

Technical Note

Imbalanced Learning in Land Cover Classification: Improving Minority Classes' Prediction Accuracy Using the Geometric SMOTE Algorithm

Georgios Douzas , Fernando Bacao ^{*}, Joao Fonseca [†] and Manvel Khudinyan [†]

NOVA Information Management School (NOVA IMS), Campus de Campolide, Universidade Nova de Lisboa, 1070-312 Lisboa, Portugal; gdouzas@novaims.unl.pt (G.D.); jpfonseca@novaims.unl.pt (J.F.); mkhudinyan@novaims.unl.pt (M.K.)

^{*} Correspondence: bacao@novaims.unl.pt; Tel.: +351-21-382-8610

[†] These authors contributed equally to this work.

Received: 23 October 2019; Accepted: 12 December 2019; Published: 17 December 2019



Abstract: The automatic production of land use/land cover maps continues to be a challenging problem, with important impacts on the ability to promote sustainability and good resource management. The ability to build robust automatic classifiers and produce accurate maps can have a significant impact on the way we manage and optimize natural resources. The difficulty in achieving these results comes from many different factors, such as data quality and uncertainty. In this paper, we address the imbalanced learning problem, a common and difficult conundrum in remote sensing that affects the quality of classification results, by proposing Geometric-SMOTE, a novel oversampling method, as a tool for addressing the imbalanced learning problem in remote sensing. Geometric-SMOTE is a sophisticated oversampling algorithm which increases the quality of the instances generated in previous methods, such as the synthetic minority oversampling technique. The performance of Geometric-SMOTE, in the LUCAS (Land Use/Cover Area Frame Survey) dataset, is compared to other oversamplers using a variety of classifiers. The results show that Geometric-SMOTE significantly outperforms all the other oversamplers and improves the robustness of the classifiers. These results indicate that, when using imbalanced datasets, remote sensing researchers should consider the use of these new generation oversamplers to increase the quality of the classification results.

Keywords: imbalanced learning; LULC classification; oversampling; Geometric-SMOTE; class imbalance

1. Introduction

The production of accurate land use/land cover (LULC) maps offers unique monitoring capabilities within the remote sensing domain [1]. LULC maps are being used for a variety of applications, ranging from environmental monitoring, land change detection, and natural hazard assessment to agriculture and water/wetland monitoring [2]; therefore, accurate and timely production of LULC maps is of great significance. LULC maps are usually produced by two main procedures: photo-interpretation by the human eye, which is time and resource consuming and is not suitable for operational LULC-mapping over large areas; and second, automatic mapping using remotely sensed data and different classification algorithms.

The availability and a swift update of high-quality satellite remote sensing data has brought tremendous progress in providing up-to-date and accurate land cover information. Multispectral images, particularly, are an essential resource for building LULC maps, allowing for the use of classification algorithms to automate their production. Although significant progress has been made

in the use of supervised learning techniques for automatic image classification [3], the acquisition of labeled training sets continues to be a bottleneck [4]. In order to build accurate and robust supervised classifiers it is crucial to have a large enough training dataset. Often, the problem is that different land cover types have very different levels of area coverage, which causes some of them to be frequent in the training dataset, while others are limited [5].

A particular case where this phenomenon happens is the LUCAS dataset: Land Use and Coverage Area frame Survey coordinated by The Statistical Office of the European Commission (Eurostat) [6]. LUCAS surveys have been carried out every three-years since 2006 and are freely accessible. For this statistical sampling survey, a 2 km regular grid is implemented, and over 1,000,000 points were observed in the European Union territory for the year of 2015. Although the LUCAS dataset is designed for statistical estimation, some existing studies used this data for training machine learning classifiers for land cover classification successfully [7,8], since each observation is empirically registered in the field (in situ). This sampling strategy is particularly interesting for this research, as it causes uneven representation of different land cover classes in the dataset for the given area.

The above-mentioned asymmetry in class distribution affects the performance of classifiers negatively. In the machine learning community, the problem is known as imbalanced learning problem [9]. The imbalanced learning problem generally refers to a skewed distribution of data across classes in both binary and multiclass problems [10]. The latter, in particular, appears to be an even more challenging task [11]. In both cases, during the learning phase, the minority class(es) contribute less to the minimization of accuracy, the typical objective function, inducing a bias towards the majority class. Consequently, as typical classification algorithms are designed to work with reasonably balanced datasets, learning the decision boundaries between different classes becomes a very difficult task [12].

The possible approaches to deal with the class imbalance problem can be divided into three main groups [13]:

1. Cost-sensitive solutions. They introduce a cost matrix that applies higher misclassification costs for the examples of the minority class.
2. Algorithmic level solutions. They modify the algorithmic procedure to reinforce the learning of the minority class.
3. Resampling solutions. They rebalance the class distribution either by removing instances from the majority class or by generating artificial data for the minority class(es).

The latter method constitutes a more general approach, since it can be used for any classification algorithm and it does not require any type of domain knowledge in order to construct a cost matrix.

There are several resampling solutions to deal with the imbalanced learning problem, which also can be divided into three categories:

1. Undersampling algorithms reduce the size of the majority class.
2. Oversampling algorithms attempt to even the distributions by generating artificial data for the minority class(es).
3. Hybrid approaches use both oversampling and undersampling techniques to ensure a balanced dataset.

In this paper, we compare the performance of various oversampling algorithms on EUROSTAT's publicly available Land Use/Cover Area Statistical Survey (LUCAS) dataset [14] with Landsat 8 data. The experimental procedure included a comparison of five oversamplers using five classifiers and three evaluation metrics. Specifically, the oversampling algorithms were Geometric-SMOTE (G-SMOTE) [15], the synthetic minority oversampling technique (SMOTE) [16], Borderline-SMOTE (B-SMOTE) [17], the adaptive synthetic sampling technique (ADASYN) [18] and random oversampling (ROS), while no oversampling was included as a baseline method. Results show that G-SMOTE outperforms every other oversampling technique, for the selected evaluation metrics.

This paper is organized in five sections: Section 2 analyzes the resampling methods, Section 3 describes the proposed methodology, Section 4 shows the results and discussion, and Section 5 presents the conclusions drawn from this study.

2. Resampling Methods

Data modification through resampling has been the most popular approach to deal with the imbalanced learning problems in machine learning in general and remote sensing in particular [5]. As mentioned above, by decoupling the imbalance problem from the classification algorithms, resampling allows the users to apply any standard algorithm once the resampling preprocessing step is done. This stratagem is especially convenient for users that are not machine learning experts and want to use several classifiers. Additionally, resampling methods can be naturally applied to multiclass imbalanced data, which is relevant for LULC classification. In this section, we present the most relevant applications of resampling methods for imbalanced remote sensing data classification.

2.1. Random Resampling

Random resampling refers to non-informed strategies that remove instances from the majority class or replicate instances from the minority class. As such, the selection of the data occurs randomly without exploiting any additional information.

Some of the existing remote sensing studies implement the random undersampling (RUS) method [19], which randomly reduces the number of the majority class training samples. However, this method has the disadvantage of information loss, as it discards samples from the majority class [5]. Contrary to RUS, ROS is a method that can be considered equivalent to Bootstrapping, as it avoids information loss. However, ROS simply replicates randomly selected instances of the minority class, increasing the risk of overfitting [20]. Reference [21] reports that balancing data with ROS affects the classification performance differently for various classifiers. In their study, land cover classification with highly imbalanced data was carried out with six different models. The application of ROS slightly improved the performance of the random forest (RF) and support vector machine (SVM) classifiers. On the other hand, it reduced the classification accuracy for ecision tree (DT), artificial neural network (ANN), k-nearest neighbors (KNN) and boosted DT classifiers.

2.2. Informed Resampling

In the above section, the disadvantages of RUS and ROS have been pointed out. Informed resampling methods aim to overcome these insufficiencies. More specifically, they use the local or global information of the class distribution to remove or generate instances. Our focus is on oversampling algorithms, since the size of the LUCAS dataset does not favor the use of undersampling approaches. Additionally, reference [22] carried out a comparative analysis of undersamplers' and oversamplers' performance for land cover classification with the rotation forest ensemble classifier, showing that oversampling methods outperform undersampling methods.

SMOTE is the most popular informed oversampling method, and it has been used to successfully deal with the class imbalance problem in land cover classification [23]. In this approach, the minority class is oversampled by randomly selecting a minority class instance and generating synthetic examples along the line segment joining it with one of its minority class neighbors. A number of studies report significant improvements in LULC mapping accuracy with the use of SMOTE oversampling. For instance, the variational semi-supervised learning (VSSL) proposed by [23] aims to deal with the imbalance problem in LULC mapping. VSSL is a semi-supervised learning framework consisting of a deep generative model. It allows learning successfully from both labeled and unlabeled samples while using SMOTE to balance the data. In [24], they used OpenStreetMap crowdsourced data and Landsat time series for LULC classification. Similarly, the application of SMOTE improved the classification results. Other examples of the successful application of SMOTE in remote sensing can be found in [25,26].

Although recent studies demonstrate the usefulness of SMOTE for remote sensing applications, it still has some drawbacks. The SMOTE algorithm has the disadvantage of generating noisy data [27]. In order to mitigate this problem, many variations of SMOTE have been developed. B-SMOTE is one of the most popular SMOTE-based oversamplers. Similarly to SMOTE, it uses the k -nearest neighbors selection strategy. The main difference to the original algorithm is that it modifies the data generation mechanism by generating samples closer to the decision boundary. B-SMOTE has also been reported to perform better than SMOTE in a number of studies [28,29]. ADASYN is another well-known variation of SMOTE. It is based on the idea of adaptively generating minority class instances according to their weighted distribution: more instances are generated for those minority class instances that are harder to learn compared to ones that are easier to learn [18].

The SMOTE algorithm can be decomposed into two parts: the selection strategy for the minority class instances and the data generation mechanism. The first part is related to the generation of noisy instances since the SMOTE selection strategy considers all the minority samples as equivalent. The above-mentioned SMOTE variations (B-SMOTE and ADASYN) aim to deal with this problem. On the other hand, the second part is responsible for the diversity of the artificial instances. There are scenarios where the linear interpolation mechanism used in SMOTE generates nearly duplicate instances that may lead to overfitting. The G-SMOTE algorithm is an extension of SMOTE that aims to deal with both problems. G-SMOTE defines a flexible geometric region around each minority class instance for synthetic data generation. The shape of this area is controlled by a set of hyperparameters. This element significantly increases the diversity of instances generated. Furthermore, G-SMOTE is designed to avoid noisy sample generation since it modifies the SMOTE selection strategy. G-SMOTE has been shown to outperform SMOTE and its above-mentioned variations across 69 imbalanced datasets for various classifiers and evaluation metrics. Figure 1 depicts the data generation mechanisms of both SMOTE and G-SMOTE using a deformed geometric region.




Figure 1. Example of minority class oversampled by SMOTE and G-SMOTE algorithms. G-SMOTE generates non-noisy samples with greater variety than SMOTE.

3. Methodology

This section describes the evaluation process of G-SMOTE's performance. A description of the study area, dataset, oversamplers, classifiers, evaluation metrics, and the experimental procedure is provided. Figure 2 represents the flowchart of the steps applied in this experiment.




Figure 2. Flowchart containing the steps applied in the entire method.

3.1. Study Area

The area of study was within north-western Portugal, corresponding to the area covered by the Landsat 8 image from track 204 and row 32, shown in Figure 3. The area contains all eight main land cover types defined by LUCAS 2015: artificial land, cropland, woodland, shrubland, grassland, bare land, water, and wetlands.




Figure 3. Study area and LUCAS 2015 reference data (coordinate system: WGS-84 UTM Zone 29, projection: Transverse-Mercator, Landsat image acquisition date: 25 May 2015).

3.2. Remote Sensing Data

The remotely sensed data includes eight images from the moderate-resolution Landsat 8 multi-spectral sensor. The images are Level-2 surface reflectance products (OLI/TIRS); one image was acquired each month from February to September 2015. The acquisition mode was descending. Data were pre-processed in order to remove pixels with cloud cover. Only bands 2, 3, 4, 5, 6, and 7 were used from each image. Accordingly, each reference point from the LUCAS dataset had 48 features, representing pixel values from each spectral band from each image.

3.3. LUCAS Dataset

The 2015 LUCAS data was used as reference data for both model training and validation. The LUCAS point label represents the corresponding land cover/use type within the radius of 1.5 m for homogeneous classes and a 20 m radius extent (“extended window”) for heterogeneous classes (e.g., shrubland), gathered by field observation and a very high-resolution photo interpretation [6]. In order to reduce the risk of having Landsat pixel information represented wrongly in the field, we only kept points observed in situ from a close distance (<100 m). With the same objective we removed the points which had linear features in the observation (e.g., roads). This procedure was solely not applied to the class of “artificial land”, as this would have removed most parts of the samples. Furthermore, points with cloudy pixels in the Landsat data were also excluded. This way, 1694 out of 2060 LUCAS points were retained. This dataset contains eight classes that represent the main land cover types for the study area.

This pixel selection excluded a large number of unacceptable reference points, and we assumed the remaining ones to be suitable enough to represent the land cover type in a Landsat pixel coverage

area of 30×30 m. Further, we surmised that classifiers are capable of overcoming the noise caused by pixels having mixed land cover representation if such pixels are still available in the dataset.

The number of samples per class and the imbalance ratio (IR), defined as the ratio of the number of samples for the majority class over the number of samples for any of the minority classes, is presented in Table 1.

Table 1. LUCAS nomenclature and classes' distributions.

LUCAS Category	Land Cover Type	Instances	IR
A	Artificial land	131	5.81
B	Cropland	270	2.81
C	Woodland	761	1.00
D	Shrubland	296	2.61
E	Grassland	185	4.11
F	Bareland	37	20.56
G	Water	10	76.10
H	Wetlands	4	190.25

Table 2 presents a description of the LUCAS dataset, including information about the majority class C and the smallest minority class H to emphasize the imbalanced character of the dataset:

Table 2. Description of the LUCAS dataset.

Dataset	LUCAS
Features	47
Instances	1694
Instances of class C	761
Instances of class H	4
IR of class H	190.25

3.4. Evaluation Metrics

Amongst the possible choices existing for a classifier's performance evaluation, *Accuracy*, user's accuracy (or *Precision*) and producer's accuracy (or *Recall*) are the most common in LULC classification [30,31]. For a binary classification task, their calculation is given in terms of the true positives TP , true negatives TN , false positives FP , and false negatives FN [30]. More specifically, $Precision = \frac{TP}{TP+FP}$ and $Recall = \frac{TP}{TP+FN}$. For the multiclass case, the average value across classes is used, as explained below.

The LUCAS dataset is highly imbalanced, having a wide range of IRs for the different minority classes. Therefore, the use of the metrics above is not an appropriate choice since they are mainly determined by the majority class contribution [32]. An appropriate evaluation metric should consider the classification accuracy of all classes. A simple approach for the multiclass case is to select a binary class evaluation metric; apply it to each binary sub-task of the multiclass problem, i.e., consider each class versus the rest; and finally, average its values. For this purpose, *F-score* and *G-mean* metrics were used as the primary evaluation methods, while *Accuracy* is provided for discussion:

- The *Accuracy* is the number of correctly classified samples divided by the sum of all samples. Assuming that the various classes are labeled by the index c , *Accuracy* is given by the following formula:

$$Accuracy = \frac{\sum_c TP_c}{\sum_c (TP_c + FP_c)}.$$

- The *F-score* is the harmonic mean of *Precision* and *Recall*. The *F-score* for the multiclass case can be calculated using their average per class values [32]:

$$F\text{-score} = 2 \frac{\overline{Precision} \times \overline{Recall}}{\overline{Precision} + \overline{Recall}}.$$

- The *G-mean* is the geometric mean of *Sensitivity* and *Specificity*. *Sensitivity* is identical to the *Recall* while *Specificity* is given by the formula $Specificity = \frac{TN}{TN+FP}$. Therefore, they are equal to the true positive and true negative rates, respectively. The *G-mean* for the multiclass case can be calculated using their average per class values:

$$G\text{-mean} = \sqrt{\overline{Sensitivity} \times \overline{Specificity}}.$$

3.5. Machine Learning Algorithms

The main objective of the paper is to show the effectiveness of G-SMOTE when it is used on multiclass, highly imbalanced data of a remote sensing application and to compare its performance to other oversampling methods. Four oversampling algorithms were used in the experiment along with G-SMOTE. ROS was chosen for its simplicity. SMOTE was selected for being the most widely used oversampler. ADASYN and B-SMOTE were selected for representing popular modifications of the original SMOTE algorithm. Finally, no oversampling was applied as an additional baseline method.

For the evaluation of the oversampling methods, the classifiers logistic regression (LR) [33], k-nearest neighbors (KNN) [34], decision tree (DT) [35], gradient Boosting classifier (GBC) [36], and random forest (RF) [37] were selected. The choice of classifiers was made according to the following criteria: learning type, training time, and popularity within the remote sensing community. All these algorithms were found to be computationally efficient and commonly used for the proposed task, with the exception of LR, which is rarely used in remote sensing applications [2,21].

3.6. Experimental Settings

In order to evaluate the performance of each oversampler, every possible combination of oversampler, classifier, and metric was formed. The evaluation score for each of the above combinations was generated through an n -fold cross-validation procedure with $n = 3$. Before starting the training of each classifier, and in each stage $i \in \{1, 2, \dots, n\}$ of the n -fold cross-validation procedure, synthetic data S_i were generated using the oversampler, based on the training data T_i of the $n - 1$ folds, such that the resulting $S_i \cup T_i$ training set became perfectly balanced. This enhanced training set, in turn, was used to train the classifier. The performance evaluation of the classifiers was done on the validation data V_i of the remaining fold, where $V_i \cup T_i = D$, $V_i \cap T_i = \emptyset$ while D represents the dataset. The process above was repeated three times, and the results were averaged.

The range of hyperparameters used for each classifier and oversampler are presented in Table 3:

Table 3. Hyperparameters grid.

Classifier	Hyperparameters	Values
LR	maximum iterations	10,000
KNN	number of neighbors	3, 5
DT	maximum depth	3, 6
GBC	maximum depth	3, 6
	number of estimators	50, 100
RF	maximum depth	None, 3, 6
	number of estimators	50, 100
Oversampler		
G-SMOTE	number of neighbors	3, 5
	selection strategy	combined, minority, majority
	truncation factor	−1.0, −0.5, 0, 0.25, 0.5, 0.75, 1.0
	deformation factor	0, 0.2, 0.4, 0.5, 0.6, 0.8, 1.0
SMOTE	number of neighbors	3, 5
BORDERLINE SMOTE	number of neighbors	3, 5
ADASYN	number of neighbors	2, 3

3.7. Software Implementation

The implementation of the experimental procedure was based on the Python programming language, using the Scikit-Learn [38], Imbalanced-Learn [39], and Geometric-SMOTE libraries. All functions, algorithms, experiments, and results reported are provided at the GitHub repository of the project. Additionally, the Research-Learn library provides a framework to implement comparative experiments, also being fully integrated with the Scikit-Learn ecosystem.

4. Results and Discussion

This section presents the results and analyses of oversamplers' comparisons on the LUCAS dataset. The classification results are shown for all combinations of oversamplers and classifiers used in the experiment. The next subsection covers their interpretation in detail.

4.1. Results

For each combination of classifier and metric, a cross-validation score for all oversamplers is provided in Table 4. The highest score for each row is highlighted:

Table 4. Cross-validation scores of oversamplers.

Classifier	Metric	NONE	ROS	SMOTE	B-SMOTE	ADASYN	G-SMOTE
LR	Accuracy	0.574	0.499	0.495	0.532	0.480	0.506
LR	F-score	0.296	0.293	0.288	0.299	0.282	0.313
LR	G-mean	0.513	0.529	0.525	0.530	0.518	0.566
KNN	Accuracy	0.558	0.445	0.426	0.491	0.419	0.557
KNN	F-score	0.274	0.243	0.248	0.263	0.244	0.280
KNN	G-mean	0.496	0.478	0.487	0.500	0.483	0.504
DT	Accuracy	0.514	0.431	0.419	0.474	0.417	0.479
DT	F-score	0.243	0.243	0.250	0.272	0.250	0.267
DT	G-mean	0.488	0.483	0.490	0.508	0.492	0.519
GBC	Accuracy	0.584	0.560	0.560	0.566	0.551	0.574
GBC	F-score	0.313	0.310	0.313	0.315	0.306	0.329
GBC	G-mean	0.532	0.537	0.540	0.545	0.537	0.559
RF	Accuracy	0.587	0.576	0.557	0.571	0.552	0.579
RF	F-score	0.306	0.313	0.317	0.315	0.314	0.341
RF	G-mean	0.528	0.542	0.545	0.550	0.542	0.572

A ranking score was assigned to each oversampling method, with the best and worst performing methods receiving scores of 1 or 6, respectively. Table 5 presents the ranking scores per classifier and evaluation metric. The highest ranking for each row is highlighted:

Table 5. Ranking of oversamplers.

Classifier	Metric	NONE	ROS	SMOTE	B-SMOTE	ADASYN	G-SMOTE
LR	Accuracy	1	4	5	2	6	3
LR	F-score	3	4	5	2	6	1
LR	G-mean	6	3	4	2	5	1
KNN	Accuracy	1	4	5	3	6	2
KNN	F-score	2	6	4	3	5	1
KNN	G-mean	3	6	4	2	5	1
DT	Accuracy	1	4	5	3	6	2
DT	F-score	5	6	4	1	3	2
DT	G-mean	5	6	4	2	3	1
GBC	Accuracy	1	4	5	3	6	2
GBC	F-score	3	5	4	2	6	1
GBC	G-mean	6	5	3	2	4	1
RF	Accuracy	1	3	5	4	6	2
RF	F-score	6	5	2	3	4	1
RF	G-mean	6	5	3	2	4	1

The percentage difference between G-SMOTE and NONE, ROS, and SMOTE, respectively, for every combination of metric and classifier, was calculated from the following formula:

$$\text{Percentage Difference} = 100 \times \frac{\text{Score}(G - \text{SMOTE}) - \text{Score}(\text{Oversampler})}{\text{Score}(\text{Oversampler})}$$

For each combination of an oversampler, classifier, and metric, a positive (negative) value of the above formula indicates the G-SMOTE's relative performance gain (loss) compared to the oversampler. Table 6 presents the results of the above calculation:

Table 6. Percentage difference between G-SMOTE and other popular methods.

Classifier	Metric	NONE	ROS	SMOTE
LR	Accuracy	−12.0	1.3	2.1
LR	F-score	5.7	6.8	8.5
LR	G-mean	10.2	7.0	7.8
KNN	Accuracy	−0.1	25.0	31.0
KNN	F-score	2.0	15.2	13.0
KNN	G-mean	1.5	5.5	3.5
DT	Accuracy	−6.9	11.0	14.2
DT	F-score	10.0	10.2	7.0
DT	G-mean	6.4	7.5	6.0
GBC	Accuracy	−1.7	2.4	2.5
GBC	F-score	5.3	6.3	5.3
GBC	G-mean	5.1	4.1	3.5
RF	Accuracy	−1.3	0.5	4.0
RF	F-score	11.6	8.9	7.6
RF	G-mean	8.4	5.6	5.0

Wilcoxon signed-rank test was used as an alternative to the paired Student's *t*-test when the distribution of the differences between the two samples cannot be assumed to be normally distributed. In our case, it was applied to test the null hypothesis that the pairwise difference between G-SMOTE's scores and the scores of the remaining oversampling methods follows a symmetric distribution around zero; i.e. G-SMOTE performs similarly to them. The values for the *Accuracy* metric are excluded in the

NONE case, while for the remaining oversampling methods all metrics are used. This choice will be justified in the next section. Table 7 presents the p -values for the Wilcoxon tests:

Table 7. Wilcoxon test.

Oversampler	p -Value	Significance
NONE	5.10×10^{-3}	True
ROS	6.50×10^{-4}	True
SMOTE	6.50×10^{-4}	True
B-SMOTE	9.00×10^{-3}	True
ADASYN	6.50×10^{-4}	True

4.2. Discussion

From Table 4, we can observe that G-SMOTE outperforms all other oversampling methods for both F -score and G -mean metrics on all classifiers. The absolute best results are achieved when G-SMOTE is combined with LR and RF. It is vital to notice that the *Accuracy* scores show the well-known bias towards the majority class, as discussed in Section 3.4. In a multiclass classification problem with an imbalanced dataset, where the prediction of all the classes are of equal importance as in many remote sensing applications, *Accuracy* should be of secondary importance compared to more robust metrics, such as F -score and G -mean. Nevertheless, even for the *Accuracy* metric, G-SMOTE shows the best performance among the oversamplers.

In Table 5, the rankings of the oversamplers are presented and show the superiority of G-SMOTE. Although ROS and SMOTE are the most popular oversampling methods in remote sensing applications, it is clear from the tables that they produce suboptimal results. Table 6 directly compares the performance of G-SMOTE with ROS and SMOTE, including also NONE as a baseline method.

Table 7, provides a statistical confirmation of the previous conclusions. Using the Wilcoxon signed-rank test, the null hypothesis that the pairwise difference of scores between G-SMOTE and any of the remaining oversampling methods follows a symmetric distribution around zero is rejected at a significance level of $\alpha = 0.01$.

This study is the first to present a systematic comparison of oversampling algorithms in remote sensing. However, several previous studies reported results consistent with our findings. Reference [25] reported an increase in F -score and G -mean when oversampling was applied, while *Accuracy* did not improve. Similarly, results obtained in [5] demonstrated increased classification performance when using SMOTE. According to our experiment, performance can be further increased by using G-SMOTE. A number of other studies [21,23] did not use specific imbalanced metrics; therefore, they cannot be directly compared to our results.

5. Conclusions

In this paper we applied G-SMOTE, a novel oversampling algorithm, on a LULC classification problem, using a highly imbalanced, multiclass dataset (LUCAS). G-SMOTE's performance was evaluated and compared with other oversampling methods. More specifically, ROS, SMOTE, B-SMOTE, and ADASYN were the selected oversamplers, while LR, KNN, DT, GBC, and RF were used as classifiers.

The experimental results show that using G-SMOTE can significantly improve the classification performance, resulting in higher values of F -score and G -mean. Therefore, readers should consider using G-SMOTE when accurately predicting the minority classes is of equal or higher importance compared to the accurate prediction of the majority class. Examples of the above case include the detection of land cover change and rare land cover type classification.

G-SMOTE can be a useful tool for remote sensing researchers and practitioners, as it systematically outperforms the currently widely used oversamplers. G-SMOTE is easily accessible to the users through an open source implementation.

Author Contributions: Conceptualization, F.B.; methodology, G.D.; software, G.D.; validation, F.B., G.D.; formal analysis, J.F. and M.K.; writing—original draft preparation, M.K., J.F.; writing—review and editing, F.B., G.D., J.F., and M.K.; supervision, F.B.; funding acquisition, F.B.

Funding: This research was funded by “Fundação para a Ciência e Tecnologia” (Portugal), grant numbers PCIF/SSI/0102/2017 and DSAIPA/AI/0100/2018—IPSTERS.

Acknowledgments: The authors would like to thank Direção Geral do Território (DGT) for supporting the data used in this study.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

OS	oversampling
CV	cross-validation
LULC	land use/land cover
LUCAS	land use/cover area statistical survey
SMOTE	synthetic Minority over-sampling Technique
ADASYN	adaptive synthetic sampling Technique
G-SMOTE	geometric synthetic minority over-sampling technique
B-SMOTE	borderline synthetic minority over-sampling technique
ROS	random oversampling
NONE	no oversampling
LR	logistic regression
KNN	k-nearest neighbors
DT	decision trees
GBC	gradient boosting classifier
RF	random forest

References

1. Mellor, A.; Boukir, S.; Haywood, A.; Jones, S. Exploring issues of training data imbalance and mislabelling on random forest performance for large area land cover classification using the ensemble margin. *ISPRS J. Photogramm. Remote Sens.* **2015**, *105*, 155–168. [[CrossRef](#)]
2. Khatami, R.; Mountrakis, G.; Stehman, S.V. A meta-analysis of remote sensing research on supervised pixel-based land-cover image classification processes: General guidelines for practitioners and future research. *Remote Sens. Environ.* **2016**, *177*, 89–100. [[CrossRef](#)]
3. Tewkesbury, A.P.; Comber, A.J.; Tate, N.J.; Lamb, A.; Fisher, P.F. A critical synthesis of remotely sensed optical image change detection techniques. *Remote Sens. Environ.* **2015**, *160*, 1–14. [[CrossRef](#)]
4. Rajan, S.; Ghosh, J.; Crawford, M. An Active Learning Approach to Hyperspectral Data Classification. *IEEE Trans. Geosci. Remote Sens.* **2008**, *46*, 1231–1242. [[CrossRef](#)]
5. Feng, W.; Huang, W.; Bao, W. Imbalanced Hyperspectral Image Classification With an Adaptive Ensemble Method Based on SMOTE and Rotation Forest With Differentiated Sampling Rates. *IEEE Geosci. Remote Sens. Lett.* **2019**, *1*–5. [[CrossRef](#)]
6. Eurostat. *LUCAS 2015 (Land Use/Cover Area Frame Survey)*; Technical Reference Document C1, Instructions for Surveyors; Eurostat: Strasbourg, France, 2015.
7. Pflugmacher, D.; Rabe, A.; Peters, M.; Hostert, P. Mapping pan-European land cover using Landsat spectral-temporal metrics and the European LUCAS survey. *Remote Sens. Environ.* **2019**, *221*, 583–595. [[CrossRef](#)]
8. Mack, B.; Leinenkugel, P.; Kuenzer, C.; Dech, S. A semi-automated approach for the generation of a new land use and land cover product for Germany based on Landsat time-series and Lucas in-situ data. *Remote Sens. Lett.* **2017**, *8*, 244–253. [[CrossRef](#)]

9. Chawla, N.V.; Japkowicz, N.; Kotcz, A. Editorial: Special issue on learning from imbalanced data sets. *ACM SIGKDD Explor. Newsl.* **2004**, *6*, 1. [[CrossRef](#)]
10. Abdi, L.; Hashemi, S. To Combat Multi-Class Imbalanced Problems by Means of Over-Sampling Techniques. *IEEE Trans. Knowl. Data Eng.* **2016**, *28*, 238–251. [[CrossRef](#)]
11. García, S.; Zhang, Z.L.; Altalhi, A.; Alshomrani, S.; Herrera, F. Dynamic ensemble selection for multi-class imbalanced datasets. *Inf. Sci.* **2018**, *445–446*, 22–37. [[CrossRef](#)]
12. Sáez, J.A.; Krawczyk, B.; Woźniak, M. Analyzing the oversampling of different classes and types of examples in multi-class imbalanced datasets. *Pattern Recognit.* **2016**, *57*, 164–178. [[CrossRef](#)]
13. Fernández, A.; López, V.; Galar, M.; del Jesus, M.J.; Herrera, F. Analysing the classification of imbalanced data-sets with multiple classes: Binarization techniques and ad-hoc approaches. *Knowl.-Based Syst.* **2013**, *42*, 97–110. [[CrossRef](#)]
14. Eurostat. *LUCAS 2015 (Land Use/Cover Area Frame Survey)*; Technical Reference Document C3 Classification (Land cover and Land Use); Eurostat: Strasbourg, France, 2015.
15. Douzas, G.; Bacao, F. Geometric SMOTE a geometrically enhanced drop-in replacement for SMOTE. *Inf. Sci.* **2019**, *501*, 118–135. [[CrossRef](#)]
16. Chawla, N.V.; Bowyer, K.W.; Hall, L.O.; Kegelmeyer, W.P. SMOTE: Synthetic Minority Over-sampling Technique. *J. Artif. Intell. Res.* **2002**, *16*, 321–357. [[CrossRef](#)]
17. Han, H.; Wang, W.Y.; Mao, B.H. Borderline-SMOTE: A New Over-Sampling Method in Imbalanced Data Sets Learning. In Proceedings of the International Conference on Intelligent Computing, Hefei, China, 23–26 August 2005; Springer: Berlin/Heidelberg, Germany, 2005; pp. 878–887. [[CrossRef](#)]
18. He, H.; Bai, Y.; Garcia, E.A.; Li, S. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In Proceedings of the 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence), Hong Kong, China, 1–8 June 2008; pp. 1322–1328. [[CrossRef](#)]
19. Azadbakht, M.; Fraser, C.; Khoshelham, K. Improved urban scene classification using full-waveform LiDAR. *Photogramm. Eng. Remote Sens.* **2016**, *82*, 973–980. [[CrossRef](#)]
20. Krawczyk, B. Learning from imbalanced data: Open challenges and future directions. *Prog. Artif. Intell.* **2016**, *5*, 221–232. [[CrossRef](#)]
21. Maxwell, A.E.; Warner, T.A.; Fang, F. Implementation of machine-learning classification in remote sensing: An applied review. *Int. J. Remote Sens.* **2018**, *39*, 2784–2817. [[CrossRef](#)]
22. Feng, W.; Huang, W.; Ye, H.; Zhao, L. Synthetic Minority Over-Sampling Technique Based Rotation Forest for the Classification of Unbalanced Hyperspectral Data. In Proceedings of the IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium, Valencia, Spain, 22–27 July 2018; pp. 2651–2654.
23. Cenggoro, T.W.; Isa, S.M.; Kusuma, G.P.; Pardamean, B. Classification of imbalanced land-use/land-cover data using variational semi-supervised learning. In Proceedings of the 2017 International Conference on Innovative and Creative Information Technology: Computational Intelligence and IoT, ICITech 2017, Salatiga, Indonesia, 2–4 November 2017; pp. 1–6. [[CrossRef](#)]
24. Johnson, B.A.; Iizuka, K. Integrating OpenStreetMap crowdsourced data and Landsat time-series imagery for rapid land use/land cover (LULC) mapping: Case study of the Laguna de Bay area of the Philippines. *Appl. Geogr.* **2016**, *67*, 140–149. [[CrossRef](#)]
25. Bogner, C.; Seo, B.; Rohner, D.; Reineking, B. Classification of rare land cover types: Distinguishing annual and perennial crops in an agricultural catchment in South Korea. *PLoS ONE* **2018**, *13*, e0190476. [[CrossRef](#)]
26. Panda, A.; Singh, A.; Kumar, K.; Kumar, A.; Uddeshya; Swetapadma, A. Land Cover Prediction from Satellite Imagery Using Machine Learning Techniques. In Proceedings of the 2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT), Coimbatore, India, 20–21 April 2018; pp. 1403–1407.
27. Douzas, G.; Bacao, F. Self-Organizing Map Oversampling (SOMO) for imbalanced data set learning. *Expert Syst. Appl.* **2017**, *82*, 40–52. [[CrossRef](#)]
28. Nguyen, H.M.; Cooper, E.W.; Kamei, K. Borderline over-sampling for imbalanced data classification. In Proceedings of the 5th International Workshop on Computational Intelligence & Applications (IWCIA2009), Hiroshima, Japan, 10–12 November 2009; Volume 2009, pp. 24–29.
29. Ramentol, E.; Caballero, Y.; Bello, R.; Herrera, F. SMOTE-RSB*: A hybrid preprocessing approach based on oversampling and undersampling for high imbalanced data-sets using SMOTE and rough sets theory. *Knowl. Inf. Syst.* **2012**, *33*, 245–265. [[CrossRef](#)]

30. Liu, C.; Frazier, P.; Kumar, L. Comparative assessment of the measures of thematic classification accuracy. *Remote Sens. Environ.* **2007**, *107*, 606–616. [[CrossRef](#)]
31. Olofsson, P.; Foody, G.M.; Stehman, S.V.; Woodcock, C.E. Making better use of accuracy data in land change studies: Estimating accuracy and area and quantifying uncertainty using stratified estimation. *Remote Sens. Environ.* **2013**, *129*, 122–131. [[CrossRef](#)]
32. He, H.; Garcia, E.A. Learning from Imbalanced Data. *IEEE Trans. Knowl. Data Eng.* **2009**, *21*, 1263–1284. [[CrossRef](#)]
33. McCullagh, P.; Nelder, J. *Generalized Linear Models*; Chapman and Hall: London, UK, 1989; p. 532. [[CrossRef](#)]
34. Cover, T.; Hart, P. Nearest neighbor pattern classification. *IEEE Trans. Inf. Theory* **1967**, *13*, 21–27. [[CrossRef](#)]
35. Salzberg, S.L. C4.5: Programs for Machine Learning by J. Ross Quinlan. Morgan Kaufmann Publishers, Inc., 1993. *Mach. Learn.* **1994**, *16*, 235–240. [[CrossRef](#)]
36. Friedman, J.H. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* **2001**, *29*, 1189–1232. [[CrossRef](#)]
37. Liaw, A.; Wiener, M. Classification and regression by randomForest. *R News* **2002**, *2*, 18–22.
38. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
39. Lemaître, G.; Nogueira, F.; Aridas, C.K. Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning. *J. Mach. Learn. Res.* **2017**, *18*, 1–5.



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).