Crop Map of England

Product Specification- v.2016.2

Non-Sensitive Information

Contents

Abbreviations

CROME – Crop Map of England

ESA – European Space Agency

GRD – Ground Range Detected

LUCODE – Land Use Code

NDVI – Normalised Difference Vegetation Index

SNAP – Sentinel Application Platform

VH – Vertical Horizontal polarisation

VV – Vertical Vertical polarisation

VV/VH – Vertical Vertical/ Vertical Horizontal dual polarisation

1 Introduction

1.1 Document Control

1.1.1 Revision History

Date	Author	Version	Change reference
28/11/2016	Sanjay Rana	0.1	Initial Version
09/12/16	Andrew McClune	0.2	Updated Version
29/12/2016	Sanjay Rana	0.3	Updated to include information from Confusion Matrix and general revisions
05/01/17	Andrew McClune	0.4	Updated Version
15/05/17	Andrew McClune	0.5	Updated Confusion Matrix
12/07/2017	Sanjay Rana	0.6	Minor revisions around document sensitive status and other updates.
22/08/2017	Sanjay Rana	0.7	Minor Revision to clarify zone names

1.1.2 Reviewers and Contributors

Name	Role	Team
Sanjay Rana	Lead Author	GI Tech Team
Andrew McClune	Co-Author	GI Tech Team
Liann Rafter	Contributor	GI Tech Team
Matauka Yeta	Reviewer	Land Technical Data Steward

1.2 Purpose Scope

This document defines the product specification for the Crop Map of England (CROME) layer delivered by the Rural Payments Agency (RPA).

CROME is a tessellated multi-temporal visualisation of the type and distribution of land covers identified using remote sensing techniques. The purpose of this specification is to define that representation plus the technical details for the reference dataset structure, data format and delivery. This document includes the technical specification and a general overview of the methodology used to produce the CROME layer.

The CROME layer covers most of the land in England (including small isles) into approximately 32 million hexagon cells. The CROME Layer is expected to be revised twice a year (April/May and July/August) to provide a snapshot of cropping variations on the land during the winter and summer seasons.

It is important to note here that this release of CROME is basically a scientific product and not a conventional cartographic product. Therefore, minimal efforts have been taken to generalise the thematic variation of crop and land use types to improve the cosmetic appearance of the layer.

1.3 Document Ownership and Distribution

This document is owned by the GI Technical Team within the RPA Design Directorate. The document is for use by RPA staff and third party suppliers. There are no restrictions on internal Defra Organisation employee access to this document, or to contractors/consultants, third parties and any other agency or body with access to Defra Organisation assets or data handling facilities.

1.4 Assumptions, Dependencies and Constraints

The document describes the derivation of a technical geospatial data product. It is not written for the layman and thus assumes that users of this document are knowledgeable in:

- Geographic Information Systems
- Earth Observation concepts

Product is released under the Open Government Licence. This requires a reference as follows:

2 Data Creation

2.1 Feature Classification

The sources for crop classification are satellite images from the Sentinel constellation. A combination of radar and multispectral imagery from Sentinel sensors was used for automatic classification. Sentinel-1 transmits and receives microwave radiation in the horizontal (H) and vertical (V) polarisation. Sentinel-2 captures multispectral imagery in 13 spectral bands at varying spatial resolutions. More information regarding the specification of each satellite can be found on the ESA website.

2.2 Geometry

Unlike conventional land cover classification maps, the classification of crops in CROME is represented as hexagon cells. Traditional thematic visualisation of geospatial information (e.g. distribution of crop types and other land use types) partitions the land into either a regular grid (e.g. CropScape-Cropland Data layer) or an agricultural/administrative polygonal representation of the land use extent (e.g. CEH Land Cover plus –Crops 2015). However, the CROME layer is based on the hexagonal grid which provides a superior opportunity to model and visualise the arbitrary arrangement of land use forms and locations.

Each hexagon cell covers an area of 4156 sq. m., or 0.41 hectares. The hexagon cells in the CROME layer are spatially distinct units and the CROME layer doesn't provide any spatial adjacency information. The vertices of adjacent cells are mostly coincident; therefore the CROME layer provides a continuous representation of the land use. The hexagon cells are not constrained by any topographic features, except the extent of the land.

Figure 1. Example of CROME hexagonal classification cells

2.3 Attributes

In the creation of the CROME layer, attribution is generated from the associated imagery and the feature classification process.

3 Reference Layer Features

3.1 Representation

The classification is provided as hexagonal polygon cells, with each cell being attributed a land cover classification, according to the classification in Annex A. An example of the landscape and the classification cells, labelled with the land cover code in Annex A, is shown in Figure 1.

It was realised early on that it would be computationally non-trivial and theoretically problematic to classify the crop types for the entire country in a single process. The reasoning behind the latter logic was that horticulturally, crops and cropping patterns in England vary dependent upon the climatic zones in England (see Figure 2).

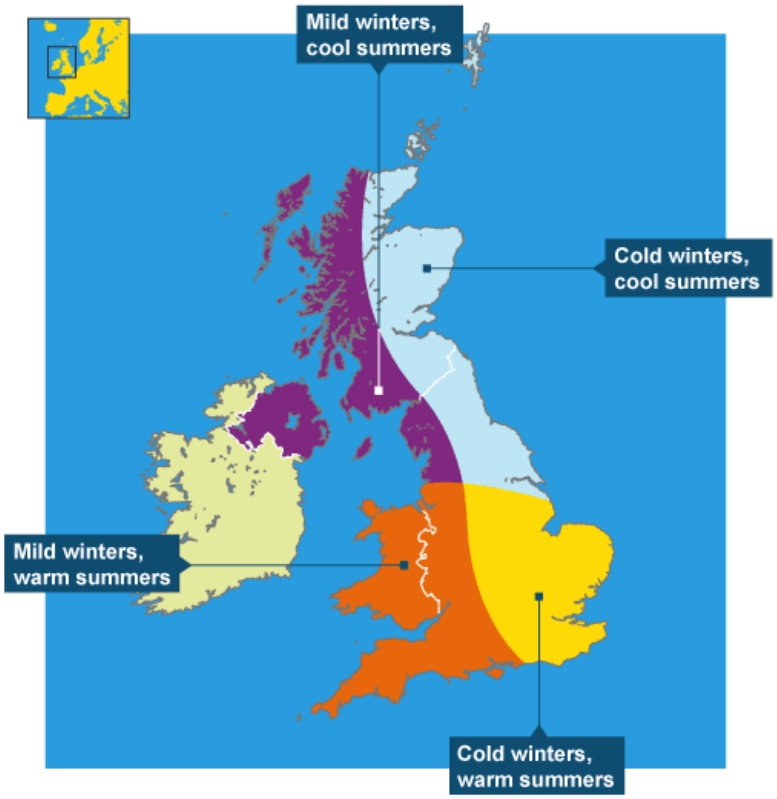

Figure 2. General climate zones in the UK (Source: BBC)

Therefore, CROME was classified by splitting England into different “agricultural-climate” zones. Major roads were used to ensure that a parcel can only belong to one of the zones.

Table 1 lists the number of CROME cells associated with each CROME zone. Please note that some zone names and their extents don’t align with other established administrative or census extents associated with zone names.

Zone Name	Zone Code	Number of Cells
South East	SE	9328065
South West	SW	7431895
East Midlands	EM	4989891
North East	NE	4618260
West Midlands	WM	3831363
North West	NW	1080887
Isle of Wight	IW	91480
Scilly Isles	SC	3947
Lundy Isles	LN	1018
Farne Isles	FR	77
Hauxley Isles	HX	20
	Total	**31376903**

Table 1. CROME classification Zones with their zone codes and cell counts respectively

Figure 3. CROME classification Zones

4 CROME Classification

The classification of the topography was performed using Random Forest classification, a supervised machine learning approach. More information on Random Forest Classification can be found on the authors' homepage. The approach learns the classification of specific crop types by associating backscatter characteristics of each radar polarisation (VV, VH, VV/VH) to known land covers, previously collected by field inspectors. The ground truth points and radar images were fed into the workflow, outlined in Figure 4, making up the four main work packages described below in further detail.

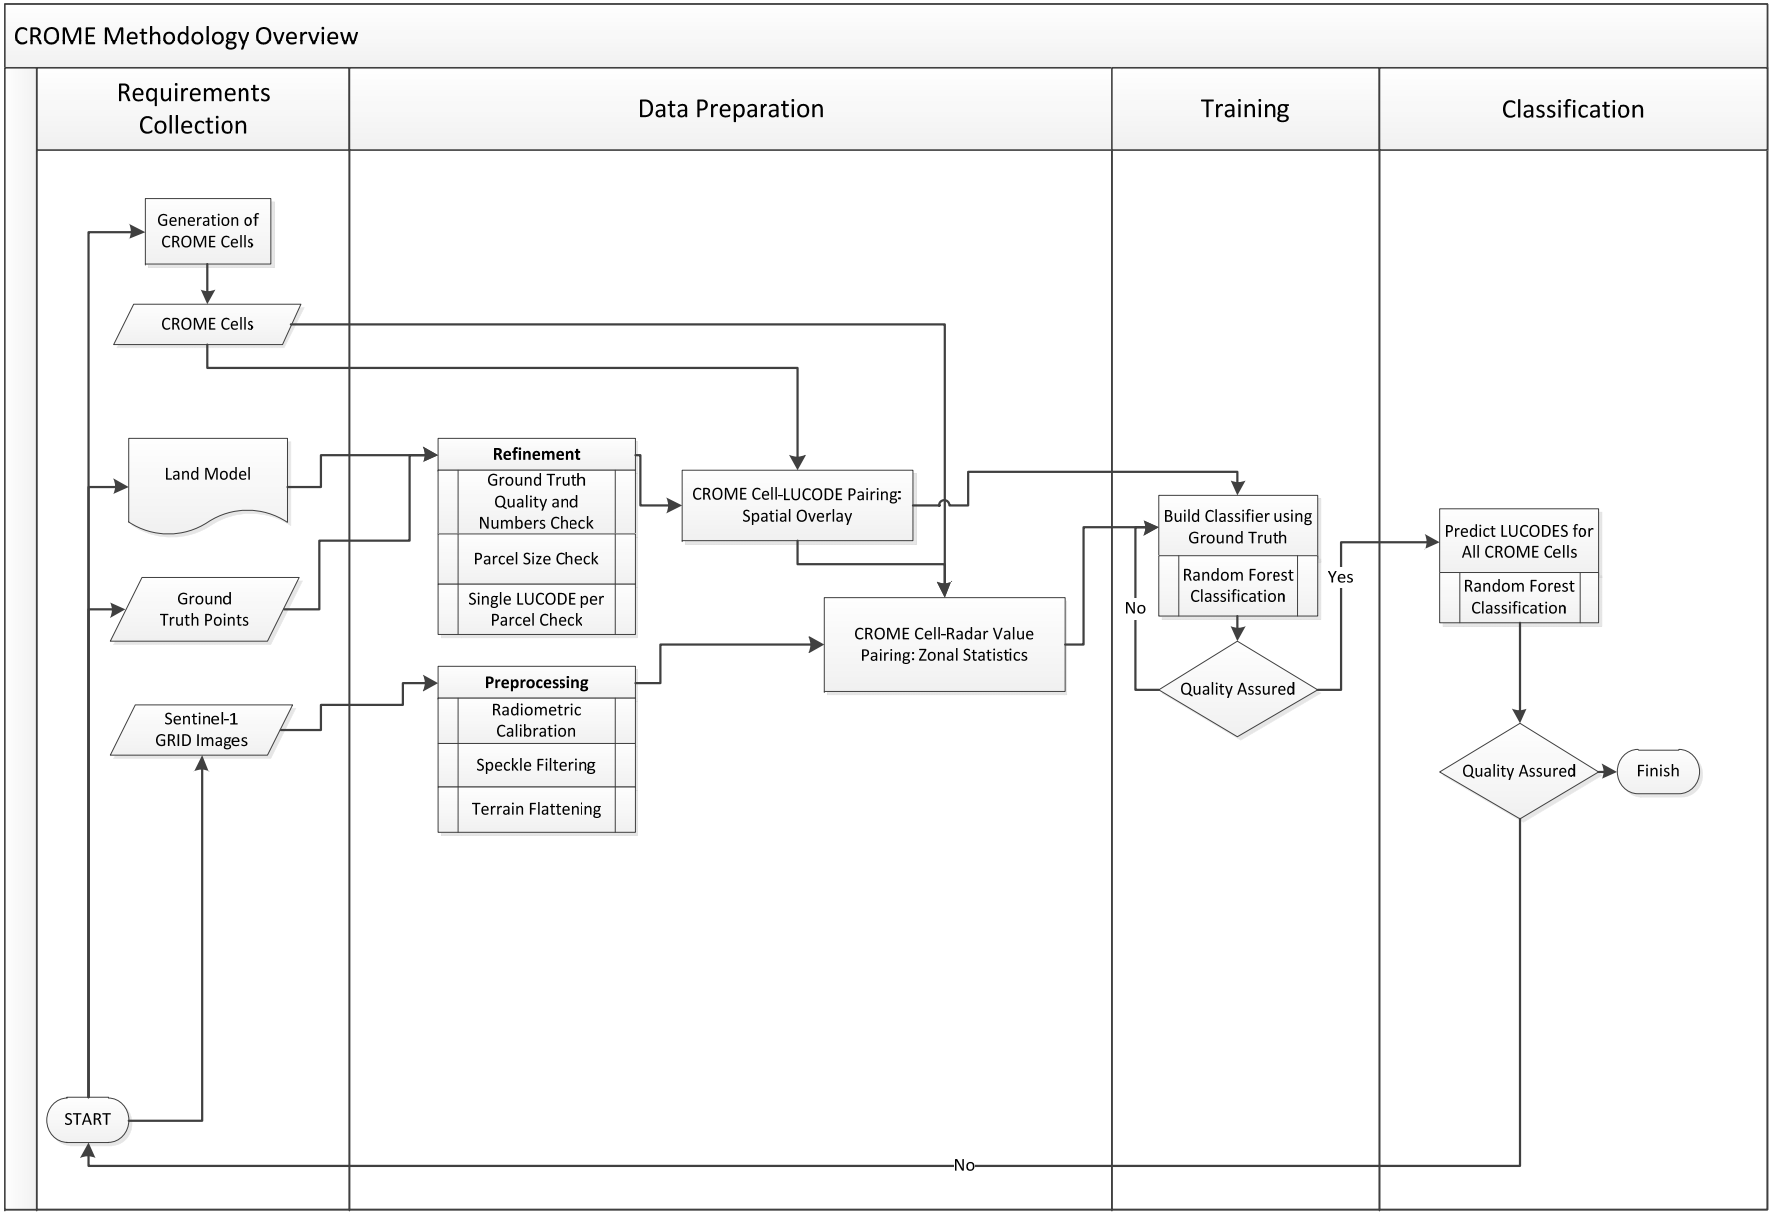

Figure 4. Workflow for CROME classification

4.1 Requirements Collection

This step involves the creation and collection of the input datasets and specifications required to perform the classification. The step involves four main inputs:

a) CROME Cells: These are hexagon cells that form the spatial framework of the CROME layer. These are essentially Voronoi polygons and were created using a combination of ESRI ArcMap and FME applications.
b) Land Model (current version 1.6.4): This is a document that contains a list and specification of the relevant land use and crop types which are eventually assigned to the CROME cells.
c) Ground Truth Points: These are a record of the spatial location and type of crop and land use found by the RPA field inspectors during late-spring and summer period of 2016. The ground truth collection process involves verification and, where required, correction of the crop and other land cover types declared by the farmers against each parcel registered under their name. In order to minimise classification errors, only the ground truth data and parcels that matched the following conditions were used:
    a. Parcels only contained a single crop/land use type.
    b. Parcel area was bigger than the CROME cell area.
    c. The crop/land use type was present in at least 10 parcels.

    In addition, the numbers and varieties for non-agricultural areas were boosted by collecting additional ground truth information, such as man-made surfaces and structures, from 2016 very high resolution satellite. A total of 18271 ground truth samples were used for building the Random Forest classifier.

d) Sentinel-1 GRD images: These are Sentinel-1 Radar Ground Range Detected (GRD) dual polarisation (VV/VH) images for the period of January till August 2016, covering the main land masses of England. A total of 672 images were used. These can be downloaded from the scihub (scihub.copernicus.eu).

4.2 Data Preparation

The objective of this step is to attribute the CROME Cells with the required LUCODE (only for cells used during supervised training) and radar values for the purposes of classification. It involves the following main processes:

a) Ground truth data points are filtered by applying a combination of simple attribute and spatial filtering using farm field polygons.
b) The Sentinel-1 GRD images are processed using the Sentinel Application Platform (more information available on SNAP homepage) software to transform the raw images into geoferenced and radiometrically corrected Sigma-0 (backscattering coefficient) images.
c) “Zonal Statistics” analysis is performed to calculate the monthly mean and standard deviation of VV, VH and VV/VH-ratio image values for the area covered by each CROME cell. Each CROME cell (32 million in total) is assigned a pair of triplet values i.e. mean VV, VH and VV/VH, and VV, VH, VV/VH standard deviation. Statistics were calculated of each polarization combinations (VV, VH and VV/VH) from 672 radar images captured from January to October 2016 (totaling 2016 images). Monthly statistics allowed a common temporal reference axis for all cells because although satellite pass dates over cells are broadly predictable as per ESA guidance, i.e. within 10 days revisit, these still varied considerably across England thereby creating data gaps.

d) Spatial overlap is tested between the selected ground truth subset points and all the CROME cell polygons to assign the known LUCODES to the training-CROME cell polygons.

Classification rules were formulated based on the known LUCODES, established in step (d), and the monthly radar backscattering coefficient values derived in step (c) from the training CROME cells.

4.3 Training

The objective of this step is to automatically build the classification rules that can be used to predict the LUCODE of a CROME cell, given its associated monthly radar backscattering coefficient values.

This work uses the Random Forest classification technique, a supervised machine learning approach, available in the R application. The unique advantage of Random Forest classification is that given a large number of ground truth points, it is able to perform robust cross-validation internally by generating numerous decision trees.

The training process involves randomly combining input variables and deriving a permutation of input variables that provides the most instances of accurate match between the known LUCODE and a predicted LUCODE. The software internally uses 2/3rds of the input variables for developing and training the ruleset, and the remaining 1/3rd for testing the classification.

4.4 Classification

Each CROME cell is then applied a LUCODE from the Random Forest classification ruleset. A random visual check using 2016 Sentinel-2 images was performed to detect obvious misclassifications. Simple cases of misclassifications (e.g. slightly rough areas of manmade surfaces conflicted with grass and vice versa) were corrected by calculating normalized difference vegetation index (NDVI).

5 Features Types and Attribution

5.1 Attribution

Table 2 shows the attribute for each CROME cell:

Name	Type	Properties	Comments
CROMEID	TEXT	Primary Key	Unique across all supply years
REFDATE	NUMBER	Not Null	The date or date range of the images captured and used to assign the classified land cover to the cell.
LUCODE	TEXT	Not Null	The land use code. See Annex A for lookup tables.
SHAPE	GEOMETRY	Not Null	Polygon representing the extent of the classified land use object.

Table 2. Attribute Schema of CROME layer

5.1.1 CROMEID

The CROMEID is the unique identifier assigned to each cell, consisting of the letters 'RPA' and a twelve digit number, which are the Easting and Northing coordinates of the centroid of the cell.

5.1.2 REFDATE

The REFDATE is the date or date range of the images referenced to assign the land use code to the CROME cell. The date format used is YYYYMMDDYYYYMMDD e.g. 20161001201610826. The value is expected to remain constant for each cell in a CROME release.

5.1.3 LUCODE

The Land Use Code is alphanumeric code of up to 5 characters long, and is intended to record land use information. At present, all non-agricultural areas will be assigned a generic non-agricultural land use code. These land use codes are based on the land model used by the RPA as part of the claim purpose. The table in Annex A contains the valid Land Use Code (Column 3) that will be used in the current release of the CROME.

5.2 Data Format and Naming

The dataset is supplied in the ESRI shapefile format, together with a valid projection file (*.prj) that defines the coordinate system as British National Grid. The data for a given year shall be supplied either as a single shapefile or as a series of shapefiles that collectively cover all CROME zones.

The dataset format shall be consistent with the specification in Section 5.1, with all fields whose source is marked as shapefile being included in the supplied dataset.

The shapefile name shall be as follows:

CROME_<YYYYMMDD>_<ZONECODE>.shp

where:

<YYYYMMDD> is the release date of the dataset

<ZONE> is the mnemonic for each CROME zone covered by the shapefile.

Each shapefile will be accompanied by a metadata xml file in the accordance to the UK GEMINI v 2.2 standard.

Please note that the final packaged filenames may be different due to other requirements.

6 Thematic Accuracy

Thematic accuracy is concerned with the degree to which the classification and attribution of features correctly match the real world features that are being modelled within the dataset.

The quality of CROME land code classification was assessed by comparing the crop/land use types predicted by the Random Forest Classification against the ground truth data collated by the field inspectors from Rural Payments Agency. The comparison is reported in the form of a confusion matrix, with the overall accuracy and Kappa Coefficient presented in Table 3 below. The top 10 declared arable land cover types, which make up 92% of all claimed areas in 2016 are indicated in the highlighted cells. The confusion matrix was measured on the majority of CROME cells within each RPA land parcel. CROME v2016.1 has an overall accuracy of 84% and a Kappa coefficient of 82%, which confirms that a high quality of land use classification is achievable with Sentinel-1 radar data, complimented with Sentinel-2 for distinguishing grass/rough surfaces. Each CROME zone will also contain a separate confusion matrix specific to that zone to provide regional overview of the quality of the layer.

The Ground Truth data didn't provide information on non-agricultural land covers such as Woodland, water bodies and general non-agricultural areas (e.g. manmade surfaces) therefore the confusion matrix doesn't provide accuracy numbers for each land use. Samples for these land uses were collected from other reliable sources e.g. non-Remote Sensing inspections, and were only used during the training to minimise confusion with similar classes.

7 Dataset Specification

7.1 Format

The CROME datasets are in ESRI shapefile format.

7.2 Topological Consistency

No topological consistency errors are known to exist with CROME, i.e. each feature is valid according to the OGC specification.

7.3 Horizontal Positional Accuracy

Not applicable

7.4 Currency

The map is to be updated biannually (April/May and July/August).

7.5 Format Consistency

N/A

7.6 Domain Consistency

N/A

7.7 Temporal Consistency

There are no known temporal consistency issues with CROME.

7.8 Temporal Validity

There are no known temporal validity issues with CROME.

7.9 Attribute Completeness

The version of the dataset has 100% attribute completeness.

7.10 Spatial Completeness

The version of the dataset has 100% spatial completeness, as verified by visual inspection.

7.11 Lineage

Crop types were derived by using Random Forest Classification on Sentinel-1 Radar data, and refined using vegetation indices from Sentinel-2 optical data.

7.12 CROME Layer Metadata

A metadata file will be supplied with each product supply. Metadata will conform to ISO 19115 and be UK GEMINI discovery level metadata.

		Ground Truth																															
		AC01	AC03	AC07	AC16	AC17	AC19	AC20	AC23	AC30	AC32	AC34	AC36	AC44	AC50	AC58	AC63	AC65	AC66	AC67	AC68	AC69	FA01	LG03	LG07	LG11	LG14	LG20	NA01	PG01	TC01	WO12	Total	User Accuarcy
CROME	AC01	561	3	0	0	2	17	0	0	5	11	0	0	1	0	1	6	0	4	2	0	2	8	1	2	0	1	2	0	1	0	0	630	89
	AC03	0	196	0	0	1	1	3	0	0	1	1	0	14	0	0	1	0	1	1	0	0	7	2	2	0	0	0	0	0	0	0	231	85
	AC07	0	0	3	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	0	0	3	100
	AC16	0	0	0	44	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	45	98
	AC17	2	7	0	4	264	0	0	0	0	1	1	0	4	0	1	0	0	0	0	0	0	10	0	0	0	1	0	0	2	0	0	297	89
	AC19	0	0	0	0	0	67	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	68	99
	AC20	0	0	0	0	0	0	10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	10	100
	AC23	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	0	0	0	0	0	0	N/A
	AC30	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	0	0	0	0	0	0	N/A
	AC32	2	0	0	1	0	8	0	0	0	95	0	0	0	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	1	0	0	111	86
	AC34	0	0	0	0	0	0	0	0	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	4	100
	AC36	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3	100
	AC44	0	12	1	0	2	0	0	0	0	0	0	0	125	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	140	89
	AC50	0	0	0	0	0	0	0	0	0	0	0	0	0	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	5	100
	AC58	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	100
	AC63	7	1	0	0	1	0	0	0	0	4	0	0	0	0	0	680	3	31	8	9	0	13	1	1	0	2	0	0	6	0	0	767	89
	AC65	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	38	0	0	0	0	1	0	0	0	0	0	0	0	0	0	39	97
	AC66	9	1	0	0	1	5	0	0	0	27	0	0	0	0	1	19	64	1822	14	0	4	21	3	0	1	1	5	0	8	1	0	2007	91
	AC67	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3	706	5	0	11	2	0	0	0	2	0	0	0	0	729	97
	AC68	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	18	0	1	0	0	0	0	0	0	0	0	0	19	95
	AC69	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1	100
	FA01	7	0	0	2	5	3	0	0	3	1	1	0	0	0	1	3	1	2	2	0	0	388	5	0	2	1	0	0	10	1	0	438	89
	LG03	4	5	1	0	4	0	0	0	0	3	2	2	3	0	0	3	0	1	1	0	0	9	441	30	0	0	8	0	4	0	0	521	85
	LG07	1	0	0	0	0	0	4	0	0	0	0	0	3	0	0	0	0	0	1	1	0	3	4	168	0	0	2	0	0	0	0	187	90
	LG11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	4	0	0	0	0	0	0	4	100
	LG14	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	11	0	0	0	0	0	11	100
	LG20	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	1	3	0	0	2	2	0	0	0	138	0	1	0	0	148	93
	NA01	17	0	2	3	15	14	0	0	0	18	0	2	3	0	1	4	0	6	1	2	0	47	12	2	3	0	0	0	39	0	0	191	N/A
	PG01	5	0	0	1	0	5	0	0	0	3	0	1	1	0	2	6	1	6	0	1	2	63	2	0	5	11	0	0	950	1	2	1068	89
	TC01	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	0	0	0	7	0	7	100
	WO12	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1	0	0	0	9	0	0	0	0	0	0	7	1	0	19	N/A
	Total	615	225	7	55	295	120	17	0	8	165	9	9	155	5	9	722	107	1882	739	36	9	594	475	205	15	28	157	0	1029	11	2	7705	
	Producer Accuracy	91	87	43	80	89	56	59	N/A	0	58	44	33	81	100	22	94	36	97	96	50	11	65	93	82	27	39	88	n/a	92	64	N/A		

Table 3. Confusion Matrix of CROME v 2016.2

Annex A – Land Cover Codes

Land Cover Description	LUCODE	Land Use Description
Cereal Crops	AC01	Barley - Spring
	AC03	Beet
	AC07	Carrot
	AC16	Linseed – Spring
	AC17	Maize
	AC19	Oats – Spring
	AC20	Onions
	AC23	Parsnips
	AC30	Triticale – Spring
	AC32	Wheat – Spring
	AC34	Cabbage – Spring
	AC36	Oilseed – Spring
	AC44	Potato
	AC50	Squash
	AC58	Mixed Crop – Group 1
	AC63	Barley - Winter
	AC65	Oats – Winter
	AC66	Wheat – Winter
	AC67	Oilseed – Winter
	AC68	Rye – Winter
	AC69	Triticale – Winter
Leguminous Crops	LG03	Field beans – Spring
	LG07	Peas – Spring
	LG11	Lucerne
	LG14	Clover
	LG20	Field beans – Winter
Grassland	FA01	Fallow Land
	PG01	Permanent Grassland
Perennial Crops	TC01	Perennial Crops
Non–Agricultural Land	NA01	Non-Agricultural Land
Water	WA01	Water
Trees	WO12	Woodland

Note that the variety of land cover codes may change during releases of CROME due to variability in ground truth however the LUCODE will remain same.

Acknowledgements

We would like to acknowledge support and guidance from the following sources and individuals:

- ESA's Scihub and Alaska Satellite Facility for download access to the Sentinel-1 Data;
- Guido Lemoine (Joint Research Centre, Italy) and Andrea Minchella (previously at Satellite Applications Catapult, UK) for pointers on how to batch process large number of Sentinel-1 images.
- SNAP Discussion forum members on very helpful tips on improving the performance of batch processing and issues surround Sentinel-1 radar data.

- Defra Centre of Excellence in Earth Observation, and in particular Gwawr Jones (Joint Nature Conservation Committee, UK) for providing Sentinel-2 images for visual quality assurance of CROME.
- Asger Petersen and Gregers Petersen (Septima.dk, Denmark) for advice on how to use GDAL and parallel processing of large amount of geospatial data.