Article

EvoSplit: An Evolutionary Approach to Split a Multi-Label Data Set into Disjoint Subsets

Francisco Florez-Revuelta

Department of Computing Technology, University of Alicante, P.O. Box 99, 03080 Alicante, Spain; francisco.florez@ua.es

Abstract: This paper presents a new evolutionary approach, EvoSplit, for the distribution of multi-label data sets into disjoint subsets for supervised machine learning. Currently, data set providers either divide a data set randomly or using iterative stratification, a method that aims to maintain the label (or label pair) distribution of the original data set into the different subsets. Following the same aim, this paper first introduces a single-objective evolutionary approach that tries to obtain a split that maximizes the similarity between those distributions independently. Second, a new multi-objective evolutionary algorithm is presented to maximize the similarity considering simultaneously both distributions (labels and label pairs). Both approaches are validated using well-known multi-label data sets as well as large image data sets currently used in computer vision and machine learning applications. EvoSplit improves the splitting of a data set in comparison to the iterative stratification following different measures: Label Distribution, Label Pair Distribution, Examples Distribution, folds and fold-label pairs with zero positive examples.

Citation: Florez-Revuelta, F. EvoSplit: Evolutionary Approach to Split a Multi-Label Data Set into Disjoint Subsets. *Appl. Sci.* **2021**, *11*, 2823. <https://doi.org/10.3390/app11062823>

Academic Editor: Sławomir Nowaczyk

Received: 3 March 2021

Accepted: 18 March 2021

Published: 22 March 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2021 by the author. 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 (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: multi-label data sets; supervised learning; machine learning; evolutionary computation; big data applications

1. Introduction

Supervised learning is the machine learning task of learning a function that maps an input to an output based on example input-output pairs [1]. An inducer receives a set of labeled examples as training data and makes predictions for unseen inputs [2,3]. Traditionally, each example is associated with a single label. However, in some problems an example might be associated with multiple labels. Multi-label machine learning has received significant attention in fields such as text categorization [4], image classification [5], health risk prediction [6], or electricity load monitoring [7], among others. In computer vision in particular, there are more and more applications and available data sets that involve multi-label learning [8–10].

Formally [11], suppose $X = \mathbb{R}^d$ (or $\mathbb{Z}^d$) denotes the d -dimension instance space, and $Y = \{y_1, y_2, \dots, y_q\}$ denotes the label space with q possible class labels. The task of multi-label learning is to learn a function $h : X \rightarrow 2^Y$ from the multi-label training set $D = \{(x_i, Y_i) | 1 \leq i \leq m\}$. For each multi-label example (x_i, Y_i) , $x_i \in X$ is a d -dimensional feature vector $(x_{i1}, x_{i2}, \dots, x_{id})$ and $Y_i \subseteq Y$ is the set of labels associated with x_i . For any unseen instance $x \in X$, the multi-label classifier $C(\cdot)$ predicts $C(x) \subseteq Y$ as the set of proper labels for x .

In supervised learning, experiments typically involve a first step of distributing the examples of a data set into two or more disjoint subsets [12]. If a large amount of training data is available, the holdout method [3] is used to distribute the examples into two mutually exclusive subsets called training set and test set, or holdout set. Sometimes the training set is also divided into two disjoint subsets to create a validation set. When training data is limited, k -fold cross-validation is used, which splits the data set into k disjoint subsets of approximately equal size.

In single-label data sets, those disjoint subsets are built by distributing equally and randomly the examples in the original data set belonging to the different class labels. However, splitting a multi-label data set is not straightforward, as an over-represented class in one subset will be an under-represented in the other/s [3].

Furthermore [12], random distribution of multi-label training examples into subsets suffers from the following practical problem: it can lead to test subsets lacking even just one example of a rare label, which in turn causes calculation problems for a number of multi-label evaluation measures. The typical way these problems get by-passed in the literature is through complete removal of rare labels. This, however, implies that the performance of the learning systems on rare labels is unimportant, which is seldom true.

As mentioned by [13], multi-label classification usually follows predetermined train/test splits set by data set providers, without the analysis in terms of how well the examples are distributed into those train/test splits. Therefore, a method called stratification [3] or stratified sampling [12] was developed, in which a data set is split so that the proportion of examples of each class label in each subset is approximately equal to that in the complete data set. Stratification improves upon standard cross-validation both in terms of bias and variance, when compared to regular cross-validation [3].

The data used for learning a classifier is often imbalanced [14,15]. Thus, the class labels assigned to each instance are not equally represented. Traditionally, imbalanced classification has been faced through techniques such as resampling, cost-sensitive learning, and algorithmic-specific adaptations [16,17]. In deep learning, data augmentation is a technique that has been successful in order to address imbalanced data sets [18,19].

Different measures have been proposed to estimate the degree of multi-labelledness and the imbalance level of a data set [11,20,21]. The *label cardinality* $Card$ indicates the average number of labels per example (Equation (1)). This measure can be normalized by the number q of possible labels to obtain the *label density* $Dens$ (Equation (2)). The *label diversity* Div is the number of distinct label combinations that appear in the data set (Equation (3)), which can also be normalized by the number of examples to indicate the proportion of distinct label sets (Equation (4)). The *Theoretical Complexity Score* TCS (Equation (5)) integrates the number of input features, the number of labels, and distinct label combinations into a single measure.

$$Card(D) = \frac{\sum_{i=1}^m |Y_i|}{m} \quad (1)$$

$$Dens(D) = \frac{Card(D)}{q} \quad (2)$$

$$Div(D) = |\{Y^* | \exists x : (x, Y^*) \in D\}| \quad (3)$$

$$PDiv(D) = \frac{Div(D)}{m} \quad (4)$$

$$TCS(D) = \log(m \times q \times Div(D)) \quad (5)$$

As mentioned in [15], in binary classification the imbalance level is measured taking into account only two classes: the majority class and the minority class. In multi-label data sets the presence of the different labels can vary considerably. The *average Imbalance Ratio* $avgIR$ is the average of the imbalance ratios ($IRLbl$) between the majority label and each label $Y_i \subseteq Y$ (Equation (6)). $IRLbl$ is equal to 1 for the most frequent label and greater for

the other labels. Therefore, a larger value of the *average Imbalance Ratio* represents a higher imbalance level in the data set.

$$\begin{aligned} avgIR(D) &= \frac{\sum_{i=1}^q IRLbl(Y_i)}{q}, \\ IRLbl(y) &= \frac{\arg \max_{i=1}^q \left(\sum_{j=1}^m h(Y_i, Y_j) \right)}{\sum_{i=1}^m h(y, Y_i)}, \\ h(y, Y_k) &= \begin{cases} 1, & y = Y_k \\ 0, & y \neq Y_k \end{cases} \end{aligned} \quad (6)$$

The *SCUMBLE* measure (Equation (7)) aims to quantify the imbalance variance among the labels present in each data sample. This measure allows the estimation of the level of co-occurrence between minority and majority labels, i.e., if minority labels appear in their own or jointly with majority ones.

$$\begin{aligned} SCUMBLE(D) &= \frac{\sum_{i=1}^m SCUMBLE_i}{m}, \\ SCUMBLE_i &= 1 - \frac{\left(\prod_{j=1}^q IRLbl(Y_j) \right)^{\frac{1}{q}}}{IRLbl(Y_i)} \end{aligned} \quad (7)$$

Table 1 presents these measures for well-known multi-label data sets, and recent large multi-label image data sets that are used in machine learning and computer vision applications. The complexity of these later data sets in terms of size, number of labels, cardinality, and diversity, is much higher than *traditional* multi-label data sets. Some of these data sets, e.g., Microsoft COCO, are labeled not only with the different classes that appear in one example but with the exact number of appearances of each class. This is the reason why the frequency of the label appearing the most in the Microsoft COCO data set is higher than one, as it appears several times, on average, per example (Figure 1).

Figure 1. Image from the Microsoft COCO data set in which there are several instances of multiple objects/class labels, e.g., sheep, people. (adapted from <https://cocodataset.org/#detection-2020> (accessed on 21 March 2021)).

The remainder of this paper is organized as follows: Next, in Section 2 a review of available methods for the stratification of multi-label data sets is presented; Section 3 introduces and evaluates an evolutionary approach to obtain a stratified sampling of a data set; Section 4 proposes a multi-objective evolutionary algorithm to obtain an improved stratification; Section 5 evaluates the effect that this new splitting algorithm has on the classification metrics; and Section 6 validates this latest evolutionary approach with large image data sets currently used in computer vision and machine learning applications, and compares the results with the official splits usually employed in the literature. Finally, Section 7 discusses the methods proposed in the paper and presents some future work.

Table 1. Imbalance measures of different multi-label data sets: size of data set (m), number of labels (q), max number of labels in an example of the data set (Max Labels), maximum frequency of a label in the data set (Max Frequency), label cardinality (Card), label density (Dens), label diversity (Div), normalized label diversity (PDiv), theoretical complexity score (TCS), average imbalance ratio per label (avgIR), and SCUMBLE. Well-known multi-label data sets (top) have smaller size and lower complexity than data sets currently used in computer vision applications (bottom). Data sets are ordered by theoretical complexity score. These data sets will be used in the validation of the different methods presented in this paper.

	m	q	Max Labels	Max Frequency	Card	Dens	Div	PDiv	TCS	avgIR	SCUMBLE
emotions [22]	593	6	3	0.45	1.87	0.311	27	0.05	9.7×10^4	1.48	0.01
scene [23]	2407	6	3	0.22	1.07	0.179	15	0.01	2.2×10^5	1.25	0.00
genbase [24]	662	27	6	0.26	1.25	0.046	32	0.05	5.7×10^5	37.31	0.03
medical [25]	978	45	3	0.27	1.25	0.028	94	0.10	4.1×10^6	89.50	0.05
yeast [26]	2417	14	11	0.75	4.24	0.303	198	0.08	6.7×10^6	7.20	0.10
enron [27]	1702	53	12	0.54	3.38	0.064	753	0.44	6.8×10^7	73.95	0.30
rcv1subset4 [28]	6000	101	11	0.25	2.48	0.025	816	0.14	4.9×10^8	89.37	0.22
rcv1subset3 [28]	6000	101	12	0.24	2.61	0.026	939	0.16	5.7×10^8	68.33	0.21
rcv1subset5 [28]	6000	101	13	0.25	2.64	0.026	946	0.16	5.7×10^8	69.68	0.24
rcv1subset2 [28]	6000	101	12	0.24	2.63	0.026	954	0.16	5.8×10^8	45.51	0.21
rcv1subset1 [28]	6000	101	13	0.23	2.88	0.029	1028	0.17	6.2×10^8	54.49	0.22
tmc2007_500 [29]	28,596	22	10	0.59	2.22	0.101	1172	0.04	7.4×10^8	17.13	0.19
bibtex [30]	7395	159	28	0.14	2.40	0.015	2856	0.39	3.4×10^9	12.50	0.09
Corel5k [31]	5000	374	5	0.22	3.52	0.009	3175	0.64	5.9×10^9	189.57	0.39
COCO [8]	123,267	133	98	2.22	11.25	0.085	100,676	0.82	1.7×10^{12}	75.33	0.40
Imagenet [9,32]	10,756,941	10,592	17	0.91	8.70	0.001	10,321	0.001	1.2×10^{15}	52,059.67	0.95
OpenImages [9,33]	6,941,550	1345	91	0.50	8.79	0.007	885,489	0.13	8.3×10^{15}	3015.02	0.42

2. Related Works

The first approach to apply stratification to multi-label data sets was proposed by Sechidis et al. [12]. They proposed a greedy algorithm, Iterative Stratification (IS), that assigns iteratively examples from the original data set to each subset. The algorithm starts by taking the label with fewer examples in the data set. Those examples are, then, iteratively distributed among the different subsets based on the expected proportion of that label in each subset, i.e., an example is assigned to the subset that deviates more from the expected proportion. This process is repeated for each label until all the examples in the data set are distributed.

Sechidis et al. compared their approach against a random splitting using different measures. Following the notation presented in [12], let's consider a data set, D , annotated with a set of labels, $L = \lambda_1, \dots, \lambda_q$, a desired number k of disjoint subsets $S_1, \dots, S_k$ of D , and a desired proportion of examples $r_1, \dots, r_k$ in each of these subsets. The desired number of examples at each subset S_j is denoted as c_j and is equal to $r_j \cdot |D|$. The subsets of D and S_j that contain positive examples of label λ_i are denoted D_i and S_j^i respectively.

The Label Distribution (*LD*) measure [12] (Equation (8)) evaluates the extent to which the distribution of positive and negative examples of each label in each subset follows the distribution of that label in the whole data set. For each label λ_i , the measure computes the absolute difference between the ratio of positive to negative examples in each subset S_j with the ratio of positive to negative examples in the whole data set D , and then averages the results across all labels. This measure is calculated as:

$$LD = \frac{1}{q} \sum_{i=1}^q \left(\frac{1}{k} \sum_{j=1}^k \left| \frac{|S_j^i|}{|S_j| - |S_j^i|} - \frac{|D^i|}{|D| - |D^i|} \right| \right) \quad (8)$$

The Examples Distribution (*ED*) measure (Equation (9)) evaluates the extent to which the final number of examples in each subset S_j deviates from the desired/expected number of examples in that subset.

$$ED = \frac{1}{k} \sum_{j=1}^k ||S_j| - c_j| \quad (9)$$

Other measures are the number of folds that contain at least one label with zero positive examples (*FZ*), and the number of fold-label pairs with zero positive examples (*FLZ*).

Using these measures, Sechidis et al. demonstrated that their iterative approach maintains better the ratio of positive to negative examples of each label in each subset (*LD*) and produces the smallest number of folds (*FZ*) and fold-label pairs (*FLZ*) with zero positive examples. However, their algorithm does not consider the desired number of examples in each subset as a hard constraint, getting worse results in terms of Examples Distribution (*ED*). They also mentioned that their approach might get worse results for multi-label classification methods that consider pairs of labels, e.g., Calibrated Label Ranking [34], as their stratification method only considers the distribution of single labels.

Similarly to the measures presented in Section 1, measures could be defined not only for single labels appearing in a data set but also to higher order relations between them, i.e., simultaneous appearance of labels (label pairs, triplets...), such as $Card^k$, $Dens^k$, Div^k , and $PDiv^k$. For instance, $Card^2(D)$ would indicate the average number of label pairs per example. Table 2 shows these measures for order 2 for the data sets previously analyzed.

Given the limitation mentioned by Sechidis et al. regarding pairs of labels, Szymański and Kajdanowicz [13] extended the Iterative Stratification approach to take into account second-order relationships between labels, i.e. label pairs, and not just single labels into account when performing stratification. The proposed algorithm, Second Order Iterative Stratification (SOIS), behaves similarly to Sechidis et al.' stratification method but considering label pairs instead of single labels. The algorithm intends to maintain the same proportion of label pairs in each subset than in the original data set, considering at each point the label pair with fewer examples, distributing them among the subsets and repeating this process iteratively until no more label pairs are available. Finally, if there are examples with no label pairs these are assigned to the different subsets to comply with their expected size.

Table 2. Measures of label pair imbalance of different multi-label data sets: label pair cardinality (Card²), label pair density (Dens²), maximum frequency of a label pair (Max Frequency²), label pair diversity (Div²) and normalized label pair diversity (PDiv²).

	Card ²	Dens ²	Max Frequency ²	Div ²	PDiv ²
emotions	1.04	0.173	0.18	14	0.02
scene	0.07	0.012	0.03	8	0.00
genbase	0.40	0.015	0.04	36	0.05
medical	0.26	0.006	0.09	63	0.06
yeast	8.09	0.578	0.74	89	0.04
enron	5.19	0.098	0.34	675	0.40
rcv1subset4	2.98	0.029	0.09	1452	0.24
rcv1subset3	3.27	0.032	0.09	1665	0.28
rcv1subset5	3.38	0.033	0.09	1680	0.28
rcv1subset2	3.34	0.033	0.09	1660	0.28
rcv1subset1	3.84	0.038	0.08	1802	0.30
tmc2007_500	1.92	0.087	0.16	217	0.01
bibtex	3.11	0.020	0.02	4173	0.56
Corel5k	4.66	0.012	0.05	4342	0.87
COCO	26.69	0.201	0.22	8130	0.07
Imagenet	36.52	0.003	0.88	80,587	0.01
OpenImages	52.79	0.039	0.45	266,470	0.04

Szymański and Kajdanowicz compared SOIS with IS and random distribution using the same measures (*LD*, *ED*, *FZ*, *FLZ*). They also included a new measure, the Label Pair Distribution (*LPD*) (Equation (10)), an extension of the *LD* measure that operates on positive and negative subsets of label pairs instead of labels. Given *E* the set of label pairs appearing in the data set, S_j^i and D^i are the sets of samples that have the *i*-th label pair from *E* assigned in subset S_j and the entire data set respectively. In most cases, SOIS obtains better results than IS.

$$LPD = \frac{1}{|E|} \sum_{i=1}^{|E|} \left(\frac{1}{k} \sum_{j=1}^k \left| \frac{|S_j^i|}{|S_j| - |S_j^i|} - \frac{|D^i|}{|D| - |D^i|} \right| \right) \quad (10)$$

3. First Approach: Single-Objective Evolutionary Algorithm

This work proposes an evolutionary algorithm (EA), EvoSplit, to obtain the distribution of a data set into disjoint subsets, considering their desired size as a hard constraint. The structure of the evolutionary algorithm follows the process presented in Algorithm 1.

Algorithm 1 EvoSplit

Input: Data set D , labels L , number of disjoint k subsets, desired number c of examples at each subset

Output: Disjoint subsets S

Initialize the population with n individuals generated randomly

Rank the population by fitness

repeat

 Generate offspring

Create c new individuals by crossover

Create m new individuals by mutation

 Evaluate the offspring

Calculate the fitness of the $c + m$ new individuals

 Generate next generation's population

Rank the $n + c + m$ individuals in the population by fitness

Select the n best individuals to generate a new population

until $generations_without_changes > gen_{max}$

3.1. Characteristics of the Algorithm

Let D be a multi-label data set, k the desired number of disjoint subsets $S_1, \dots, S_k$ of D , and $c_1, \dots, c_k$ the desired number of examples at each subset. Each individual is encoded as an integer vector of size $|D|$, in which each gene represents the subset to which each example is assigned.

Different possibilities can be used to generate new individuals by crossover and mutation. EvoSplit selects parents by ranking, recombination is performed using 1-point crossover, and a mutation is carried out by reassigning randomly 1% of the genes to a different subset. This process to generate new individuals would produce in most cases individuals that do not comply with the constraint of having c_i examples in subset S_i , $i = 1 \dots k$. Therefore, a repairing process is applied to randomly reassign examples/genes to other subsets to fully comply with the constraint.

This work considers two different fitness functions: a variant of the Label Distribution (LD) and the Label Pair Distribution (LPD), which were introduced in Section 2. The Label Distribution is appropriate for data sets in which a specific label can appear only once in an example. However, for data sets that might include in a particular example several instances of the same label, Equation (8) is not appropriate. This is the case, as it was shown earlier of well-known data sets in computer vision, as Microsoft COCO [8]. Therefore, the LD has been modified to consider also data sets with this characteristic.

Let's consider $\lambda_i^{S_j}$ and λ_i^D the number of appearances of label λ_i in subset S_j and data set D respectively, L^{S_j} and L^D the total number of labels in subset S_j and data set D respectively. The modified Label Distribution measure, which is used as fitness function in EvoSplit, is then calculated following Equation (11).

$$LD' = \frac{1}{q} \sum_{i=1}^q \left(\frac{1}{k} \sum_{j=1}^k \left| \frac{\lambda_i^{S_j}}{L^{S_j} - \lambda_i^{S_j}} - \frac{\lambda_i^D}{L^D - \lambda_i^D} \right| \right) \quad (11)$$

We could proceed similarly with the LPD measure. However, EvoSplit does not consider in its calculation the number of appearances of each label in each example, but only the co-occurrence of labels, as in the original LPD measure. In this case, a variant of the LPD would increase considerably the number of pair combinations and make difficult that different examples share the same pair.

3.2. Constraints

The application of evolutionary computation allows the introduction of constraints that all the individuals in the population must fulfill to be feasible. In Section 1, it was mentioned that the distribution of labels with few examples might lead to subsets lacking examples having that label, which can difficult validation and test of multi-label classifiers. Therefore, EvoSplit introduces an optional constraint to ensure that, if possible, all subsets contain, at least, one example of each label. For instance, if a data set has to be split into three subsets and a label only appears in three examples, each of those examples will be distributed to a different subset. The constraint is not considered for those class labels which number of examples is lower than the number of subsets. Some other constraints could also be considered, if needed.

In case of generating an individual that does not fulfill the constraint, a repairing process similar to that explained before would be applied.

3.3. Results

Next, this work presents a comparison of the performance of the proposed evolutionary approach with other alternatives to split a data set into disjoint subsets, i.e., random and stratified (SOIS) (The stratified alternative has been obtained using the algorithm provided by scikit-multilearn [35]). Similarly to the literature [12,13], this work has evaluated the different methods considering 10-fold cross-validation of well-know multi-label data sets.

For the evolutionary approaches, the parameters have been selected experimentally:

- Size of the population (n): 50
- Individuals created by crossover (c): 10
- Individuals created by mutation (m): 10
- Number of generations without changes in the best individual (gen_{max}): 25

The splitting has been carried out using both the Label Distribution and the Label Pair Distribution as fitness functions. For each of these alternatives, results have been obtained without and with the constraint presented in Section 3.2 that tries to ensure that, if possible, all folds have examples for all the labels. For each alternative, given the probabilistic behavior of evolutionary algorithms, the best result of five runs of the algorithm has been selected.

Table 3 shows the results in the case of using the Label Distribution as fitness function. In this case, the optimization algorithm tries to create subsets in which the proportion of examples of each class is close to that in the complete data set. In both cases, whether considering or not the constraint, the evolutionary approach obtains better results than any other method. Something similar happens (Table 4) when the Label Pair Distribution is employed as fitness function, i.e., the algorithm tries to approximate the proportion of label pairs. However, in most cases, when the evolutionary algorithm tries to improve the distribution of single labels (by using *LD*) fails to distribute label pairs better than the Stratification method. Something similar happens when the fitness function is the *LPD* measure and the results are measured in terms of *LD*.

It is worth mentioning that the Stratification method does not consider the desired number of samples per fold as a hard constraint. Therefore, the final sizes of the subsets might deviate from the pre-established ones, as measured by the Examples Distribution and shown in Table 5. This is not allowed for the evolutionary approaches.

Following [12], besides using the Label Distribution and the Label Pair Distribution, the result of each alternative is also measured (see Tables 6 and 7) in terms of the number of folds that contain at least one label with zero positive examples (*FZ*), and the number of fold-label pairs with zero positive examples (*FLZ*). The *FLZ* measure shows the effect of introducing the constraint of ensuring that subsets contain, if possible, one example of each label. In the constrained version of EvoSplit, the *FLZ* values are lower than with all the other methods.

Table 3. Label Distribution of different splitting algorithms. The Stratification method is considered as the baseline. In bold the results that are better than the Stratification method. In all cases, the evolutionary algorithm that uses *LD* as fitness function obtains the best results. This does not happen when the *LPD* is used as fitness function.

	Random	Stratification	Evolutionary Algorithm			
			Label Distribution		Label Pair Distribution	
			Unconstrained	Constrained	Unconstrained	Constrained
emotions	3.77×10^{-2}	6.45×10^{-3}	2.87×10^{-3}	2.85×10^{-3}	1.84×10^{-2}	1.79×10^{-2}
scene	2.68×10^{-2}	1.69×10^{-3}	1.27×10^{-3}	1.42×10^{-3}	2.46×10^{-2}	2.38×10^{-2}
genbase	1.38×10^{-2}	4.75×10^{-3}	4.54×10^{-3}	4.48×10^{-3}	1.03×10^{-2}	1.30×10^{-2}
medical	7.68×10^{-3}	2.96×10^{-3}	2.88×10^{-3}	2.79×10^{-3}	7.56×10^{-3}	7.61×10^{-3}
yeast	5.72×10^{-3}	1.47×10^{-3}	5.24×10^{-4}	5.12×10^{-4}	1.35×10^{-3}	1.29×10^{-3}
enron	2.89×10^{-3}	1.79×10^{-3}	8.90×10^{-4}	8.34×10^{-4}	1.46×10^{-3}	1.45×10^{-3}
rcv1subset4	1.55×10^{-3}	4.24×10^{-4}	3.72×10^{-4}	3.30×10^{-4}	5.87×10^{-4}	5.73×10^{-4}
rcv1subset3	1.60×10^{-3}	4.96×10^{-4}	3.42×10^{-4}	3.30×10^{-4}	5.84×10^{-4}	5.41×10^{-4}
rcv1subset5	1.55×10^{-3}	4.21×10^{-4}	3.20×10^{-4}	3.14×10^{-4}	5.14×10^{-4}	4.95×10^{-4}
rcv1subset2	1.53×10^{-3}	4.07×10^{-4}	3.39×10^{-4}	3.21×10^{-4}	5.30×10^{-4}	5.05×10^{-4}
rcv1subset1	1.67×10^{-3}	4.32×10^{-4}	3.26×10^{-4}	2.95×10^{-4}	5.04×10^{-4}	4.59×10^{-4}
tmc2007_500	1.76×10^{-3}	2.23×10^{-4}	1.74×10^{-4}	1.61×10^{-4}	8.05×10^{-4}	7.43×10^{-4}
bibtex	1.35×10^{-3}	3.62×10^{-4}	2.40×10^{-4}	2.39×10^{-4}	7.81×10^{-4}	7.94×10^{-4}
Corel5k	6.36×10^{-4}	4.39×10^{-4}	2.68×10^{-4}	2.48×10^{-4}	3.71×10^{-4}	3.44×10^{-4}

Table 4. Label Pair Distribution of different splitting algorithms. In bold the results that are better than the stratification method. In all cases, the evolutionary algorithm that uses *LPD* as fitness function obtains the best results. This does not happen when the *LD* is used as fitness function.

	Random	Stratification	Evolutionary Algorithm			
			Label Distribution		Label Pair Distribution	
			Unconstrained	Constrained	Unconstrained	Constrained
emotions	2.64×10^{-2}	6.61×10^{-3}	1.54×10^{-2}	1.93×10^{-2}	4.81×10^{-3}	4.63×10^{-3}
scene	1.08×10^{-1}	2.83×10^{-2}	7.84×10^{-2}	8.11×10^{-2}	1.97×10^{-2}	2.11×10^{-2}
genbase	1.98×10^{-2}	1.73×10^{-2}	1.71×10^{-2}	1.74×10^{-2}	1.37×10^{-2}	1.37×10^{-2}
medical	1.61×10^{-2}	1.28×10^{-2}	1.53×10^{-2}	1.46×10^{-2}	1.17×10^{-2}	1.17×10^{-2}
yeast	1.57×10^{-3}	6.45×10^{-4}	8.39×10^{-4}	8.81×10^{-4}	3.60×10^{-4}	3.35×10^{-4}
enron	6.68×10^{-4}	5.61×10^{-4}	5.66×10^{-4}	5.57×10^{-4}	4.16×10^{-4}	4.19×10^{-4}
rcv1subset4	3.35×10^{-4}	2.10×10^{-4}	2.63×10^{-4}	2.54×10^{-4}	1.87×10^{-4}	1.86×10^{-4}
rcv1subset3	2.98×10^{-4}	2.29×10^{-4}	2.40×10^{-4}	2.41×10^{-4}	1.73×10^{-4}	1.72×10^{-4}
rcv1subset5	2.85×10^{-4}	1.96×10^{-4}	2.34×10^{-4}	2.30×10^{-4}	1.69×10^{-4}	1.69×10^{-4}
rcv1subset2	2.92×10^{-4}	1.85×10^{-4}	2.40×10^{-4}	2.38×10^{-4}	1.73×10^{-4}	1.70×10^{-4}
rcv1subset1	2.70×10^{-4}	1.78×10^{-4}	2.11×10^{-4}	2.05×10^{-4}	1.51×10^{-4}	1.52×10^{-4}
tmc2007_500	4.69×10^{-4}	1.94×10^{-4}	3.82×10^{-4}	3.71×10^{-4}	1.30×10^{-4}	1.18×10^{-4}
bibtex	1.96×10^{-4}	1.68×10^{-4}	1.82×10^{-4}	1.80×10^{-4}	1.44×10^{-4}	1.45×10^{-4}
Corel5k	1.72×10^{-4}	1.43×10^{-4}	1.59×10^{-4}	1.59×10^{-4}	1.29×10^{-4}	1.29×10^{-4}

Table 5. Examples Distribution of different splitting algorithms. Only the stratification method deviates from zero.

	Random	Stratification	All EAs
emotions	0	1.2	0
scene	0	1.4	0
genbase	0	0.8	0
medical	0	1.2	0
yeast	0	4.2	0
enron	0	3.8	0
rcv1subset4	0	3.8	0
rcv1subset3	0	4.8	0
rcv1subset5	0	6.6	0
rcv1subset2	0	3.6	0
rcv1subset1	0	7.2	0
tmc2007_500	0	19.8	0
bibtex	0	14	0
Corel5k	0	6.2	0

Table 6. Number of folds that contain at least one label with zero positive examples (FZ). All the splitting algorithms obtain the same results.

	All Cases
emotions	0
scene	0
genbase	10
medical	10
yeast	0
enron	10
rcv1subset4	10
rcv1subset3	10
rcv1subset5	10
rcv1subset2	10
rcv1subset1	10
tmc2007_500	0
bibtex	0
Corel5k	10

Table 7. Number of fold-label pairs with zero positive examples (*FLZ*) of different splitting algorithms. In bold the results that are better than the stratification method. Between brackets those results that are worse.

	Random	Stratification	Evolutionary Algorithm			
			Label Distribution		Label Pair Distribution	
			Unconstrained	Constrained	Unconstrained	Constrained
emotions	0	0	0	0	0	0
scene	0	0	0	0	0	0
genbase	85	76	74	75	74	73
medical	202	174	(176)	(175)	(200)	(180)
yeast	0	0	0	0	0	0
enron	87	77	59	53	62	55
rcv1subset4	106	68	(78)	63	66	62
rcv1subset3	77	57	54	40	44	39
rcv1subset5	83	50	49	44	48	47
rcv1subset2	69	35	(36)	30	34	33
rcv1subset1	74	59	56	43	50	44
tmc2007_500	0	0	0	0	0	0
bibtex	2	0	0	0	0	0
Corel5k	1157	1019	925	843	948	866

Tables 3 and 4 show that the evolutionary approach gets better results than Stratification only for the metric that is employed as fitness function. There are only some few cases in which good results are obtained for both metrics, the Label Distribution and the Label Pair Distribution simultaneously. On the other hand, the Stratification method obtains splits that behave well for both metrics, even if the best results for each metric are obtained with the evolutionary approaches. This is probably due to the process that the Stratification method follows to distribute examples to subsets: first, considering label pairs in the distribution and, later, assigning remaining examples based on single labels, i.e., taking into consideration both labels and label pairs in the splitting.

4. Second Approach: Multi-Objective Evolutionary Algorithm

Then, it seems appropriate to consider both statistical measures, the Label Distribution and the Label Pair Distribution, in the optimization algorithm to split a data set into disjoint subsets. Multi-objective optimization problems are those problems where the goal is to optimize simultaneously several objective functions. These different functions have conflicting objectives, i.e., optimizing one affects the others. Therefore, there is not a unique solution but a set of solutions. The set of solutions in which the different objective components cannot be simultaneously improved constitute a Pareto front. Each solution in the Pareto front represents a trade-off between the different objectives. Similarly to evolutionary algorithms for single objective problems, multi-objective evolutionary algorithms (MOEA) [36] are heuristic algorithms to solve problems with multiple objective functions. The three goals of an MOEA are [37]: (1) to find a set of solutions as close as possible to the Pareto front (known as convergence); (2) to maintain a diverse population that contains dissimilar individuals to promote exploration and to avoid poor performance due to premature convergence (known as diversity); and (3) to obtain a set of solutions that spreads in a more uniform way over the Pareto front (known as coverage). Several MOEAs have been proposed in the literature. This work employs the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [38].

NSGA-II has the three following features: (1) it uses an elitist principle, i.e., the elites of a population are given the opportunity to be carried to the next generation; (2) it uses an

explicit diversity preserving mechanism (Crowding distance); and (3) it emphasizes the non-dominated solutions.

Therefore, this work employs NSGA-II to distribute the data set into subsets optimizing simultaneously both the Label Distribution and the Label Pair Distribution. The algorithm will obtain a set of solutions, some of them optimizing one over the other objective and vice versa. From these set of solutions, EvoSplit selects the solution closer (using Euclidean distance) to the coordinates origin (Figure 2).

Figure 2. Final set of solutions for a run of the MOEA algorithm to split the Corel5k data set. Selected global optimum is marked with a circle.

This work has employed the implementation of NSGA-II offered by pymoo [39], a multi-objective optimization framework in Python, using the same parameters in terms of individuals, size of offspring and ending condition presented in Section 3.3.

Results

Table 8 shows the results obtained for the different measures of the splits obtained using the MOEA unconstrained approach. The Examples Distribution measure is not shown as it is always zero, as with the previous evolutionary approaches. The obtained results are, in most cases, better than those obtained with the Stratification method. In this approach, unlike the previous single-objective evolutionary alternatives, results are good in terms of both *LD* and *LPD*. The MOEA approach obtains results in terms of *LD* close to those obtained by the single-objective approach optimizing only *LD* (see Table 3), and close in terms of *LPD* when the optimization is based only in *LPD* (see Table 4). These are more balanced results than those obtained with the single-objective evolutionary approaches, i.e., a good result in one of the measures does not affect a good result in the other one. Additionally, results are also quite similar in terms of *FZ* and *FLZ*. For those data sets with *FLZ* different to zero, Table 9 shows the results obtained with the constrained alternative. For some data sets, *FLZ* is reduced without affecting considerably the *LD* and *LPD* measures.

Table 8. Evaluation of the MOEA approach without considering the constraint. In bold the results that are better than the Stratification method.

	LD	LPD	FZ	FLZ
emotions	5.28×10^{-3}	7.43×10^{-3}	0	0
scene	1.77×10^{-3}	2.18×10^{-2}	0	0
genbase	4.70×10^{-3}	1.42×10^{-2}	10	74
medical	3.39×10^{-3}	1.19×10^{-2}	10	179
yeast	6.56×10^{-4}	4.80×10^{-4}	0	0
enron	8.28×10^{-4}	4.60×10^{-4}	10	58
rcv1subset4	2.93×10^{-4}	2.04×10^{-4}	10	71
rcv1subset3	2.93×10^{-4}	1.87×10^{-4}	10	42
rcv1subset5	2.70×10^{-4}	1.81×10^{-4}	10	47
rcv1subset2	2.96×10^{-4}	1.89×10^{-4}	10	36
rcv1subset1	2.94×10^{-4}	1.68×10^{-4}	10	50
tmc2007_500	1.27×10^{-4}	1.89×10^{-4}	0	0
bibtex	2.68×10^{-4}	1.53×10^{-4}	0	0
Corel5k	2.84×10^{-4}	1.41×10^{-4}	10	907

Table 9. Evaluation of the MOEA approach considering the constraint. In bold the results that are better than the Stratification method.

	LD	LPD	FZ	FLZ
genbase	5.12×10^{-3}	1.40×10^{-2}	10	73
medical	2.96×10^{-3}	1.19×10^{-2}	10	175
enron	8.29×10^{-4}	4.63×10^{-4}	10	53
rcv1subset4	3.11×10^{-4}	2.06×10^{-4}	10	61
rcv1subset3	2.70×10^{-4}	1.85×10^{-4}	10	38
rcv1subset5	2.69×10^{-4}	1.84×10^{-4}	10	45
rcv1subset2	2.87×10^{-4}	1.89×10^{-4}	10	31
rcv1subset1	2.54×10^{-4}	1.63×10^{-4}	10	41
Corel5k	2.52×10^{-4}	1.39×10^{-4}	10	841

5. Evaluation of Classification

Next, this work evaluates the effect that the distribution of examples in a multi-label data set into subsets has on the classification metrics. Similarly to [13], the evaluation is carried out employing two standard multi-label classification algorithms: Binary Relevance (BR) and Label Powerset (LP). The results of the classification are measured using recall, precision, F_1 score and Jaccard index. These metrics are calculated using both micro-average (i.e., calculate metrics globally by counting the total true positives, false negatives and false positives) and macro-average (i.e., calculate metrics for each label, and then their unweighted mean). The variance in the classification across different folds is also analyzed, which provides information about the generalization stability of the different approaches.

Following [12,13], the results are discussed based on the average ranking of the different methods. The best method for a particular measure obtains a rank of 1, the next

one a rank of 2, and so on. Figures 3 and 4 show the evaluations for both classification algorithms, BR and LP. Evolutionary approaches obtain, in general, better classification results than the Stratification method for all the metrics (left part of the figures). Even in the case of the constrained approaches, in which it is supposed to favor better results in macro-averaged metrics maybe in prejudice of micro-averaged metrics, evolutionary approaches behave better than Stratification. Using BR the best classification metrics are obtained by the unconstrained EA using LPD as fitness function followed by the constrained MOEA. Regarding variance over folds, these two methods are also the best ones. Using LP the best method is, for almost all the metrics, the constrained MOEA followed by the unconstrained EA using LPD. In terms of variance over folds, all the evolutionary approach have globally a similar behavior, being better than Stratification.

Figure 3. Average ranks of the different methods when classification is performed using Binary Relevance. Left figure shows the ranks regarding classification metrics. Right figure shows the ranks regarding the variance over folds.

Figure 4. Average ranks of the different methods when classification is performed using Label Powerset. Left figure shows the ranks regarding classification metrics. Right figure shows the ranks regarding the variance over folds.

6. Application to Large Image Data Sets

Given the recent relevance of multi-label computer vision applications using deep learning techniques, EvoSplit has also been validated using some large multi-label data sets widely employed by the research community: Microsoft COCO, ImageNet, and Open-Images.

6.1. Microsoft COCO

The Microsoft Common Objects in Context (MS COCO) data set [8] contains 91 common object categories with 82 of them having more than 5000 labeled instances. In total, the data set has 2,500,000 labeled instances in 328,000 images. This work has used the subset considered in the COCO Panoptic Segmentation Task, which includes 123,167

images. From these, 5000 are selected for validation, and the remaining ones for training. The panoptic segmentation task involves assigning a semantic label and instance id for each pixel of an image, which requires generating dense, coherent scene segmentations. From all the data sets considered in this work, MS COCO is the one with the highest cardinality, an average of more than 11 labels per image. It is the only data set in which examples are not labeled stating if a label appears or not in an image but with the number of times that the label appears. As shown in Table 1, this makes possible that the label (class 1 = person) appearing the most in the data set (Max Frequency) does it more than once, on average, in each image.

6.2. Tencent ML-Images

The Tencent ML-Images database [9] is a multi-label image database with 18M images and 11K categories, collected from ImageNet [32] and OpenImages [33]. After a process of removal and inclusion of images, and relabeling of the data set, 10,756,941 images, covering 10,032 categories, are included from Imagenet. From these, 50,000 are randomly selected as validation set. Following a similar process, 6,902,811 training images and 38,739 validation images are selected from OpenImages, covering 1134 unique categories. Finally, these images and categories from ImageNet and OpenImages are merged to construct the Tencent ML-Images database, which includes 17,609,752 training and 88,739 validation images (50,000 from ImageNet and 38,739 from OpenImages), covering 11,166 categories.

6.3. Results

The measures of the application of the different methods to split these data sets are shown in Tables 10–12. In almost all the cases (in bold), any evolutionary approach, either single-objective or multi-objective, either constrained or unconstrained, performs better than the official or stratified splitting methods. Similarly to the results shown for traditional smaller data sets, MOEA shows the best combined results for the Label Distribution and the Label Pair Distribution in almost all the cases. For the Microsoft COCO and the OpenImages data sets the results for those measures are improved by one or more orders of magnitude with respect to the official splits, i.e., those offered by the providers of the data set.

With these data sets it is even clearer the effect of using the constrained approach, in which the goal is to include all the labels in every fold. The introduction of the constraint allows to obtain *FZ* and *FLZ* equal to zero for the OpenImages data set, while their values for the official split are 1 and 9 respectively. *FLZ* is dramatically reduced for the Imagenet data set (361 vs. 19 using MOEA).

Table 10. Evaluation of the different splitting approaches of the Microsoft COCO data set. In bold the results that are better than the official method, usually employed in the literature. The results of the constrained alternative is not presented as the unconstrained approach obtains *FZ* and *FLZ* measures equal to zero.

	Official	Random	Stratification	Evolutionary Algorithm		MOEA
				LD	LPD	
LD	2.42×10^{-4}	5.06×10^{-4}	2.54×10^{-4}	3.10×10^{-5}	2.76×10^{-4}	1.51×10^{-5}
LPD	7.24×10^{-6}	1.28×10^{-5}	5.66×10^{-6}	1.22×10^{-5}	6.48×10^{-6}	3.96×10^{-6}
ED	0	0	934	0	0	0
FZ	0	1	0	0	0	0
FLZ	0	1	0	0	0	0

Table 11. Evaluation of the different splitting approaches of the Imagenet subset in the Tencent ML-Images database. In bold the results that are better than the official method, usually employed in the literature. NOTE: The results for the Stratification method are not shown as no result was obtained after 10 days of processing. Therefore, *ED* is not shown as it only deviates from zero for that method.

	Official	Random	Evolutionary Algorithm				MOEA	
			Label Distribution		Label Pair Distribution		Unconst.	Const.
			Unconst.	Const.	Unconst.	Const.		
LD	2.04×10^{-6}	2.86×10^{-6}	1.33×10^{-6}	1.35×10^{-6}	9.17×10^{-7}	8.02×10^{-7}	1.04×10^{-6}	1.48×10^{-6}
LPD	4.21×10^{-7}	6.45×10^{-7}	3.34×10^{-7}	3.36×10^{-7}	2.23×10^{-7}	1.94×10^{-7}	2.58×10^{-7}	3.49×10^{-7}
FZ	1	1	1	1	1	1	1	1
FLZ	361	495	466	36	415	36	429	19

Table 12. Evaluation of the different splitting approaches of the OpenImages subset in the Tencent ML-Images database. In bold the results that are better than the official method, usually employed in the literature. NOTE: The results for the Stratification method are not shown as no result was obtained after 10 days of processing. Therefore, *ED* is not shown as it only deviates from zero for this method.

	Official	Random	Evolutionary Algorithm				MOEA	
			Label Distribution		Label Pair Distribution		Unconst.	Const.
			Unconst.	Const.	Unconst.	Const.		
LD	2.25×10^{-4}	9.03×10^{-6}	2.89×10^{-6}	2.45×10^{-6}	2.63×10^{-6}	2.79×10^{-6}	2.15×10^{-6}	2.59×10^{-6}
LPD	1.18×10^{-6}	1.34×10^{-7}	1.15×10^{-7}	1.18×10^{-7}	6.89×10^{-8}	7.22×10^{-8}	8.73×10^{-8}	9.06×10^{-8}
FZ	1	1	1	0	1	0	1	0
FLZ	9	27	25	0	11	0	21	0

7. Conclusions

This paper presents EvoSplit, a novel evolutionary method to split a multi-label data set into disjoint subsets. Different proposals, single-objective and multi-objective, using diverse measures as fitness function, have been proposed. A constraint has also been introduced in order to ensure that, if possible, labels are distributed among all subsets. In almost all the cases, the multi-objective proposal obtains state-of-the-art results, improving or matching the quality of the splits officially provided or obtained with iterative stratification methods. The improvement of EvoSplit over previous methods is highlighted when applied to very large data sets, as those currently used in machine learning and computer vision applications.

Moreover, the introduction of the constrained optimization decreases the chance of producing subsets with zero positive examples for one or more labels. This should have an effect on the training as there will be fewer labels for which there are no training or validation examples. A very relevant result is that EvoSplit is able to find splits that fulfill the constrain without affecting too much the distribution of labels and label pairs.

EvoSplit is able to obtain better distributions of the original data sets considering the desired size of the subsets as a hard constraint, i.e., ensuring that the Examples Distribution is equal to zero. This is not the case for the iterative stratification method. The subsets obtained with EvoSplit allow better classification results and reduce the variance in the performance across subsets/folds.

Only in the case of the Imagenet data set, the best results are not obtained by the multi-objective EA but by the single-objective EA optimizing the Label Pair Distribution measure. An explanation to this might be related to the relation in the diversity between labels and pair labels for this data set. This data set has a particular characteristic: there are

almost eight times more different pair labels than different labels in the data set (see the Diversity measure in Tables 1 and 2). For all the other data sets, the relation is close to one. Therefore, this larger proportion in the diversity of label pairs might have an influence, benefiting the optimization based only on label pairs.

In conclusion, EvoSplit supports researchers in the process of creating a data set by providing different evolutionary alternatives to split that data set by optimizing the distribution of examples into the different subsets. EvoSplit can, in the future, be extended to higher levels of relationship between labels, e.g., triplets, by implementing a many-objective evolutionary algorithm [40].

8. Availability of Splits and Code

The splits obtained with EvoSplit for the different data sets employed in this paper are freely available at <https://github.com/FranciscoFlorezRevuelta/EvoSplit> (accessed on 21 March 2021) for their use by the research community. The EvoSplit code will be also available at the same repository.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The author declares no conflict of interest. Furthermore, 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.

References

1. Russell, S.; Norvig, P. *Artificial Intelligence: A Modern Approach*; Prentice Hall: Hoboken, NJ, USA, 2002.
2. Mohri, M.; Rostamizadeh, A.; Talwalkar, A. *Foundations of Machine Learning*; MIT Press: Cambridge, MA, USA, 2018.
3. Kohavi, R. A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection. In Proceedings of the 14th International Joint Conference on Artificial Intelligence—Volume 2, Montreal, QC, Canada, 20–25 August 1995; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 1995; IJCAI'95, pp. 1137–1143.
4. Liu, J.; Chang, W.C.; Wu, Y.; Yang, Y. Deep Learning for Extreme Multi-Label Text Classification. In Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Tokyo, Japan, 7–11 August 2017; Association for Computing Machinery: New York, NY, USA, 2017; SIGIR '17, pp. 115–124. [\[CrossRef\]](#)
5. Wang, J.; Yang, Y.; Mao, J.; Huang, Z.; Huang, C.; Xu, W. CNN-RNN: A Unified Framework for Multi-Label Image Classification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016.
6. Maxwell, A.; Li, R.; Yang, B.; Weng, H.; Ou, A.; Hong, H.; Zhou, Z.; Gong, P.; Zhang, C. Deep learning architectures for multi-label classification of intelligent health risk prediction. *BMC Bioinform.* **2017**, *18*, 523. [\[CrossRef\]](#)
7. Tabatabaei, S.M.; Dick, S.; Xu, W. Toward Non-Intrusive Load Monitoring via Multi-Label Classification. *IEEE Trans. Smart Grid* **2017**, *8*, 26–40. [\[CrossRef\]](#)
8. Lin, T.Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; Zitnick, C.L. Microsoft COCO: Common Objects in Context. In *Computer Vision—ECCV 2014*; Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T., Eds.; Springer International Publishing: Cham, Switzerland, 2014; pp. 740–755.
9. Wu, B.; Chen, W.; Fan, Y.; Zhang, Y.; Hou, J.; Liu, J.; Zhang, T. Tencent ml-images: A large-scale multi-label image database for visual representation learning. *IEEE Access* **2019**, *7*, 172683–172693. [\[CrossRef\]](#)
10. Bustos, A.; Pertusa, A.; Salinas, J.M.; de la Iglesia-Vayá, M. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Med Image Anal.* **2020**, *66*, 101797. [\[CrossRef\]](#)
11. Zhang, M.; Zhou, Z. A Review on Multi-Label Learning Algorithms. *IEEE Trans. Knowl. Data Eng.* **2014**, *26*, 1819–1837. [\[CrossRef\]](#)
12. Sechidis, K.; Tsoumakas, G.; Vlahavas, I. On the stratification of multi-label data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 145–158.
13. Szymański, P.; Kajdanowicz, T. A network perspective on stratification of multi-label data. *arXiv* **2017**, arXiv:1704.08756.
14. Tahir, M.A.; Kittler, J.; Bouridane, A. Multilabel classification using heterogeneous ensemble of multi-label classifiers. *Pattern Recognit. Lett.* **2012**, *33*, 513–523. [\[CrossRef\]](#)
15. Charte, F.; Rivera, A.J.; del Jesus, M.J.; Herrera, F. Addressing imbalance in multilabel classification: Measures and random resampling algorithms. *Neurocomputing* **2015**, *163*, 3–16. [\[CrossRef\]](#)

16. López, V.; Fernández, A.; García, S.; Palade, V.; Herrera, F. An insight into classification with imbalanced data: Empirical results and current trends on using data intrinsic characteristics. *Inf. Sci.* **2013**, *250*, 113–141. [\[CrossRef\]](#)
17. Liu, B.; Tsoumakas, G. Synthetic Oversampling of Multi-label Data Based on Local Label Distribution. In *Machine Learning and Knowledge Discovery in Databases*; Brefeld, U., Fromont, E., Hotho, A., Knobbe, A., Maathuis, M., Robardet, C., Eds.; Springer International Publishing: Cham, Switzerland, 2020; pp. 180–193.
18. Shorten, C.; Khoshgoftaar, T.M. A survey on image data augmentation for deep learning. *J. Big Data* **2019**, *6*, 60. [\[CrossRef\]](#)
19. Leng, B.; Yu, K.; QIN, J. Data augmentation for unbalanced face recognition training sets. *Neurocomputing* **2017**, *235*, 10–14. [\[CrossRef\]](#)
20. Charte, F.; Rivera, A.; del Jesus, M.J.; Herrera, F. A First Approach to Deal with Imbalance in Multi-label Datasets. In *Hybrid Artificial Intelligent Systems*; Pan, J.S., Polycarpou, M.M., Woźniak, M., de Carvalho, A.C.P.L.F., Quintián, H., Corchado, E., Eds.; Springer: Berlin/Heidelberg, Germany, 2013; pp. 150–160.
21. Charte, F.; Rivera, A.J.; Charte, D.; del Jesus, M.J.; Herrera, F. Tips, guidelines and tools for managing multi-label datasets: The mldr.datasets R package and the Cometa data repository. *Neurocomputing* **2018**, *289*, 68–85. [\[CrossRef\]](#)
22. Tsoumakas, G.; Katakis, I.; Vlahavas, I. Effective and efficient multilabel classification in domains with large number of labels. In *Proceedings of the ECML/PKDD 2008 Workshop on Mining Multidimensional Data (MMD'08)*, Antwerp, Belgium, 19 September 2008; Volume 21, pp. 53–59.
23. Boutell, M.R.; Luo, J.; Shen, X.; Brown, C.M. Learning multi-label scene classification. *Pattern Recognit.* **2004**, *37*, 1757–1771. [\[CrossRef\]](#)
24. Diplaris, S.; Tsoumakas, G.; Mitkas, P.A.; Vlahavas, I. Protein Classification with Multiple Algorithms. In *Advances in Informatics*; Bozanis, P., Houstis, E.N., Eds.; Springer: Berlin/Heidelberg, Germany, 2005; pp. 448–456.
25. Pestian, J.; Brew, C.; Matykiewicz, P.; Hovermale, D.J.; Johnson, N.; Cohen, K.B.; Duch, W. A shared task involving multi-label classification of clinical free text. In *Proceedings of the Workshop on BioNLP 2007: Biological, Translational, and Clinical Language Processing*, Prague, Czech Republic, 29 June 2007; pp. 97–104.
26. Elisseeff, A.; Weston, J. A kernel method for multi-labelled classification. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, Vancouver, BC, Canada, 3–8 December 2001; pp. 681–687.
27. Read, J.; Pfahringer, B.; Holmes, G. Multi-label Classification Using Ensembles of Pruned Sets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, Pisa, Italy, 15–19 December 2008; pp. 995–1000.
28. Lewis, D.D.; Yang, Y.; Rose, T.G.; Li, F. Rcv1: A new benchmark collection for text categorization research. *J. Mach. Learn. Res.* **2004**, *5*, 361–397.
29. Srivastava, A.N.; Zane-Ulman, B. Discovering recurring anomalies in text reports regarding complex space systems. In *Proceedings of the 2005 IEEE Aerospace Conference, Big Sky, MT, USA*, 5–12 March 2005; pp. 3853–3862.
30. Katakis, I.; Tsoumakas, G.; Vlahavas, I. Multilabel text classification for automated tag suggestion. In *Proceedings of the ECML/PKDD*, Antwerp, Belgium, 15–19 September 2008; Volume 18, p. 5.
31. Duygulu, P.; Barnard, K.; de Freitas, J.F.G.; Forsyth, D.A. Object Recognition as Machine Translation: Learning a Lexicon for a Fixed Image Vocabulary. In *Computer Vision—ECCV 2002*; Heyden, A., Sparr, G., Nielsen, M., Johansen, P., Eds.; Springer: Berlin/Heidelberg, Germany, 2002; pp. 97–112.
32. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*, Miami Beach, FL, USA, 20–25 June 2009; pp. 248–255.
33. Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. The Open Images Dataset V4. *Int. J. Comput. Vis.* **2020**, *128*, 1956–1981. [\[CrossRef\]](#)
34. Fürnkranz, J.; Hüllermeier, E.; Mencía, E.L.; Brinker, K. Multilabel classification via calibrated label ranking. *Mach. Learn.* **2008**, *73*, 133–153. [\[CrossRef\]](#)
35. Szymanski, P.; Kajdanowicz, T. Scikit-Multilearn: A Scikit-Based Python Environment for Performing Multi-Label Classification. *J. Mach. Learn. Res.* **2019**, *20*, 209–230.
36. Coello, C.A.C.; Lamont, G.B.; Van Veldhuizen, D.A. *Evolutionary Algorithms for Solving Multi-Objective Problems*; Springer: New York, NY, USA, 2007; Volume 5.
37. Trivedi, A.; Srinivasan, D.; Sanyal, K.; Ghosh, A. A Survey of Multiobjective Evolutionary Algorithms Based on Decomposition. *IEEE Trans. Evol. Comput.* **2017**, *21*, 440–462. [\[CrossRef\]](#)
38. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. [\[CrossRef\]](#)
39. Blank, J.; Deb, K. Pymoo: Multi-Objective Optimization in Python. *IEEE Access* **2020**, *8*, 89497–89509. [\[CrossRef\]](#)
40. Li, B.; Li, J.; Tang, K.; Yao, X. Many-objective evolutionary algorithms: A survey. *ACM Comput. Surv. (CSUR)* **2015**, *48*, 1–35. [\[CrossRef\]](#)