in partial fulfilment of the requirements for the degree of Master of Information Technology at Massey University, Albany, New Zealand. Master of Information Science (MInfSc) Masters thesis, Massey University.

Bolster, S. J. S. (1999). Regolith Mapping: Is It Really Necessary? AEG-AIG-GSA Conference, Perth, Australian Institute of Geoscientists.

Boria, R. A., L. E. Olson, S. M. Goodman and R. P. Anderson (2014). "Spatial filtering to reduce sampling bias can improve the performance of ecological niche models." Ecological Modelling **275**: 73-77.

Brenning, A., D. Bangs, M. Becker, P. Schratz and F. Polakowski (2018). Package 'RSAGA'. (Version 1.3.0). CRAN. https://cran.r-project.org/web/packages/RSAGA/RSAGA.pdf.

Canada Science and Technology Museum. (2017). "Social License to Operate - Jim Cooney." Retrieved 21 January, 2022, from https://www.youtube.com/watch?v=NkQMq0gIEYU.

Cohen, J. (1960). "A coefficient of agreement for nominal scales." Educational and psychological measurement **20**(1): 37-46.

Conrad, O., B. Bechtel, M. Bock, H. Dietrich, E. Fischer, L. Gerlitz, J. Wehberg, V. Wichmann and J. Böhner (2015). "System for Automated Geoscientific Analyses (SAGA) v. 2.1.4." Geosci. Model Dev. **8**(7): 1991-2007.

Côte, M. (2016, 25 March 2016). "After the gold rush: exploration on the permanent mining frontier of Burkina Faso." Retrieved 26 October, 2016, from https://politicalecologynetwork.com/2016/04/26/after-the-gold-rush-exploration-on-the-permanent-mining-frontier-of-burkina-faso-2/.

Cruz-Cárdenas, G., L. López-Mata, J. L. Villaseñor and E. Ortiz (2014). "Potential species distribution modeling and the use of principal component analysis as predictor variables." Revista Mexicana de Biodiversidad **85**(1): 189-199.

Dare, M., J. Schirmer and F. Vanclay (2014). "Community engagement and social licence to operate." Impact Assessment and Project Appraisal **32**(3): 188-197.

De Longueville, F., Y.-C. Hountondji, I. Kindo, F. Gemenne and P. Ozer (2016). "Long-term analysis of rainfall and temperature data in Burkina Faso (1950–2013)." International Journal of Climatology **36**(13): 4393-4405.

Dowd-Uribe, B., C. Roncoli and B. Orlove (2012). "Water Grows Food: Dry Season Farming, Food Sovereignty, and Integrated Water Resource Management in Burkina Faso." Special Bulletin Editors: Noah Zerbe and Brian Dowd-Uribe Bulletin Editors: Timothy Scarnecchia and Peter Limb: 11.

Dudík, M., S. J. Phillips and R. E. Schapire (2004). Performance Guarantees for Regularized Maximum Entropy Density Estimation, Berlin, Heidelberg, Springer Berlin Heidelberg.

Elith, J., S. J. Phillips, T. Hastie, M. Dudik, Y. E. Chee and C. J. Yates (2011). "A statistical explanation of MaxEnt for ecologists." Diversity and Distributions **17**(1): 43.

Elsheikh, R., A. Shariff, F. Amiri, N. B. Ahmad, S. K. Balasundram and M. A. M. Soom (2013). "Agriculture Land Suitability Evaluator (ALSE): A decision and planning support tool for tropical and subtropical crops." Computers and Electronics in Agriculture **93**: 98-110.

Elsheikh, R., A. R. B. M. Shariff, F. Amiri, N. B. Ahmad, S. K. Balasundram and M. A. M. Soom (2013). "Agriculture Land Suitability Evaluator (ALSE): A decision and planning support tool for tropical and subtropical crops." Computers and Electronics in Agriculture **93**: 98-110.

Engels, B. (2014). "Beyond scarcity: conflicts over land and social relations in south-western Burkina Faso." Modern Africa: politics, history and society **2**(1): 75-93.

Estes, L. D., B. A. Bradley, H. Beukes, D. G. Hole, M. Lau, M. G. Oppenheimer, R. Schulze, M. A. Tadross and W. R. Turner (2013). "Comparing mechanistic and empirical model projections of crop suitability and productivity: implications for ecological forecasting." Global Ecology and Biogeography **22**(8): 1007-1018.

FAO (1976). A framework for land evaluation. Rome, Food and Agriculture Organization of the United Nations. Available at http://www.fao.org/docrep/X5310E/x5310e00.HTM. Accessed 6 October 2022.

FAO (2007). Land evaluation Towards a revised framework. Rome, Food and Agriculture Organization of the United Nations. Available at https://mylupis.xoom.it/doc/Pub/LandEvaluation.pdf. Accessed 6 October 2022.

FAO (2012). Voluntary Guidelines on the responsible Governance of tenure of land, fisheries and forests in the Context of national food security. Rome, Food and Agriculture Organization of the United Nations. Available at http://www.fao.org/docrep/016/i2801e/i2801e.pdf. Accessed 6 October 2022.

Gallant, J. C. and T. I. Dowling (2003). "A multiresolution index of valley bottom flatness for mapping depositional areas." Water Resources Research **39**(12).

Garrity, D. P., F. K. Akinnifesi, O. C. Ajayi, S. G. Weldesemayat, J. G. Mowo, A. Kalinganire, M. Larwanou and J. Bayala (2010). "Evergreen Agriculture: a robust approach to sustainable food security in Africa." Food security **2**(3): 197-214.

George, H. (2005). An overview of land evaluation and land use planning at FAO. Rome, Food and Agriculture Organization of the United Nations. Available at https://nanopdf.com/download/land-evaluation-at-fao-food-and-agriculture-organization-of-the_pdf. Accessed 6 October 2022.

Gongo, S. and P. Bax. (2016, 30 August 2016). "Africa's Fastest Growing Gold Producer Says More to Come." Retrieved 30 November 2016, from https://www.bloomberg.com/news/articles/2016-08-31/africa-s-fastest-growing-gold-producer-says-more-to-come.

Griffin, A. L., T. White, C. Fish, B. Tomio, H. Huang, C. R. Sluter, J. V. M. Bravo, S. I. Fabrikant, S. Bleisch, M. Yamada and P. Picanço (2017). "Designing across map use contexts: a research agenda." International Journal of Cartography **3**(sup1): 90-114.

Hahn, H. and L. Kibora (2008). "The Domestication of the Mobile Phone: Oral Society and New ICT in Burkina Faso." The Journal of Modern African Studies **46**: 87-109.

Hall, J., D. Maschmedt and B. Billing (2009). The soils of southern South Australia. Adelaide, Department of Water, Land and Biodiversity Conservation.

Hallgren, W., L. Beaumont, A. Bowness, L. Chambers, E. Graham, H. Holewa, S. Laffan, B. Mackey, H. Nix, J. Price, J. Vanderwal, R. Warren and G. Weis (2016). "The Biodiversity and Climate Change Virtual Laboratory: Where ecology meets big data." Environmental Modelling & Software **76**: 182-186.

Hengl, T., J. Mendes de Jesus, G. B. M. Heuvelink, M. Ruiperez Gonzalez, M. Kilibarda, A. Blagotić, W. Shangguan, M. N. Wright, X. Geng, B. Bauer-Marschallinger, M. A. Guevara, R. Vargas, R. A. MacMillan, N. H. Batjes, J. G. B. Leenaars, E. Ribeiro, I. Wheeler, S. Mantel and B. Kempen (2017). "SoilGrids250m: Global gridded soil information based on machine learning." PLOS ONE **12**(2): e0169748.

Heumann, B. W., S. J. Walsh and P. M. McDaniel (2011). "Assessing the application of a geographic presence-only model for land suitability mapping." Ecological Informatics **6**(5): 257-269.

Hijmans, R. J. (2012). "Cross-validation of species distribution models: removing spatial sorting bias and calibration with a null model." Ecology **93**(3): 679-688.

IAEA (2003). Guidelines for Radioelement Mapping Using Gamma Ray Spectrometry Data. Vienna, International Atomic Energy Agency.

IETI (2014). 2012 Burkina Faso EITI Report. Burkina Faso EITI. Available at https://eiti.org/document/2012-burkina-faso-eiti-report. Accessed 6 October 2022.

IFC (2012). Performance Standards on Environmental and Social Sustainability. International Finance Corporation. Available at www.ifc.org/performancestandards. Accessed 6 October 2022.

IFC (2016). IFC The First Six Decades : Leading the Way in Private Sector Development--A History. Washington, DC, International Finance Corporation. Available at https://openknowledge.worldbank.org/handle/10986/25402. Accessed 6 October 2022.

Ingram, K. T., M. C. Roncoli and P. H. Kirshen (2002). "Opportunities and constraints for farmers of west Africa to use seasonal precipitation forecasts with Burkina Faso as a case study." Agricultural Systems **74**(3): 331-349.

Intersocial Consulting (2014). Banfora Gold Project Social Impact Assessment. Société Minière Gryphon SA.

Jiménez-Valverde, A. (2012). "Insights into the area under the receiver operating characteristic curve (AUC) as a discrimination measure in species distribution modelling." Global Ecology and Biogeography **21**(4): 498-507.

Jones, A., H. Breuning-Madsen, M. Brossard, A. Dampha, J. Deckers, O. Dewitte, T. Gallali, S. Hallett, R. Jones, M. Kilasara, P. Le Roux, E. Micheli, L. Montanarella, O. Spaargaren, T. Lamourdia, E. Van Ranst, M. Yemefack and R. Zougmore (2013). Soil Atlas of Africa.

Jordan, G. (2007). Digital Terrain Analysis in a GIS Environment. Concepts and Development. Digital Terrain Modelling: Development and Applications in a Policy Support Environment. R. J. Peckham and G. Jordan. Berlin, Heidelberg, Springer Berlin Heidelberg**:** 1-43.

Kemp, K. K. (2012). The Hawaii Island Crop Probability Map: An Update of the Crop Growth Parameters for the Hawaii County Crop Model. County of Hawai`i Research and Development Department. Available at https://spatial.usc.edu/wp-content/uploads/2014/03/gislabtr13b.pdf. Accessed 6 October 2022.

Kiepe, P. (2006). Characterization of three key environments for integrated irrigation-aquaculture and their local names. Integrated irrigation and aquaculture in West Africa: concepts, practices and potential. M. Halwart and A. A. van Dam. Rome, FAO.

Kimani, M. (2009) "Mining to profit Africa's people: Governments bargain for 'fair deals' that enhance development." African Renewal Online **April 2009**. http://www.un.org/africarenewal/magazine/april-2009/mining-profit-africa%E2%80%99s-people.

Komnitsas, K. (2020). "Social License to Operate in Mining: Present Views and Future Trends." Resources Conservation and Recycling **9**(79).

Konate, M. (2016). Personal conversation, Adelaide, South Australia.

Kone, A. a. (2010). "Politics of Language: The Struggle for Power in Schools in Mali and Burkina Faso." International Education **39**(2): 6-20.

Kramer-Schadt, S., J. Niedballa, J. D. Pilgrim, B. Schröder, J. Lindenborn, V. Reinfelder, M. Stillfried, I. Heckmann, A. K. Scharf, D. M. Augeri, S. M. Cheyne, A. J. Hearn, J. Ross, D. W. Macdonald, J. Mathai, J. Eaton, A. J. Marshall, G. Semiadi, R. Rustam, H. Bernard, R. Alfred, H. Samejima, J. W. Duckworth, C. Breitenmoser-Wuersten, J. L. Belant, H. Hofer and A. Wilting (2013). "The importance of correcting for sampling bias in MaxEnt species distribution models." Diversity and Distributions **19**(11): 1366-1379.

Kumah, A. (2006). "Sustainability and gold mining in the developing world." Journal of Cleaner Production **14**(3–4): 315-323.

Liu, H., Q. Zhan and M. Zhan (2017). The Uncertainties on the GIS-Based Land Suitability Assessment for Urban and Rural Planning ISPRS Geospatial Week 2017. Wuhan, China, The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. **XLII-2/W7**.

Lobo, J. M., A. Jiménez-Valverde and R. Real (2008). "AUC: a misleading measure of the performance of predictive distribution models." Global Ecology and Biogeography **17**(2): 145-151.

Luning, S. (2008). "Liberalisation of the Gold Mining Sector in Burkina Faso." Review of African Political Economy **35**(117): 387-401.

Luning, S. (2012). "Processing promises of gold: a minefield of company-community relations in Burkina Faso." Africa today **58**(3): 22-39.

Malczewski, J. (2004). "GIS-based land-use suitability analysis: a critical overview." Progress in Planning **62**(1): 3-65.

Maschmedt, D. J. (2002). Assessing Agricultural Land. South Australia, Department of Water Land and Biodiversity Conservation. Available at https://data.environment.sa.gov.au/Content/Publications/Assessing-Agricultural-Lands.pdf. Accessed 6 October 2022.

Mattivi, P., F. Franci, A. Lambertini and G. Bitelli (2019). "TWI computation: a comparison of different open source GISs." Open Geospatial Data, Software and Standards **4**(1): 6.

McHugh, M. L. (2012). "Interrater reliability: the kappa statistic." Biochemia medica **22**(3): 276-282.

McMurray, A. J., R. W. Pace and D. Scott (2004). Research: a commonsense approach. Victoria, Australia, Thomson Social Science Press.

Medhi, I., A. Sagar and K. Toyama (2006). Text-Free User Interfaces for Illiterate and Semi-Literate Users.

Mercer-Mapstone, L., W. Rifkin, W. Louis and K. Moffat (2017). "Meaningful dialogue outcomes contribute to laying a foundation for social licence to operate." Resources Policy **53**: 347-355.

Merow, C., M. J. Smith and J. A. Silander, Jr (2013). "A practical guide to MaxEnt for modeling species' distributions: what it does, and why inputs and settings matter." Ecography **30**(10): 1058-1069.

Mighty, M. A. (2015). "Site suitability and the analytic hierarchy process: How GIS analysis can improve the competitive advantage of the Jamaican coffee industry." Applied Geography **58**: 84-93.

Mining technology. (2018). "Wahgnion Gold Project." Retrieved 27 August, 2019, from https://www.mining-technology.com/projects/wahgnion-gold-project/.

Moffat, K. and A. Zhang (2014). "The paths to social licence to operate: An integrative model explaining community acceptance of mining." Resources Policy **39**: 61-70.

Morales, N. S., I. C. Fernández and V. Baca-González (2017). "MaxEnt's parameter configuration and small samples: are we paying attention to recommendations? A systematic review." PeerJ **5**.

Müller, M. J. (1982). Selected climatic data for a global set of standard stations for vegetation science. The Hague ; Boston, Dr. W. Junk Publishers : Hingham Distributors for the U.S. and Canada, Kluwer Boston.

NALRC (2010). Fulfude, Moore & Dioula. Indiana University, National African Language Resource Center. Available at https://nalrc.indiana.edu/doc/brochures/fulfude-moore-diaula.pdf. Accessed 6 October 2022.

National Research Council (1996). Lost Crops of Africa: Volume I: Grains. Washington, DC, The National Academies Press.

Nature Methods (2011). "Ground-truth data cannot do it alone." Nature Methods **8**(11): 885-885.

Peterson, A. T., J. Soberón, R. G. Pearson, R. P. Anderson, E. Martínez-Meyer, M. Nakamura and M. B. Araújo (2011). Ecological Niches and Geographic Distributions (MPB-49), Princeton University Press.

Phillips, S. J. (2017). A Brief Tutorial on Maxent AT&T Research. Available at https://biodiversityinformatics.amnh.org/open_source/maxent/Maxent_tutorial2017.pdf. Accessed 6 October 2022.

Phillips, S. J., R. P. Anderson, M. Dudík, R. E. Schapire and M. E. Blair (2017). "Opening the black box: an open-source release of Maxent." Ecography **40**(7): 887-893.

Phillips, S. J., R. P. Anderson and R. E. Schapire (2006). "Maximum entropy modeling of species geographic distributions." Ecological Modelling **190**(3): 231-259.

Phillips, S. J. and M. Dudík (2008). "Modeling of Species Distributions with Maxent: New Extensions and a Comprehensive Evaluation." Ecography **31**(2): 161-175.

Phillips, S. J., M. Dudík and R. E. Schapire (2004). A maximum entropy approach to species distribution modeling. Proceedings of the twenty-first international conference on Machine learning. Banff, Alberta, Canada, Association for Computing Machinery**:** 83.

Pike, R. J., I. S. Evans and T. Hengl (2009). Chapter 1 Geomorphometry: A Brief Guide. Developments in Soil Science, Elsevier**:** 3-30.

Poushter, J. (2016). Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies. Pew Research Center. Available at https://www.pewresearch.org/global/2016/02/22/smartphone-ownership-and-internet-usage-continues-to-climb-in-emerging-economies/. Accessed 6 October 2022.

Prno, J. and D. S. Slocombe (2012). "Exploring the origins of 'social license to operate' in the mining sector: Perspectives from governance and sustainability theories." Resources Policy **37**(3): 346-357.

Pulse Australia. (2016). "Chickpea Production: Northern Region." Retrieved 7 February, 2022, from https://www.pulseaus.com.au/growing-pulses/bmp/chickpea/northern-guide.

Radosavljevic, A. and R. P. Anderson (2014). "Making better Maxent models of species distributions: complexity, overfitting and evaluation." Journal of Biogeography **41**(4): 629-643.

Rodenburg, J., S. J. Zwart, P. Kiepe, L. T. Narteh, W. Dogbe and M. C. S. Wopereis (2014). "Sustainable rice production in African inland valleys: Seizing regional potentials through local approaches." Agricultural Systems **123**: 1-11.

Rossiter, D. G. (1990). "ALES: A framework for land evaluation using a microcomputer." Soil use and management **6**(1): 7-20.

Rossiter, D. G. (1996). "A theoretical framework for land evaluation." Geoderma **72**(3-4): 165-190.

Rowland, J., D. Maschmedt and C. Liddicoat (2016). Land use potential for agricultural crops in southern South Australia: Summary of assessment and mapping methodology. Adelaide, South Australia, Department of Environment Water and Natural Resources. Available at https://data.environment.sa.gov.au/Content/Publications/LandUsePotential_Descriptions_MappingAndSpatialData.pdf. Accessed 6 October 2022.

Rupley, L., L. Bangali and B. Diamitani (2013). Historical Dictionary of Burkina Faso, Scarecrow Press.

Schueler, V., T. Kuemmerle and H. Schröder (2011). "Impacts of Surface Gold Mining on Land Use Systems in Western Ghana." AMBIO **40**(5): 528-539.

Shannon, C. E. and W. Weaver (1949). The Mathematical Theory of Communication, University of Illinois Press.

Sonter, L. J., C. J. Moran, D. J. Barrett and B. S. Soares (2014). "Processes of land use change in mining regions." Journal of Cleaner Production **84**: 494-501.

Tagliarino, N. K. (2017). "The Status of National Legal Frameworks for Valuing Compensation for Expropriated Land: An Analysis of Whether National Laws in 50 Countries/Regions across Asia, Africa, and Latin America Comply with International Standards on Compensation Valuation." Land **6**(2): 37.

Tagoe, N. D., S. Mantey, S. Adjei and M. Soakodan (2012). The Role of the Land Surveyor in Land Acquisition and Compensation–A Case Study of the Tarkwa Mining Communities, Ghana. FIG Working Week 2012, Rome, Italy, International Federation of Surveyors (FIG).

Thompson, I. (1999, January 1999). "A social license to operate: essential for success in international exploration." BC&Yukon Chamber of Mines Cordilleran Roundup Retrieved 24 January, 2022, from https://socialicense.com/publications/ROUNDUP1999%20SLO.pdf.

Tincani, L. S. (2012). Resilient Livelihoods: Adaptation, Food Security and Wild Foods in Rural Burkina Faso. Ph D Doctoral dissertation, University of London, School of Oriental and African Studies.

UNESCO Institute of Statistics. (2022). "Burkina Faso." Retrieved 2 May, 2022, from http://uis.unesco.org/en/country/bf?theme=education-and-literacy.

USAID (2017). USAID Country Profile: Property Rights and Resource Governance Burkina Faso. United States Agency for International Development. Available at https://land-links.org/wp-content/uploads/2010/08/USAID_Land_Tenure_Burkina_Faso_Country_Profile.pdf. Accessed 6 October 2022.

USAID. (2021). "BURKINA FASO AGRICULTURE AND FOOD SECURITY." Retrieved 19 January, 2022, from https://www.usaid.gov/burkina-faso/agriculture-and-food-security.

Ventura, L. (2022). "The World's Richest and Poorest Countries 2021." Retrieved 19/1/2022, from https://www.gfmag.com/global-data/economic-data/worlds-richest-and-poorest-countries.

Via Campesina (2007). Declaration of Nyéléni. Forum for Food Sovereignty, Sélingué. Mali.

Wang, L. and H. Liu (2006). "An efficient method for identifying and filling surface depressions in digital elevation models for hydrologic analysis and modelling." International Journal of Geographical Information Science **20**(2): 193-213.

Warren, D., R. Glor and M. Turelli (2010). "ENMTools: A toolbox for comparative studies of environmental niche models." Ecography **33**: 607-611.

Westwood, M. N. (1988). Temperate-zone pomology. Portland, Or, Timber Press.

Wiley, E. O., K. M. McNyset, A. T. Peterson, C. R. Robins and A. M. Stewart (2003). "Niche Modeling Perspective on Geographic Range Predictions in the Marine Environment Using a Machine-learning Algorithm." Oceanography **16**.

Williams, P. K. (2016). Personal conversation, Adelaide, South Australia.

Woodhouse, I. (2021). "On 'ground' truth and why we should abandon the term." Journal of Applied Remote Sensing **15**(4): 041501.

World Bank Group. (2022). "GDP per capita (current US$) - Burkina Faso." Retrieved 19 January, 2022, from https://data.worldbank.org/indicator/NY.GDP.PCAP.CD?locations=BF.

Appendices

Appendix A – Mining exploration data supplied for the project

Description	Date	Data type	Coverage
Community land use (x4)	2014	Shape files (ESRI)	Within mining lease
Regional soil map	2013	Shape file (ESRI)	Project area
Drainage channels	2013	Shape files (Mapinfo)	Project area
Regolith	2011	Raster tileset (TNTmips)	Project area
Radiometrics (K, Th, U)	2010 (Apr)	Raster (25m resolution)	Project area
ALOS satellite image – DSM	2011 (Jul 29)	Raster (5m resolution)	c. 85% project area (north)
Contour lines – 2m, 5m	2010	Shape files (Mapinfo)	c. 40% project area (south)
Digital Elevation Model	2010 2012	Raster (1m resolution)	c. 40% project area (south) + c. 10% project area (north)
Landsat 7 satellite image – 9 bands	2003 (Feb 8)	Raster (30m resolution)	Project area
SRTM satellite image – DSM	2000	Raster (30m resolution)	Project area
Worldview2 satellite image – 4 bands	2010 (Apr)	Raster (0.5m resolution)	Project area
Worldview2 satellite image – 8 bands	2010 (Apr)	Raster (2m multispectral)	c. 8% project area
Worldview2 satellite image – 4 bands	2012 (Apr)	Raster (0.5m panchromatic, 2m multispectral)	c. 70% project area (includes most of mining lease)
Worldview2 satellite image – 8 bands	2013	Raster (0.5m panchromatic, 2m multispectral)	c. 60% project area (includes mining lease)
Worldview2 satellite image – 4 bands	2014 (Nov)	Raster (0.5m panchromatic, 2m multispectral)	c. 25% project area (includes mining lease)

Appendix B – Land and crop analysis of mining lease communities

	FOURKOURA	NOGBELE	SAMAVOGO	STINGER	Total
Area (hectares)	286.0	1382.9	547.6	340.8	2557.3
% total area	11.2%	54.1%	21.4%	13.3%	*100.0%*
Plots (number)	308	1186	434	302	2230
% total plots	13.8%	53.2%	19.5%	13.5%	*100.0%*
Average plot size (hectares)	0.93	1.17	1.26	1.13	1.15
Land ownership (% area)					
-	4.9%	3.8%	3.9%	2.1%	3.8%
Owner and user	87.4%	93.9%	64.8%	89.1%	86.3%
Users only	7.6%	2.2%	31.3%	8.8%	9.9%
Total	*100.0%*	*100.0%*	*100.0%*	*100.0%*	*100.0%*
Cultivation state (% area)					
-	4.9%	3.7%	3.9%	2.1%	3.7%
In preparation	0.0%	0.0%	0.5%	0.0%	0.1%
Cultivated	60.8%	64.8%	58.6%	70.2%	63.8%
Harvest	0.8%	7.3%	1.1%	1.2%	4.4%
Fallow < 2 years	0.7%	1.1%	4.3%	0.5%	1.7%
Fallow > 2 years	30.9%	18.6%	21.3%	25.4%	21.4%
Never cultivated	1.9%	4.4%	9.4%	0.2%	4.7%
Uncultivable	0.0%	0.0%	0.8%	0.3%	0.2%
Total	*100.0%*	*100.0%*	*100.0%*	*100.0%*	*100.0%*
Type of soil (% area)					
-	4.9%	3.8%	3.9%	2.1%	3.8%
Bas-fonds	4.2%	5.4%	4.0%	3.5%	4.7%
Gravelly	73.9%	71.9%	88.6%	82.0%	77.1%
Sandy	1.2%	14.7%	3.4%	12.0%	10.4%
Clayey	15.8%	4.1%	0.0%	0.3%	4.1%
Total	*100.0%*	*100.0%*	*100.0%*	*100.0%*	*100.0%*
Amenities (% area)					
-	4.9%	3.8%	3.9%	2.1%	3.8%
Irrigation	0.0%	0.4%	0.0%	0.6%	0.3%
Without amenities	91.8%	95.5%	96.1%	97.2%	95.4%
Stoney cords	3.2%	0.2%	0.0%	0.0%	0.5%
Total	*100.0%*	*100.0%*	*100.0%*	*100.0%*	*100.0%*

	FOURKOURA	NOGBELE	SAMAVOGO	STINGER	Total
Area (hectares)	286.0	1382.9	547.6	340.8	2557.3
Maize	73.6	346.6	125.2	70.2	615.6
Cotton	30.9	368.0	44.5	67.5	510.9
Sesame	40.9	84.6	91.4	52.3	269.1
Peanut	9.1	44.3	52.1	17.2	122.8
Rice	12.7	52.5	31.4	6.7	103.2
Millet	5.9	38.5	21.8	19.8	86.0
Cashew	0.0	23.7	37.8	3.9	65.5
Sorghum	0.9	2.2	54.0	3.6	60.7
Cowpea	0.9	7.3	26.7	1.6	36.5
Earth pea	2.9	3.7	24.2	3.3	34.1
Beans	6.7	15.0	0.2	5.3	27.3
Potato/Yam	3.4	2.3	0.8	1.3	7.9
Fallow	*94.7*	*318.9*	*179.1*	*90.8*	*683.5*
Number of plots	308	1186	434	302	2230
Maize	50	194	57	45	346
Cotton	25	157	13	29	224
Sesame	20	91	45	51	207
Peanut	23	61	30	18	132
Rice	91	329	199	55	674
Millet	7	41	7	16	71
Cashew		18	12	2	32
Sorghum	3	5	27	7	42
Cowpea	5	17	15	4	41
Earth pea	7	8	18	6	39
Beans	17	34	1	12	64
Potato/Yam	19	14	2	5	40
Fallow	*44*	*164*	*56*	*50*	*314*
Average plot size (hectares)	0.93	1.17	1.26	1.13	1.15
Maize	1.47	1.79	2.20	1.56	1.78
Cotton	1.24	2.34	3.42	2.33	2.28
Sesame	2.04	0.93	2.03	1.02	1.30
Peanut	0.40	0.73	1.74	0.95	0.93
Rice	0.14	0.16	0.16	0.12	0.15
Millet	0.84	0.94	3.11	1.24	1.21
Cashew		1.32	3.15	1.95	2.05
Sorghum	0.29	0.44	2.00	0.51	1.44
Cowpea	0.19	0.43	1.78	0.40	0.89
Earth pea	0.41	0.46	1.34	0.56	0.87
Beans	0.39	0.44	0.23	0.44	0.43
Potato/Yam	0.18	0.17	0.41	0.27	0.20
Fallow	*2.15*	*1.94*	*3.20*	*1.82*	*2.18*

Appendix C – Geographical location of test sites

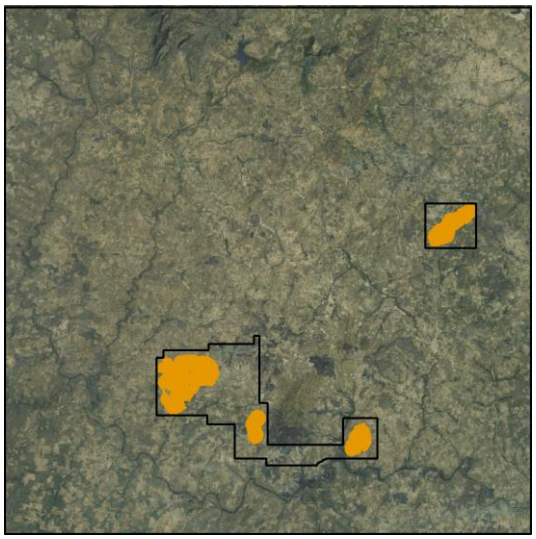

Site B – Burkina Faso project area

Size: 40km x 40km

Coordinates

UTM Zone 30P Northern Hemisphere

Eastings: 224,010 to 264,000

Northings: 1,136,010 to 1,176,000

Longitude: -5.51966904 to -5.1571874

-5° 31' 10.8084" to -5° 9' 25.8732"

Latitude: 10.26686104 to 10.63100622

10° 16' 0.6996" to 10° 37' 51.621"

WRS1: Path=212, Row=53

WRS2: Path=197, Row=53

EPSG: 32630

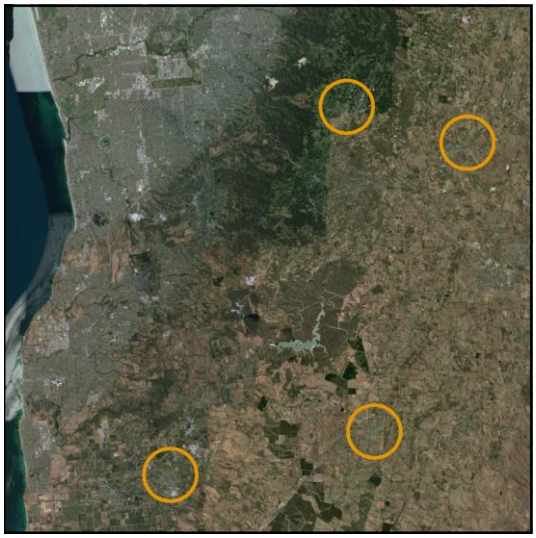

Site A – Adelaide project area

Size: 40km x 40km

Coordinates

UTM Zone 54H Southern Hemisphere

Eastings: 268,010 to 308,000

Northings: 6,097,010 to 6,137,000

Longitude: 138.45029 to 138.89887

138° 27' 1.0" to 138° 53' 55.9"

Latitude: -35.24326 to -34.89127

-35° 14’ 35.7" to -34° 53’ 28.6”

WRS1: Path=103, Row=84

WRS2: Path=97, Row=84

EPSG: 32754

Site M – Marrabel project area

Size: 40km x 40km

Coordinates

UTM Zone 54H Southern Hemisphere

Eastings: 285,010 to 325,000

Northings: 6,200,010 to 6,240,000

Longitude: 138.66342 to 139.10588

138° 39' 48.3" to 139° 6' 21.2"

Latitude: -34.31892 to -33.96598

-34° 19’ 8.1" to -33° 57’ 57.5”

WRS1: Path=103, Row=84?

WRS2: Path=97, Row=84?

EPSG: 32754

Appendix D – SoilGrids250m data

SoilGrids250m layers

Filenames: *XXXXXX*_M_sl*N*_250m.tif, where *XXXXXX* is the file and *N* is the level for seven standard depths: (1) 0cm ; (2) 5cm; (3) 15cm; (4) 30cm; (5) 60cm; (6) 100cm; (7) 200cm

Soil quality	File	Description	Units	Range	Levels
Acid grade	ACDWRB	Grade of a sub-soil being acid e.g. having a pH < 5 and low BS	grade	0-5	-
Bedrock depth	BDRICM	Depth to bedrock (R horizon) up to 200cm	cm	0-200	-
Bedrock depth	BDRLOG	Probability of occurrence (0-100%) of R horizon	percent	0-100	-
Bedrock depth	BDTICM	Absolute depth to bedrock (in cm)	cm	0-5000	-
Bulk density	BLDFIE	Bulk density (fine earth) in kg / cubic-meter	kg / cubic-m	50-2650	1-7
Carbon - density	OCDENS	Soil organic carbon density in kg per cubic-m	kg / cubic-m	0-1000	1-7
Carbon - organic	ORCDRC	Soil organic carbon content (fine earth fraction) in g per kg	g / kg	0-800	1-7
Carbon - stock	OCSTHA	Soil organic carbon stock in tons per ha	tonnes / ha	0-3000	1-7, 0-0.3m, 0-1m, 0-2m
Cation capacity	CECSOL	Cation exchange capacity of soil in cmolc/kg	cmol / kg	0-2200	1-7
pH acidity	PHIHOX	Soil pH x 10 in water	index*10	20-110	1-7
pH KCl	PHIKCL	Soil pH x 10 in Potassium chloride	index*10	20-110	1-7
Texture - clay	CLYPPT	Clay content (0-2 micrometre) mass fraction in %	percent	0-100	1-7
Texture - coarse	CRFVOL	Coarse fragments volumetric in %	percent	0-100	1-7
Texture - sand	SNDPPT	Sand content (50-2000 micrometre) mass fraction in %	percent	0-100	1-7
Texture - silt	SLTPPT	Silt content (2-50 micrometre) mass fraction in %	percent	0-100	1-7
Water available	AWCh1	Available soil water capacity (volumetric fraction) for h1	percent	0-100	1-7
Water available	AWCh2	Available soil water capacity (volumetric fraction) for h2	percent	0-100	1-7
Water available	AWCh3	Available soil water capacity (volumetric fraction) for h3	percent	0-100	1-7
Water available	AWCtS	Saturated water content (volumetric fraction) for tS	percent	0-100	1-7
Water available	WWP	Available soil water capacity (volumetric fraction) until wilting point	percent	0-100	1-7

Maps of SoilGrids layers (resampled to 30m pixels) – Burkina Faso project area

Levels 1-4 consolidated to provide an estimate for top 30cm

BLDFIE – Bulk density (fine earth) in kg / cubic-meter

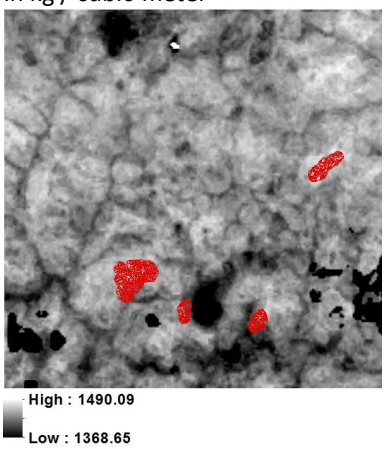

OCDENS – Soil organic carbon density in kg per cubic-m

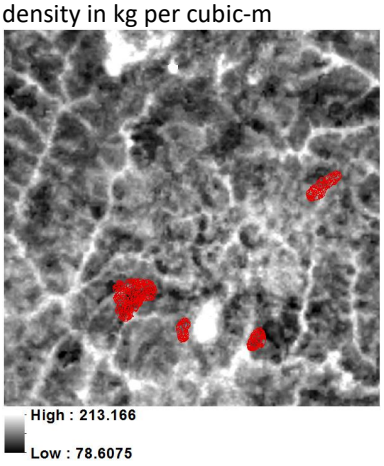

CECSOL – Cation exchange capacity of soil in cmolc/kg

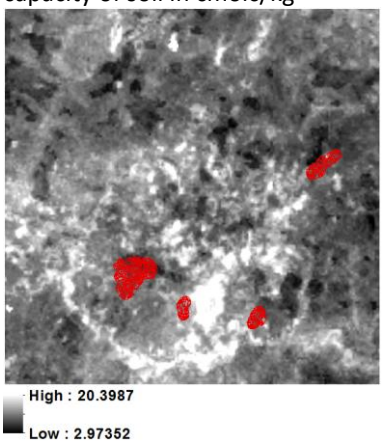

PHIHOX – Soil pH x 10 in water

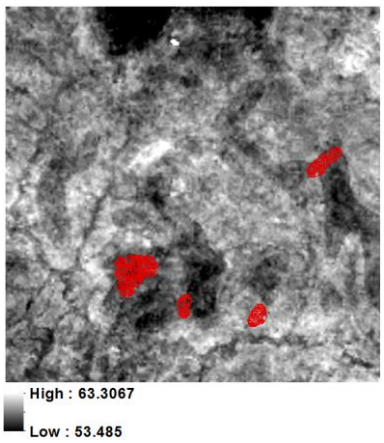

CLYPPT – Clay content (0-2 µm) mass fraction (%)

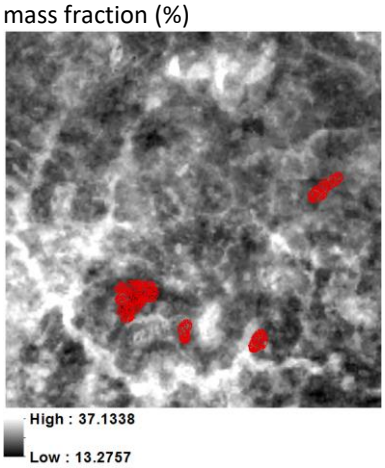

CRFVOL – Coarse fragments volumetric (%)

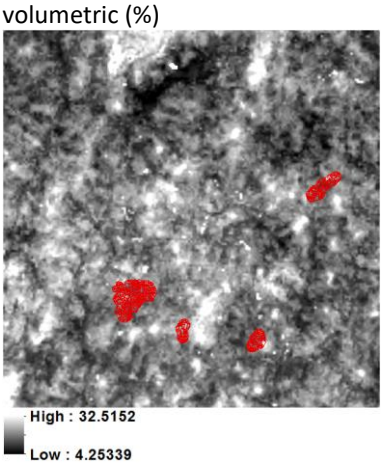

SNDPPT – Sand content (50-2000 µm) mass fraction (%)

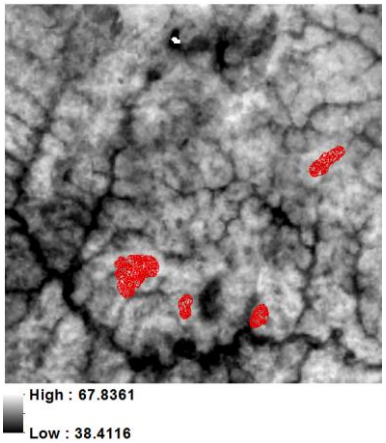

SLTPPT – Silt content (2-50 µm) mass fraction (%)

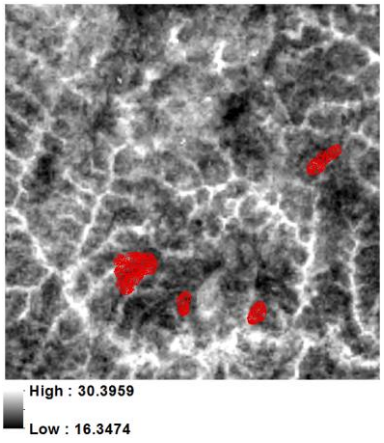

WWP – Available soil water capacity (%) until wilting point

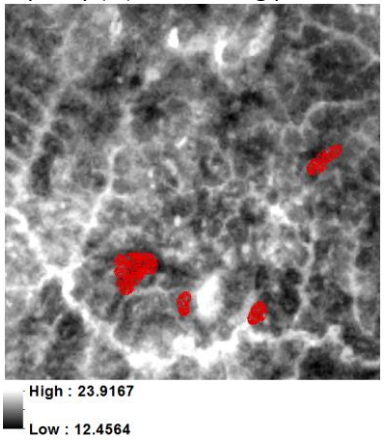

Maps of SoilGrids layers (resampled to 30m pixels) – Adelaide project area

Levels 1-4 consolidated to provide an estimate for top 30cm

BLDFIE – Bulk density (fine earth) in kg / cubic-meter

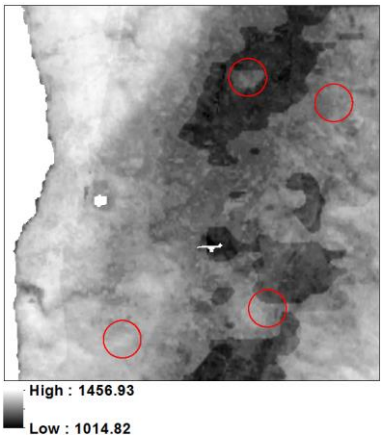

OCDENS – Soil organic carbon density in kg per cubic-m

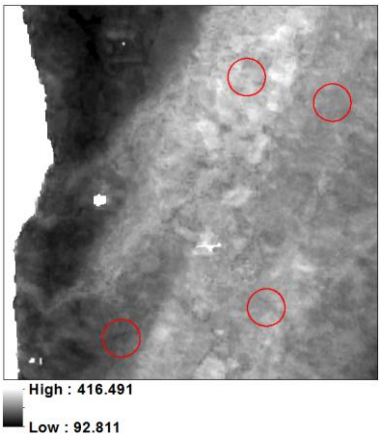

CECSOL – Cation exchange capacity of soil in cmolc/kg

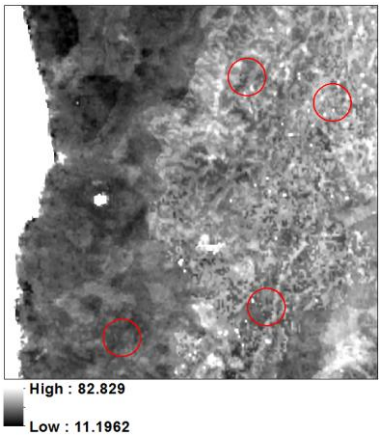

PHIHOX – Soil pH x 10 in water

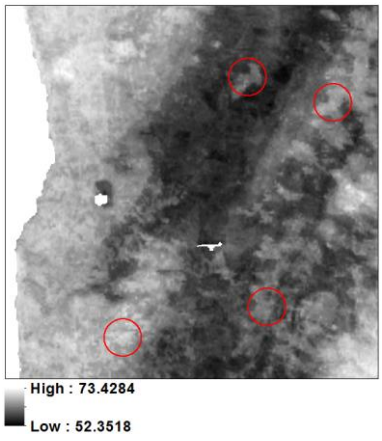

CLYPPT – Clay content (0-2 µm) mass fraction (%)

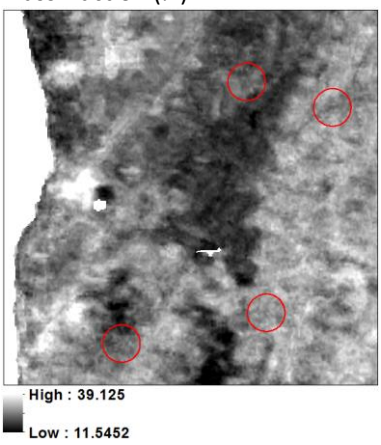

CRFVOL – Coarse fragments volumetric (%)

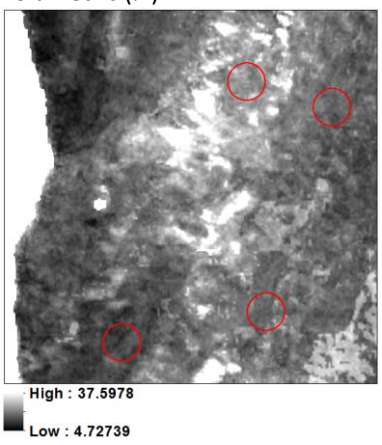

SNDPPT – Sand content (50-2000 µm) mass fraction (%)

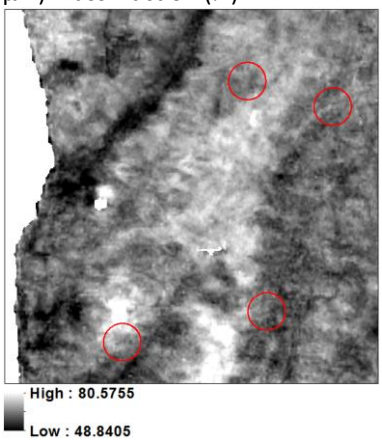

SLTPPT – Silt content (2-50 µm) mass fraction (%)

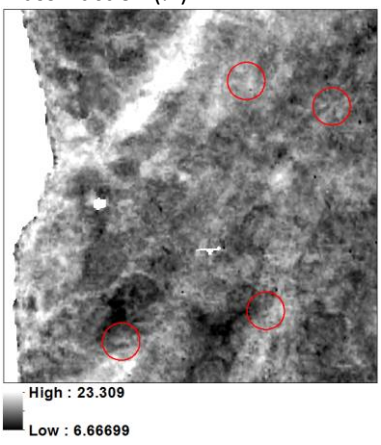

WWP – Available soil water capacity (%) until wilting point

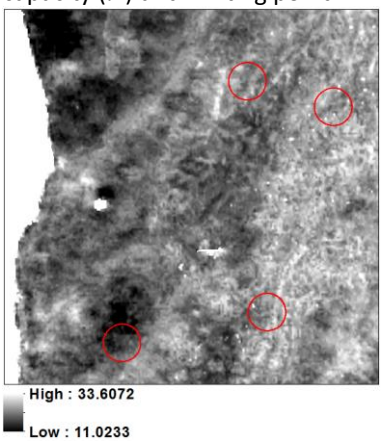

Maps of SoilGrids layers (resampled to 30m pixels) – Marrabel project area

Levels 1-4 consolidated to provide an estimate for top 30cm

BLDFIE – Bulk density (fine earth) in kg / cubic-meter

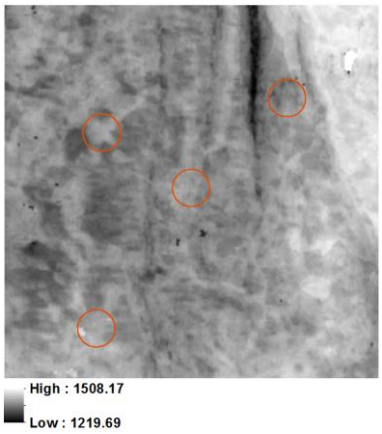

OCDENS – Soil organic carbon density in kg per cubic-m

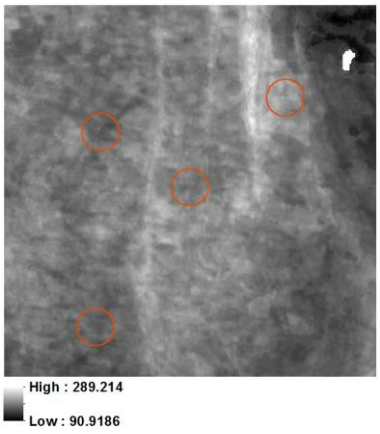

CECSOL – Cation exchange capacity of soil in cmolc/kg

PHIHOX – Soil pH x 10 in water

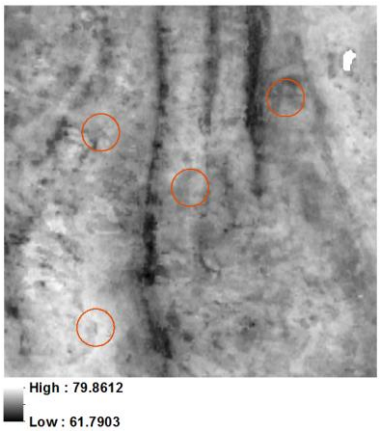

CLYPPT – Clay content (0-2 µm) mass fraction (%)

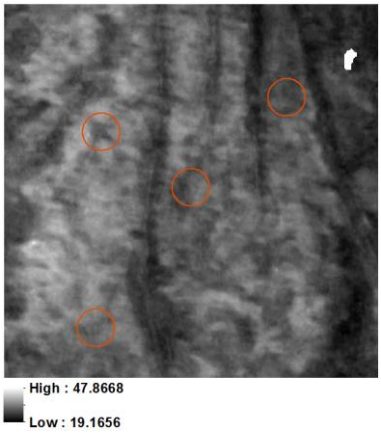

CRFVOL – Coarse fragments volumetric (%)

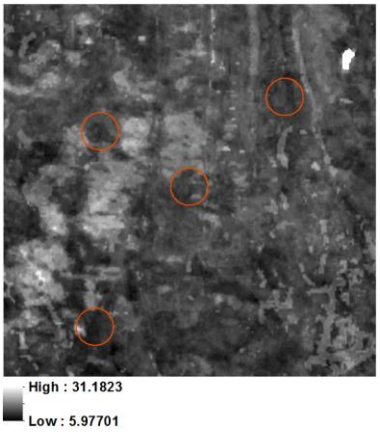

SNDPPT – Sand content (50-2000 µm) mass fraction (%)

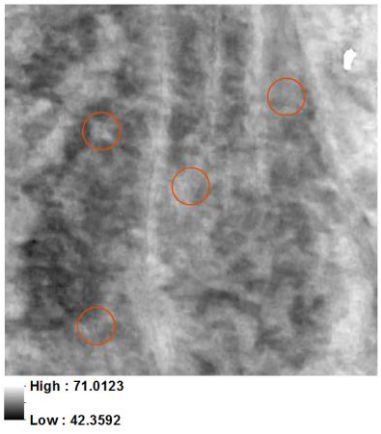

SLTPPT – Silt content (2-50 µm) mass fraction (%)

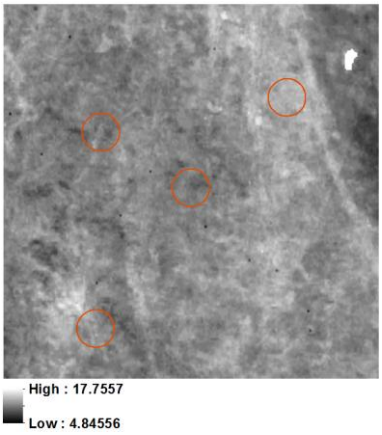

WWP – Available soil water capacity (%) until wilting point

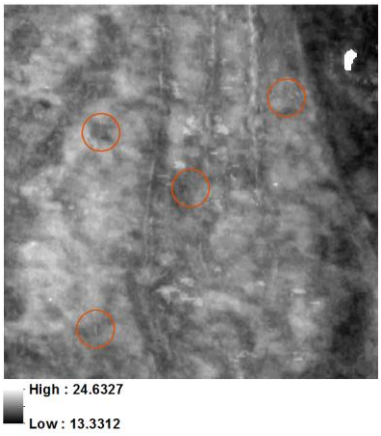

Appendix E – Correlations between environmental layers

Burkina Faso project area

Table E-1 Pearson correlation coefficients for pairwise comparisons of environmental layers for the Burkina Faso project region (absolute values > 0.9 shaded green and > 0.5 shaded yellow)

	Layer	0	1	2	3	A	B	C	D	E	F	G	H	I	M	N	W	X	Y	Z
0	DEM	1	0.48	-0.81	0.02	-0.08	0.02	0.02	-0.02	-0.02	0.1	0.04	-0.03	-0.01	-0.45	0.02	-0.22	0	-0.08	-0.04
1	Slope	0.48	1	-0.53	-0.15	-0.19	0.12	0.04	-0.01	0.01	0.08	0	-0.01	0.01	-0.33	-0.29	-0.09	0	-0.04	-0.02
2	Wetness index	-0.81	-0.53	1	0	0.05	0	-0.01	0.01	0.03	-0.1	-0.06	0.04	0.01	0.55	0.1	0.24	0	0.09	0.04
3	Solar radiation	0.02	-0.15	0	1	0.02	-0.02	0	0	0	0	0.01	0	0	0.01	0.04	0.02	0	0.01	0
A	BLDFIE	-0.08	-0.19	0.05	0.02	1	0.45	0.71	0.77	0.71	0.72	0.8	0.74	0.73	-0.03	0.05	0.49	0.76	0.72	0.77
B	OCDENS	0.02	0.12	0	-0.02	0.45	1	0.85	0.81	0.85	0.81	0.72	0.84	0.84	0.06	-0.03	0.49	0.82	0.76	0.8
C	CECSOL	0.02	0.04	-0.01	0	0.71	0.85	1	0.99	0.98	0.97	0.94	0.99	0.99	0	-0.01	0.61	0.99	0.92	0.98
D	PHIHOX	-0.02	-0.01	0.01	0	0.77	0.81	0.99	1	0.98	0.97	0.96	0.99	0.99	0	0	0.63	1	0.94	0.99
E	CLYPPT	-0.02	0.01	0.03	0	0.71	0.85	0.98	0.98	1	0.96	0.92	0.98	0.99	0.03	0	0.62	0.98	0.92	0.97
F	CRFVOL	0.1	0.08	-0.1	0	0.72	0.81	0.97	0.97	0.96	1	0.95	0.96	0.97	-0.07	-0.02	0.58	0.98	0.91	0.96
G	SNDPPT	0.04	0	-0.06	0.01	0.8	0.72	0.94	0.96	0.92	0.95	1	0.94	0.94	-0.07	-0.01	0.59	0.96	0.9	0.95
H	SLTPPT	-0.03	-0.01	0.04	0	0.74	0.84	0.99	0.99	0.98	0.96	0.94	1	0.99	0.03	0	0.63	0.99	0.93	0.98
I	WWP	-0.01	0.01	0.01	0	0.73	0.84	0.99	0.99	0.99	0.97	0.94	0.99	1	0.02	0	0.62	0.99	0.93	0.98
M	MRVBF	-0.45	-0.33	0.55	0.01	-0.03	0.06	0	0	0.03	-0.07	-0.07	0.03	0.02	1	0.33	0.13	0	0.05	0.02
N	MRRTF	0.02	-0.29	0.1	0.04	0.05	-0.03	-0.01	0	0	-0.02	-0.01	0	0	0.33	1	0.02	0	0.02	0
W	Dose rate	-0.22	-0.09	0.24	0.02	0.49	0.49	0.61	0.63	0.62	0.58	0.59	0.63	0.62	0.13	0.02	1	0.61	0.83	0.7
X	K percent	0	0	0	0	0.76	0.82	0.99	1	0.98	0.98	0.96	0.99	0.99	0	0	0.61	1	0.93	0.99
Y	Th ppm	-0.08	-0.04	0.09	0.01	0.72	0.76	0.92	0.94	0.92	0.91	0.9	0.93	0.93	0.05	0.02	0.83	0.93	1	0.96
Z	U ppm	-0.04	-0.02	0.04	0	0.77	0.8	0.98	0.99	0.97	0.96	0.95	0.98	0.98	0.02	0	0.7	0.99	0.96	1

Marrabel test region

Table E-1 Pearson correlation coefficients for pairwise comparisons of environmental layers for the Marrabel test region (absolute values > 0.9 shaded green and > 0.5 shaded yellow)

	Layer	0	1	2	3	A	B	C	D	E	F	G	H	I	M	N	X	Y	Z
0	DEM	1	0.29	-0.41	0.27	0.02	0.06	0.07	-0.17	-0.20	-0.04	0.16	0.17	-0.12	-0.25	0.00	-0.06	-0.01	0.17
1	Slope	0.29	1	-0.67	-0.31	0.00	0.05	0.07	-0.19	-0.24	-0.01	0.19	0.17	-0.15	-0.58	-0.33	-0.02	-0.03	0.00
2	Wetness index	-0.41	-0.67	1	0.09	0.02	-0.06	-0.08	0.17	0.17	0.01	-0.14	-0.14	0.10	0.71	0.20	0.02	0.01	-0.09
3	Solar radiation	0.27	-0.31	0.09	1	0.01	0.00	0.00	0.00	-0.01	0.00	0.01	0.01	-0.01	0.10	0.14	0.00	0.00	0.06
A	BLDFIE	0.02	0.00	0.02	0.01	1	-0.39	-0.32	0.13	-0.40	-0.12	0.43	-0.14	-0.48	0.03	-0.03	0.11	-0.05	-0.07
B	OCDENS	0.06	0.05	-0.06	0.00	-0.39	1	0.32	-0.23	0.25	0.04	-0.30	0.20	0.36	-0.05	0.01	-0.08	0.05	0.08
C	CECSOL	0.07	0.07	-0.08	0.00	-0.32	0.32	1	-0.22	0.16	0.03	-0.21	0.19	0.27	-0.06	0.00	-0.07	0.03	0.05
D	PHIHOX	-0.17	-0.19	0.17	0.00	0.13	-0.23	-0.22	1	0.17	0.11	-0.10	-0.25	0.02	0.12	0.05	0.03	0.01	-0.02
E	CLYPPT	-0.20	-0.24	0.17	-0.01	-0.40	0.25	0.16	0.17	1	0.16	-0.62	-0.08	0.59	0.10	0.09	-0.07	0.10	0.10
F	CRFVOL	-0.04	-0.01	0.01	0.00	-0.12	0.04	0.03	0.11	0.16	1	-0.16	0.00	0.14	0.00	0.02	-0.05	-0.01	-0.01
G	SNDPPT	0.16	0.19	-0.14	0.01	0.43	-0.30	-0.21	-0.10	-0.62	-0.16	1	0.01	-0.60	-0.08	-0.08	0.08	-0.10	-0.10
H	SLTPPT	0.17	0.17	-0.14	0.01	-0.14	0.20	0.19	-0.25	-0.08	0.00	0.01	1	0.04	-0.09	-0.04	-0.04	-0.01	0.03
I	WWP	-0.12	-0.15	0.10	-0.01	-0.48	0.36	0.27	0.02	0.59	0.14	-0.60	0.04	1	0.05	0.07	-0.09	0.10	0.11
M	MRVBF	-0.25	-0.58	0.71	0.10	0.03	-0.05	-0.06	0.12	0.10	0.00	-0.08	-0.09	0.05	1	-0.04	-0.02	-0.03	-0.09
N	MRRTF	0.00	-0.33	0.20	0.14	-0.03	0.01	0.00	0.05	0.09	0.02	-0.08	-0.04	0.07	-0.04	1	-0.01	0.02	0.05
X	K percent	-0.06	-0.02	0.02	0.00	0.11	-0.08	-0.07	0.03	-0.07	-0.05	0.08	-0.04	-0.09	-0.02	-0.01	1	0.50	0.40
Y	Th ppm	-0.01	-0.03	0.01	0.00	-0.05	0.05	0.03	0.01	0.10	-0.01	-0.10	-0.01	0.10	-0.03	0.02	0.50	1	0.61
Z	U ppm	0.17	0.00	-0.09	0.06	-0.07	0.08	0.05	-0.02	0.10	-0.01	-0.10	0.03	0.11	-0.09	0.05	0.40	0.61	1

Adelaide test region

Table E-1 Pearson correlation coefficients for pairwise comparisons of environmental layers for the Burkina Faso project region (absolute values > 0.9 shaded green and > 0.5 shaded yellow)

	Layer	0	1	2	3	A	B	C	D	E	F	G	H	I	M	N
0	DEM	1	0.32	-0.61	0.12	-0.76	0.79	0.23	-0.27	0.02	0.23	0.13	0.01	0.11	-0.30	-0.45
1	Slope	0.32	1	-0.02	0.14	-0.38	0.41	0.09	-0.17	-0.02	0.12	0.02	0.00	0.03	-0.33	-0.53
2	Wetness index	-0.61	-0.02	1	0.00	0.52	-0.56	-0.13	0.22	0.01	-0.15	-0.04	0.02	-0.05	0.00	0.00
3	Solar radiation	0.12	0.14	0.00	1	-0.04	0.04	0.04	0.01	0.03	0.03	0.04	0.02	0.03	0.19	-0.43
A	BLDFIE	-0.76	-0.38	0.52	-0.04	1	-0.82	-0.08	0.45	0.15	-0.09	0.03	0.16	0.05	0.33	0.43
B	OCDENS	0.79	0.41	-0.56	0.04	-0.82	1	0.38	-0.18	0.15	0.39	0.27	0.14	0.25	-0.35	-0.45
C	CECSOL	0.23	0.09	-0.13	0.04	-0.08	0.38	1	0.80	0.94	0.97	0.96	0.95	0.97	-0.08	-0.11
D	PHIHOX	-0.27	-0.17	0.22	0.01	0.45	-0.18	0.80	1	0.91	0.79	0.86	0.92	0.88	0.14	0.18
E	CLYPPT	0.02	-0.02	0.01	0.03	0.15	0.15	0.94	0.91	1	0.93	0.96	0.98	0.98	0.01	0.01
F	CRFVOL	0.23	0.12	-0.15	0.03	-0.09	0.39	0.97	0.79	0.93	1	0.95	0.94	0.97	-0.08	-0.12
G	SNDPPT	0.13	0.02	-0.04	0.04	0.03	0.27	0.96	0.86	0.96	0.95	1	0.96	0.98	-0.02	-0.03
H	SLTPPT	0.01	0.00	0.02	0.02	0.16	0.14	0.95	0.92	0.98	0.94	0.96	1	0.99	0.01	0.01
I	WWP	0.11	0.03	-0.05	0.03	0.05	0.25	0.97	0.88	0.98	0.97	0.98	0.99	1	-0.03	-0.04
M	MRVBF	-0.30	-0.33	0.00	0.19	0.33	-0.35	-0.08	0.14	0.01	-0.08	-0.02	0.01	-0.03	1	0.13
N	MRRTF	-0.45	-0.53	0.00	-0.43	0.43	-0.45	-0.11	0.18	0.01	-0.12	-0.03	0.01	-0.04	0.13	1

Appendix F – Maxent runs for the Burkina Faso project area

Tests of predictors and parameters

	Environmental layers	Feature classes	Beta multiplier	Hybrid soils	Test description	Mean test AUC: 12 crops	5 main crops	4 dryland crops	Rice
1	0123	LQPH	1		Terrain only	0.63	0.64	0.63	0.66
2	0123ABCDEFGHIMNWXYZ	LQPH	1	Means + jitter	All layers	0.63	0.65	0.63	0.71
3	0123ABCDEFGHIMNWXYZ	LQPH	1	Means + jitter	Linear trend removed from DEM	0.66	0.68	0.66	0.77
4	0123ABCDEFGHIMNWXYZ	LQPH	1	Means + jitter	Quadratic trend removed from DEM	0.62	0.63	0.61	0.70
5	12	LQPH	1		Terrain only	0.61	0.63	0.60	0.75
6	123	LQPH	1		Terrain only	0.63	0.65	0.62	0.75
7	123	LQPH	1		Terrain – prototype layers	0.63	0.65	0.62	0.76
8	123ABCDEFGHI	LQPH	1	Means + jitter	Terrain + soils	0.62	0.67	0.62	0.89
9	123ABCDEFGHI	L	1	Means + jitter	Terrain + soils	0.60	0.64	0.59	0.86
10	123ABCDEFGHI	L	1	Means	Terrain + soils	0.69	0.71	0.67	0.89
11	123ABCDEFGHI	LQ	1	Means + jitter	Terrain + soils	0.61	0.67	0.61	0.88
12	123ABCDEFGHI	LQP	1	Means + jitter	Terrain + soils	0.62	0.69	0.64	0.89
13	123ABCDEFGHI	LQP	1	Means	Terrain + soils	0.69	0.73	0.68	0.91
14	123ABCDEFGHI	LQP	1	Means + SG (50:50)	Terrain + soils	0.72	0.76	0.74	0.81
15	123ABCDEFGHI	LQP	1	Means + SG (75:25)	Terrain + soils	0.70	0.75	0.72	0.86
16	123ABCDEFGHIMNX	LQH	1	Means + jitter	Mixed layers	0.61	0.65	0.61	0.81
17	123ABCDEFGHIMNX	Q	1	Means + jitter	Mixed layers	0.60	0.65	0.60	0.88
18	123ABCDEFGHIMNX	Q	1	Means + jitter	Rerun of test	0.60	0.65	0.60	0.88
19	123ABEFGHIX	LQPH	1	Means + jitter	Mixed layers	0.64	0.69	0.64	0.91
20	123ABMNWX	LQPH	1	Means + jitter	Uncorrelated layers	0.61	0.65	0.61	0.83
21	123ABMWX	LQPH	1	Means + jitter	Uncorrelated layers	0.62	0.68	0.62	0.91
22	123ABWX	LQPH	1	Means + jitter	Uncorrelated layers	0.64	0.68	0.63	0.90
23	123ABWX	LQPH	2	Means + jitter	Test beta-multiplier	0.62	0.68	0.63	0.90
24	123ABWX	LQPH	4	Means + jitter	Test beta-multiplier	0.60	0.67	0.61	0.90
25	123M	LQPH	1		Terrain only	0.61	0.65	0.60	0.83
26	123MN	LQPH	1		Terrain only	0.60	0.62	0.59	0.73
27	123MN	L	1		Terrain only	0.58	0.59	0.54	0.77
28	123MN	LQP	1		Terrain only	0.60	0.63	0.59	0.80
29	123MNWXYZ	LQP	1		Terrain + radiometrics	0.61	0.65	0.60	0.85

	Environmental layers	Feature classes	Beta multiplier	Hybrid soils	Test description	Mean test AUC: 12 crops	5 main crops	4 dryland crops	Rice
30	123PQ	L	1	Means + jitter	Terrain + PCA soils	0.57	0.59	0.53	0.86
31	123PQ	LQP	1	Means + jitter	Terrain + PCA soils	0.61	0.65	0.59	0.89
32	123PQ	LQP	1	Means + SG (50:50)	Terrain + PCA soils	0.79	0.80	0.80	0.77
33	123PQ	LQP	1	Means + SG (75:25)	Terrain + PCA soils	0.80	0.84	0.83	0.87
34	123PQRS	LQP	1	Means + SG (75:25)	Terrain + PCA soils	0.76	0.81	0.79	0.86
35	123WXYZ	LQPH	1		Terrain + radiometrics	0.65	0.68	0.64	0.85
36	12ABCDEFGHI	LQPH	1	Means + jitter	Terrain + soils	0.61	0.66	0.60	0.89
37	12ABCDEFGHI	LQPH	1	Means	Terrain + soils	0.68	0.70	0.65	0.91
38	12ABCDEFGHIMNWXYZ	LQPH	1	Means + jitter	Mixed layers	0.58	0.62	0.57	0.80
39	13ABCDEGHIMX	LQPH	1	Means + jitter	Mixed layers	0.63	0.70	0.65	0.90
40	ABCDEFGHI	LQPH	1	Means + jitter	Soils only	0.59	0.65	0.60	0.85
41	PQ	LQPH	1	Means + jitter	PCA of 17 layers	0.64	0.66	0.61	0.87
42	PQR_SPCA2	LQPH	1	Means + jitter	PCA of 123ABMWX	0.60	0.64	0.58	0.89
43	PQRSTUV	LQPH	1	Means + jitter	PCA of all 19 layers	0.66	0.67	0.62	0.87
44	PQRSTUV	LQPH	1	Means + jitter	PCA of 17 layers	0.65	0.68	0.63	0.89
45	PQRSTUV	LQPH	1	Means + jitter	PCA of 123ABMWX	0.62	0.66	0.60	0.89
46	PQRSTUV	LQPH	1	Means + jitter	PCA of 123ABWX	0.62	0.68	0.63	0.89
47	WXYZ	LQPH	1		Radiometrics only	0.60	0.62	0.57	0.80
	Mean					**0.64**	**0.67**	**0.63**	**0.84**

Replication runs

	Replication method	Feature classes	Beta multiplier	Hybrid soils	Environmental layers	Mean test AUC: 12 crops	5 main crops	4 dryland crops	Rice
	5-fold cross validation	LQPH	1	Means + jitter	123ABCDEFGHIMNWXYZ	0.75	0.80	0.76	0.95
	25-fold cross validation	LQPH	1	Means + jitter	123ABCDEFGHIMNWXYZ	0.76	0.80	0.76	0.94
	Subsample – 25 iterations	LQPH	1	Means + jitter	123ABCDEFGHIMNWXYZ	0.76	0.80	0.77	0.95
	Bootstrap – 25 iterations	LQPH	1	Means + jitter	123ABCDEFGHIMNWXYZ	0.81	0.82	0.79	0.95
	Mean					**0.77**	**0.80**	**0.77**	**0.95**

Appendix G – Land use potential assessment criteria

Land use potential **class definitions**

Class	Potential	Definition
Class 1	High	Land with high productive potential and requiring no more than standard management practices to sustain productivity.
Class 2	Moderately high	Land with moderately high productive potential and / or requiring specific, but widely accepted and used, management practices to sustain productivity.
Class 3	Moderate	Land with moderate productive potential and / or requiring specialized management practices to sustain productivity.
Class 4	Moderately low	Land with marginal productive potential and / or requiring very highly specialized management skills to sustain productivity.
Class 5	Low	Land with low productive potential and /or permanent limitations which effectively preclude its use.
Class X	Not applicable *	Urban, evaporation pans, quarry, water, rock, saline soil, reservoir, cliff, reef etc.

Rules used to assign *Land use potential* mapping categories to *Soil landscape map units*

Mapping category	Rules to assign *Land use potential* mapping categories (*Mapping data*)	Order of assignment
Aa	If *Class 1* > 60%	2
Ab	If sum (*Classes 1+2*) > 60% *(and mapping category not already assigned)*	3
Ac	If sum (*Classes 1+2+3*) > 60% *(and mapping category not already assigned)*	5
Ad	*Class 3* > 60% and (*Classes 1+2*) < 1% *(and mapping category not already assigned)*	4
B	If sum (*Classes 1+2+3*) > 30-60% *(and mapping category not already assigned)*	6
C	If sum (*Classes 1+2+3*) > 10-30% *(and mapping category not already assigned)*	7
D	If sum (*Classes 1+2+3*) > 1-10% *(and mapping category not already assigned)*	8
Ea	If *Class 4* > 50% *(and mapping category not already assigned)*	9
Eb	Any remaining map unit with mapping category not already assigned	10
X	If *Class X* > 70%	1

Land use potential **mapping categories**

Mapping category	Proportion of land with moderate to high potential	Most common potential class
Aa	More than 60%	High potential (**mostly *Class 1***)
Ab	More than 60%	Moderately high potential (**mostly *Class 2***)
Ac	More than 60%	Moderate to high (mixed)
Ad	More than 60%	Moderate potential (**mostly *Class 3***)
B	30-60%	Low to high potential (mixed)
C	10-30%	Moderately low to low potential (mixed)
D	1-10%	Moderately low to low potential (mixed)
Ea	Less than 1%	Moderately low potential (**mostly *Class 4***)
Eb	Less than 1%	Low potential (**mostly *Class 5***)
X	–	–

Appendix H – Multi-page set of printable maps for crop suitability predictions

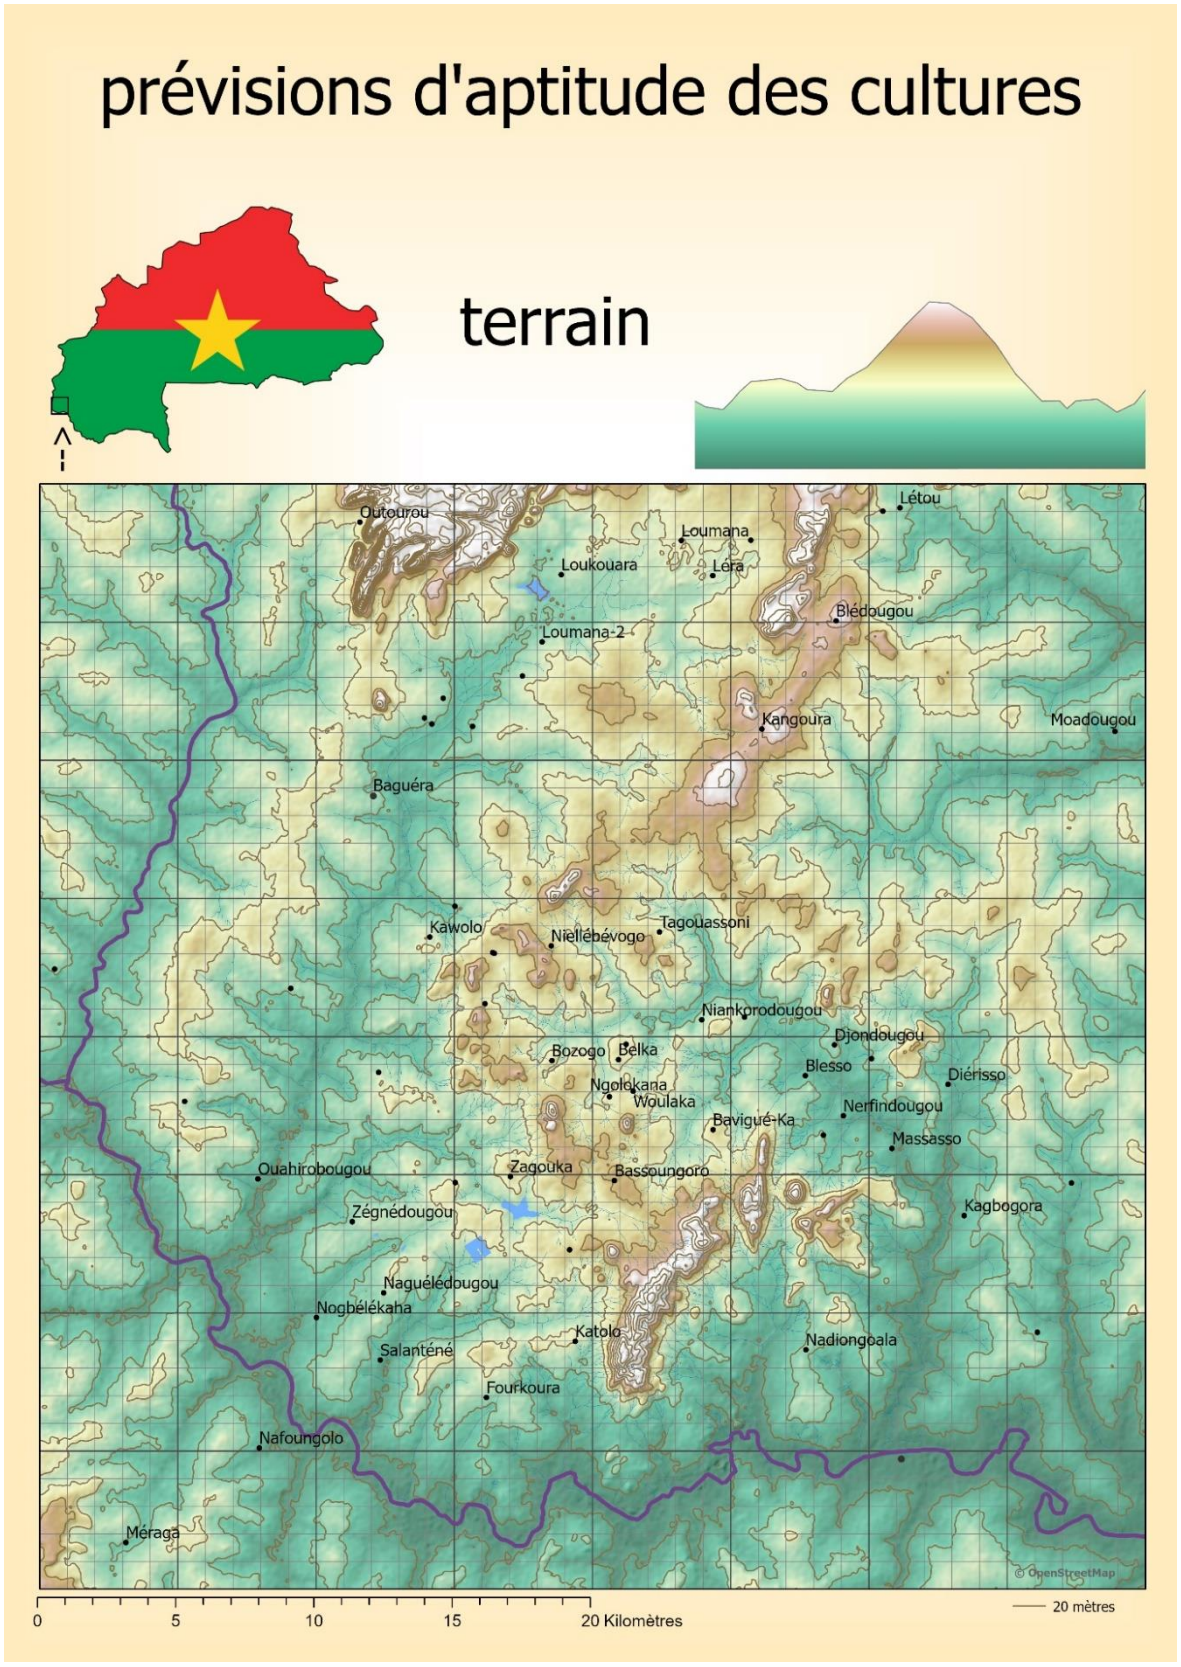

prévisions d'aptitude des cultures

mɔnsɔŋɔ

maïs

maize

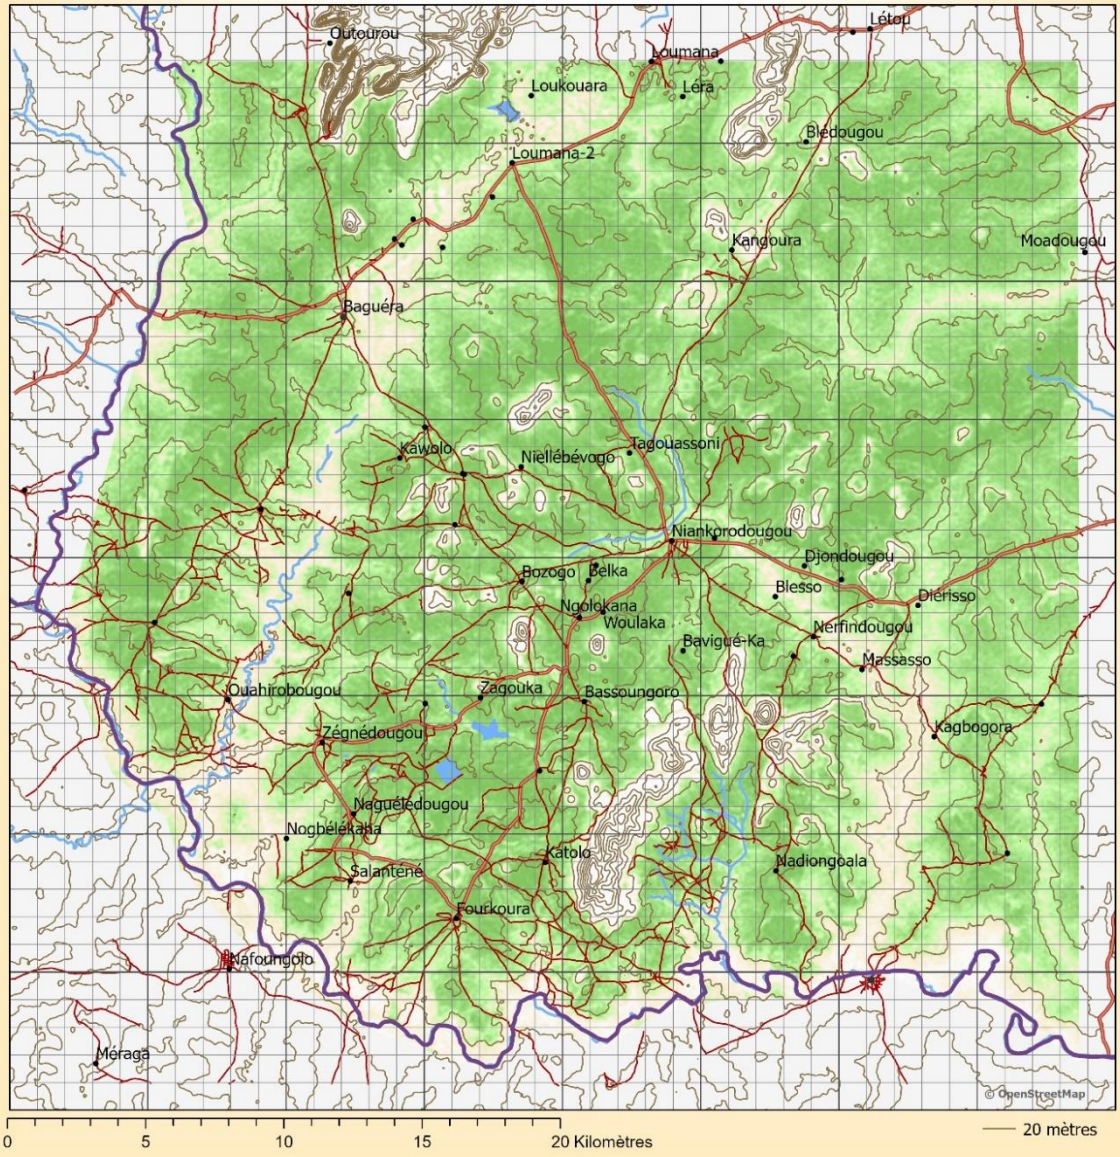

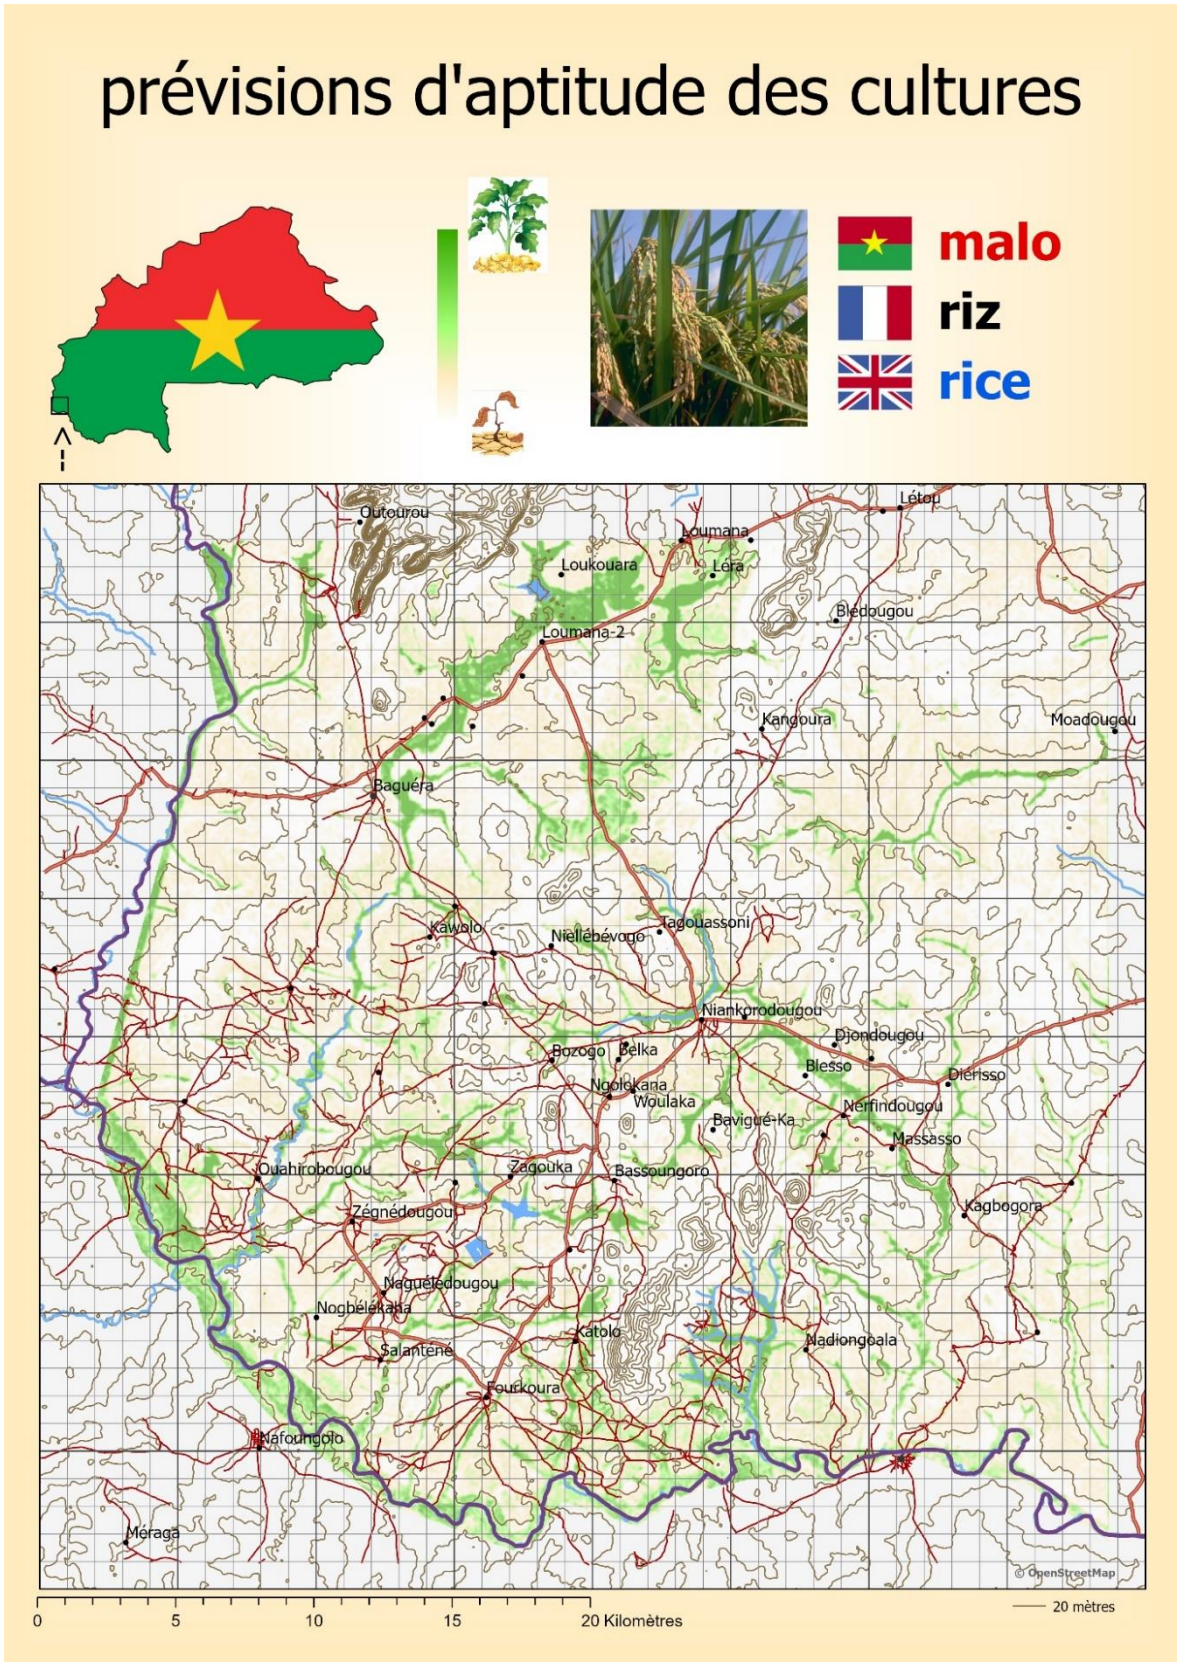
prévisions d'aptitude des cultures
malo
riz
rice
Outourou
Loumana
Loukouara
Léra
Léton
Bledougou
Loumana-2
Kangoura
Moadougou
Baguéra
Kawolo
Niellébévogo
Tagouassoni
Niankorodougou
Djondougou
Bozogo
Belka
Blesso
Dierisso
Ngolokaria
Woulaka
Nerfindougou
Bavigué-Ka
Massasso
Ouahirobougou
Zagouka
Bassoungoro
Kagbogora
Zégnédougou
Naguélédougou
Noghélékaha
Katolo
Salanténé
Nadiongoala
Fourkoura
Nafoungolo
Méraga
© OpenStreetMap
0
5
10
15
20 Kilomètres
20 mètres

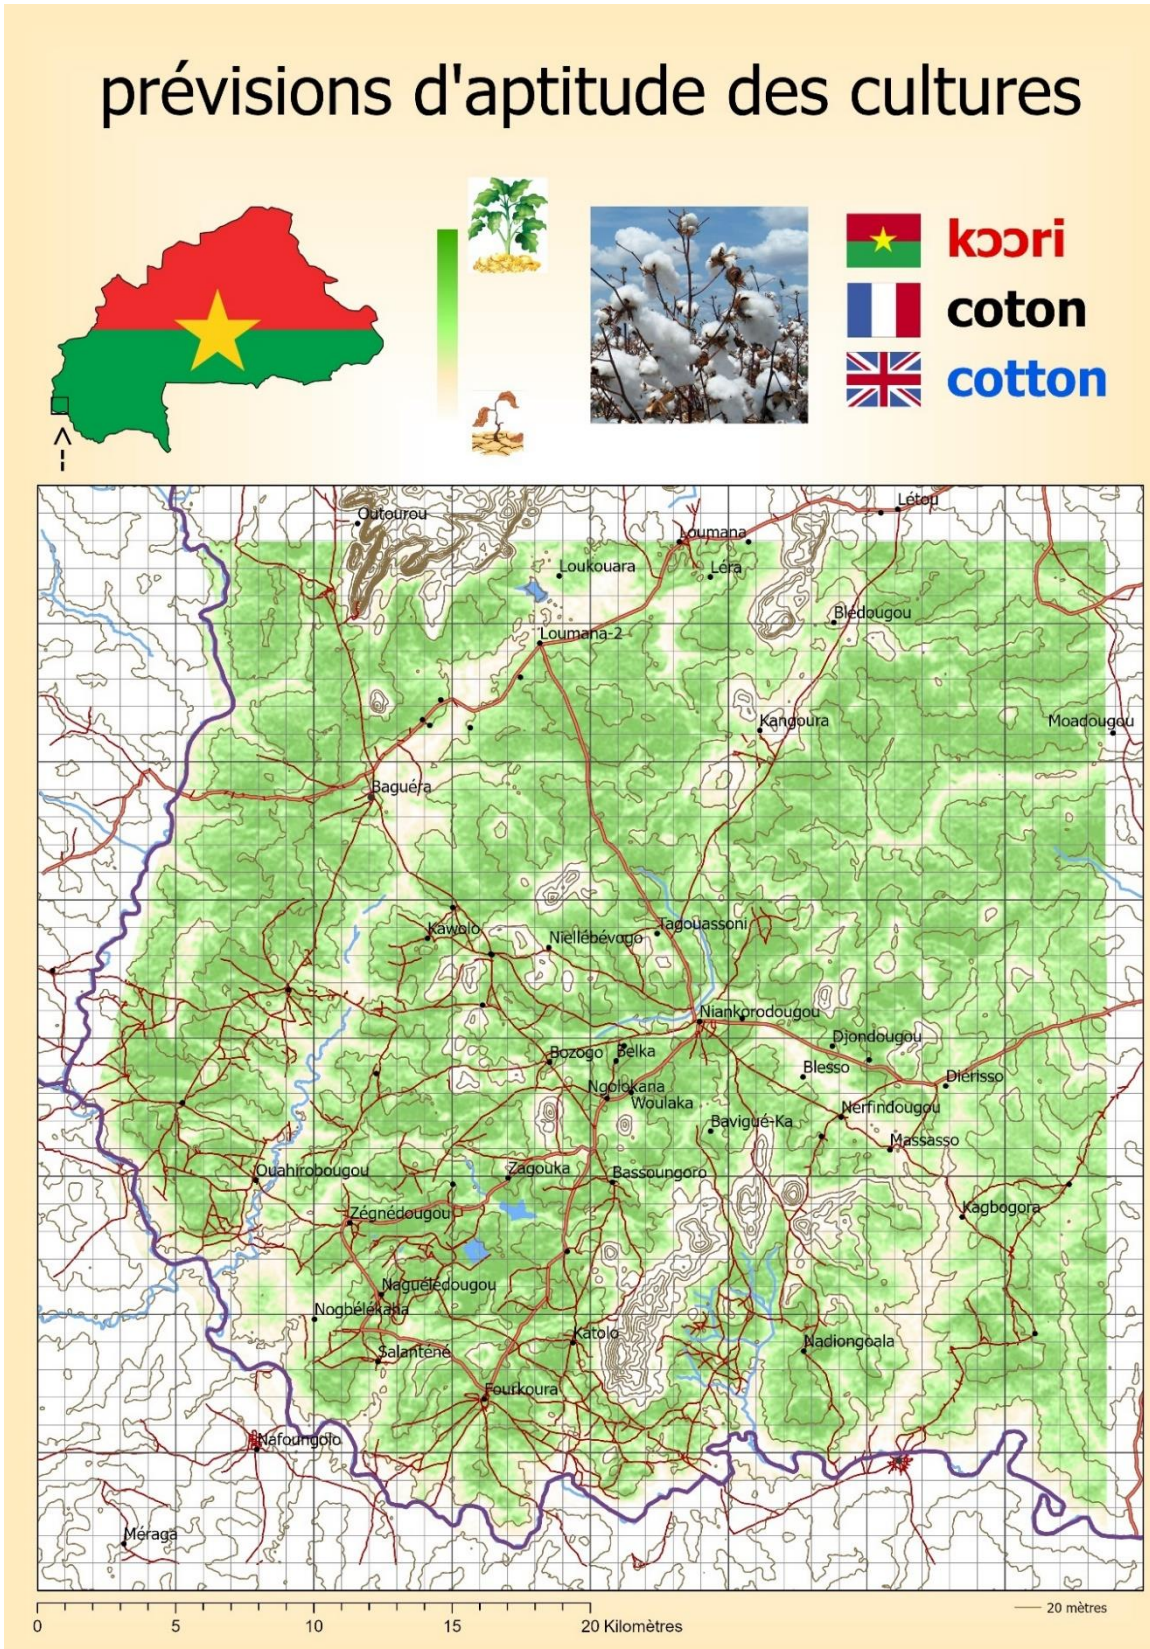
prévisions d'aptitude des cultures
kɔɔri
coton
cotton
Outourou
Loumana
Léto
Loukouara
Léra
Bledougou
Loumana-2
Kangoura
Moadougou
Baguéra
Kawolo
Niellébévogo
Tagouassoni
Niankorodougou
Djondougou
Bozogo
Belka
Blesso
Dierisso
Ngolokana
Woulaka
Nerfindougou
Bavigué-Ka
Massasso
Ouahirobougou
Zagouka
Bassoungoro
Zégnédougou
Kagbogora
Naguédédougou
Noghélékaha
Katolo
Nadiongoala
Salantené
Fourkoura
Nafoungolo
Méraga
0
5
10
15
20 Kilomètres
20 mètres

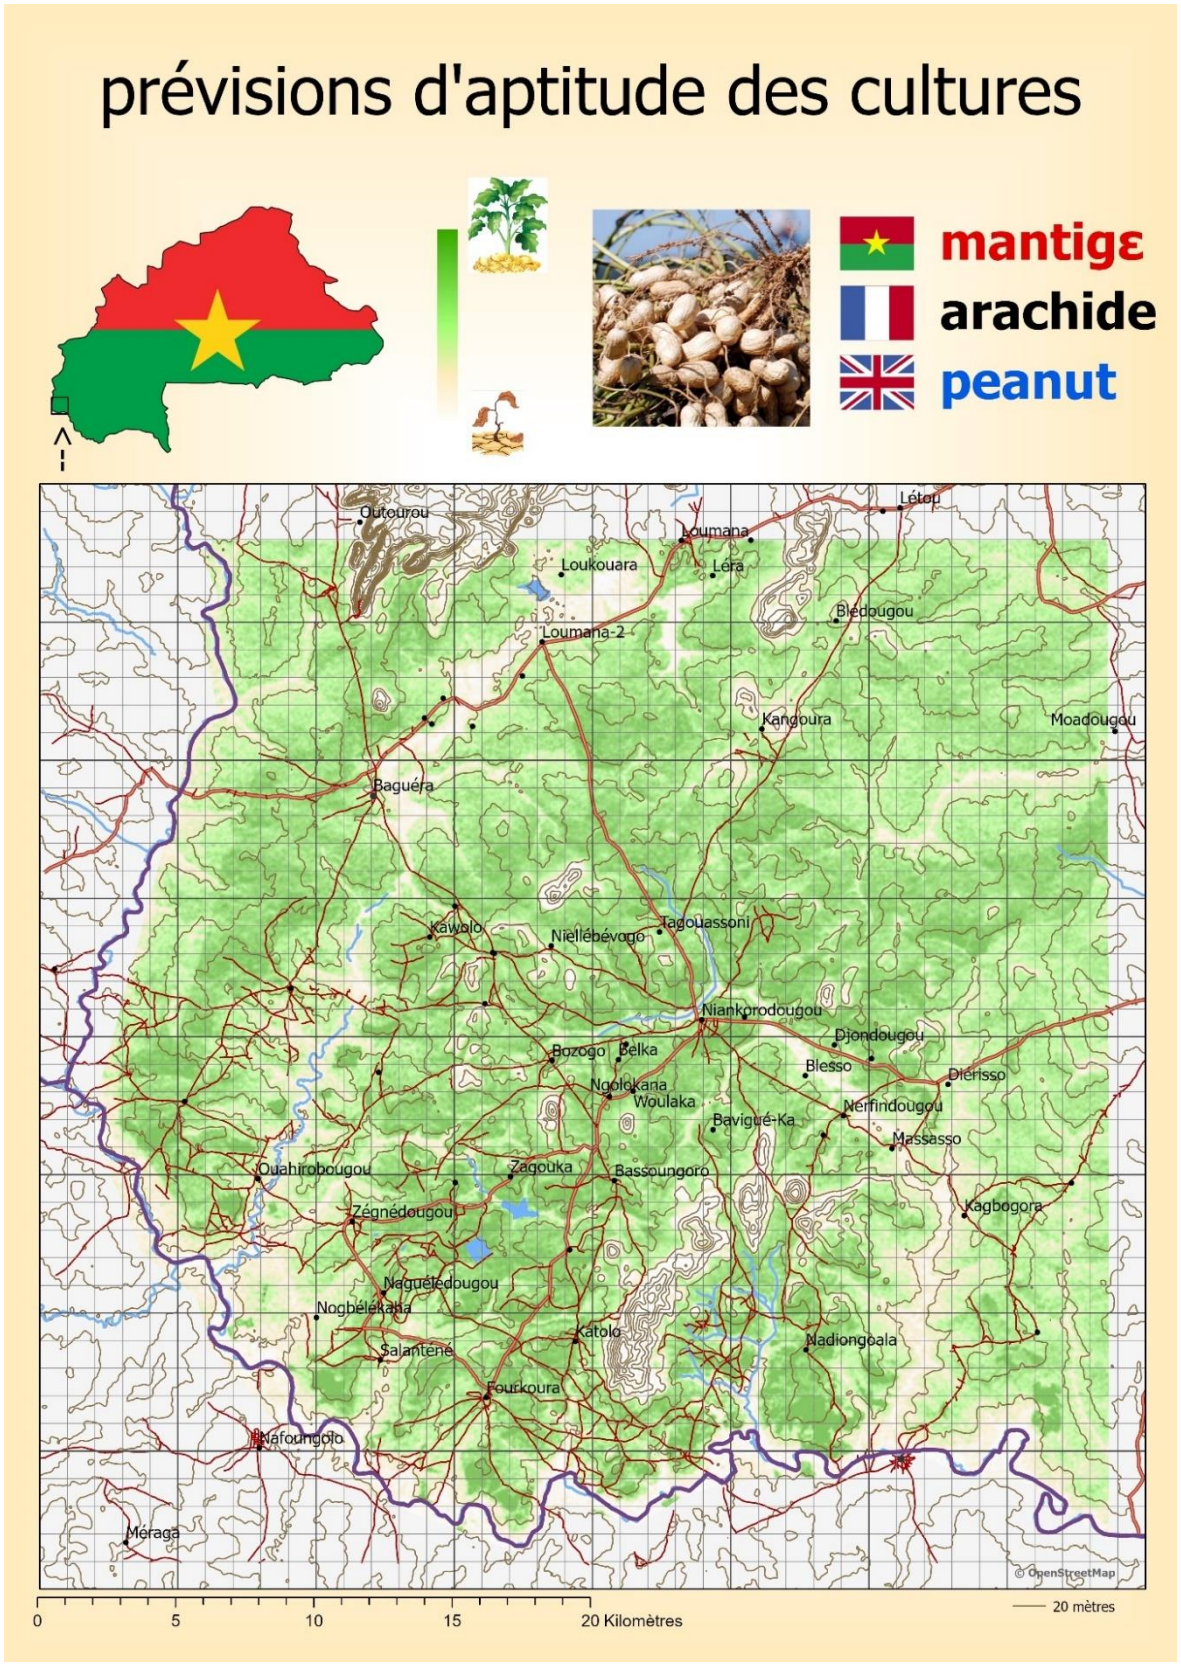
prévisions d'aptitude des cultures
mantigɛ
arachide
peanut
Outourou
Loumana
Léto
Loukouara
Léra
Bledougou
Loumana-2
Kangoura
Moadougou
Baguéra
Kawolo
Niellébévogo
Tagouassoni
Niankorodougou
Djondougou
Bozogo
Belka
Blesso
Dierisso
Ngolokana
Woulaka
Nerfindougou
Bavigué-Ka
Massasso
Ouahirobougou
Zagouka
Bassoungoro
Zégnédougou
Kagbogora
Naguélédougou
Nogbélékaha
Katolo
Nadiongoala
Salantené
Fourkoura
Nafoungolo
Méraga
© OpenStreetMap
0
5
10
15
20 Kilomètres
20 mètres

prévisions d'aptitude des cultures

bɛnɛ
sésame
sesame