

# On the Impact of Dataset Complexity and Sampling Strategy in Multilabel Classifiers Performance

Francisco Charte, Antonio Rivera, María José del Jesus, Francisco Herrera  
Dep. of Computer Science and Artificial Intelligence, University of Granada, Granada, Spain.  
Dep. of Computer Science, University of Jaén, Jaén, Spain

## Abstract

Multilabel classification (MLC) is an increasingly widespread data mining technique. Its goal is to categorize patterns in several non-exclusive groups, and it is applied in fields such as news categorization, image labeling and music classification. Comparatively speaking, MLC is a more complex task than multiclass and binary classification, since the classifier must learn the presence of various outputs at once from the same set of predictive variables. The own nature of the data the classifier has to deal with implies a certain complexity degree. How to measure this complexness level strictly from the data characteristics would be an interesting objective. At the same time, the strategy used to partition the data also influences the sample patterns the algorithm has at its disposal to train the classifier. In MLC random sampling is commonly used to accomplish this task.

This paper introduces TCS (*Theoretical Complexity Score*), a new characterization metric aimed to assess the intrinsic complexity of a multilabel dataset, as well as a novel stratified sampling method specifically designed to fit the traits of multilabeled data. A detailed description of both proposals is provided, along with empirical results of their suitability for their respective duties.

## Introduction

Unlike multiclass and binary classification, where the classifier has to predict only one output, multilabel classification (MLC) must learn the associations between the patterns' features and several outputs at once. Each output indicates if a certain label is relevant to the data sample or not, thus the algorithms have to work with a set of binary predictions. Nowadays MLC is being applied to automate tag suggestion [CRdJH15b], categorize text documents [KY04], label incoming images [DBdFF02], etc. A introduction to MLC and a recent review on MLC techniques and related topics can be found in [GV15] and [GV14], respectively.

Most of the aforementioned tasks involve working with large multilabel datasets (MLDs) having disparate numbers of input features, instances, labels, combinations of labels, etc. Undoubtedly some of these traits, such as the number of instances, determine at some extent the time necessary to train a classifier. Beyond this fact, it would be desirable to know in advance the difficulties a certain MLD can present and how its complexity can affect the classifier performance.

A second circumstance which potentially affects MLC algorithms performance is the way MLDs have been partitioned. There are MLDs containing only a few samples, sometimes only one, as representatives of rare labels. Random sampling, which is the mainstream strategy used in the multilabel field, can throw these few samples all on either the training or the test partition. Both cases will probably decrease the performance of the classifier.

The main aim of this paper is to study how the complexity of MLDs and the sampling strategy impacts classification results. For doing so, two proposals are introduced:

- A new characterization metric, called TCS (*Theoretical Complexity Score*), will allow to know the complexity of an MLD in advance, prior to use it to train a classifier. It is computed from the basic MLD traits.
- A novel stratified sampling method for partitioning datasets, aiming to improve label distribution among training and test partitions, thus providing the classifier a fairer representation of each label. It is built upon an stratification strategy, grouping instances containing labels with similar frequencies.

This paper is structured as follows. Section explains how different complexity factors influence classification results and introduces the TCS metric. In Section the problems of random sampling are described, and a new stratified sampling method is presented. The suitability of these two proposals is experimentally tested in Section . Lastly, in Section some conclusions are drawn.

## Assessing a Multilabel Dataset Complexity

Data complexity [HB02] is an intensively studied aspect in different fields, including classification. How to measure it and its influence in specific problems, such as imbalanced [LFGH11] learning and noise filtering [SLH13], have been already faced in traditional classification. Regarding MLC, some studies related to imbalance [CRdJH15a] measurement and other complexity facts, such as the concurrence among frequent and infrequent labels [CRdJH14], have been also published.

The interest here is to determine an intrinsic and easily interpretable complexity metric for MLDs. In this context, complexity has to be understood as the set of traits of the MLD that will make the learned model both more ineffective and inefficient.

## Factors Influencing the Complexity of MLDs

In order to design such a metric firstly the main traits of any MLD and its implications in learning a MLC model have to be analyzed. The considered factors are the following:

- **Number of data instances:** The number of rows in an MLD determine the amount of available patterns to train and then test any model. While it is true that a larger quantity of data samples also implies more time devoted to training, having more instances does not necessarily means that the resulting model will be more complex. In fact, training a classifier with enough representative patterns is usually associated to a better performance.
- **Number of input features:** In machine learning the curse of dimensionality [Bel56] is a very well-known problem. As the number of input features grows, so does the dimensions of the space where the patterns are located. Working in a high-dimensional space makes more difficult tasks such as measuring distances among patterns and finding analytical solutions. Most MLDs have a large number of features, thus it is a factor to be taken into account.
- **Number of labels:** In traditional classification the algorithms only have one output to learn, whether it is binary or multiclass. By contrast, MLDs have hundreds or thousands of labels. The larger is the number of labels the more complex would be the model to generate. There are many MLC methods based on binarization techniques [GS04, HFCB08, FHLMB08, RPHF11], and the number of labels has a direct impact in both the time used to train each binary classifier and the complexity of the overall algorithm. This, no doubt, is another factor to consider.
- **Number of labelsets:** The labels in an MLD appear producing different combinations, usually known as labelsets. The number of distinct labelsets is another aspect to bear in mind, since there are many MLC methods [BLSB04, TKV08, Rea08, TV07] based on training multiclass classifiers using the labelsets as class identifiers. Some of them produce simpler subsets of labels through pruning, random combinations and clustering approaches. In general, the larger is the amount of different combinations the more complex will be the final solution.

## Theoretical Complexity Score

Building on the premises just enumerated, the proposed TCS metric is computed as indicated in (1). Let  $f$  be the number of input features,  $k$  the number of labels and  $ls$  the number of distinct labelsets. The logarithm of the product of these three factors will provide a theoretical complexity score, based only on the basic traits of the MLD<sup>1</sup> and easier to interpret than the raw product.

$$TCS(D) = \log(f \times k \times ls) \tag{1}$$

The main goals in defining this metric, in addition to assess the complexity of an MLD, were ease of computation and interpretation, providing a straightforward measurement.

If there were an extremely simple MLD, having only one input attribute, two labels (on the contrary it would not be multilabel), and four different labelsets (the number of combinations two labels can produce), its TCS value would be  $\log(1 \times 2 \times 4) \approx 2$ . Table 1 shows the number of attributes, labels, labelsets and TCS for twenty MLDs commonly used on the literature, ordered according to their TCS value. As can

---

<sup>1</sup>In practice there would be other factors also influencing the classifiers performance, such as data sparseness, imbalance levels, concurrence among rare and frequent labels, etc.

Table 1: MLDs ordered according to their theoretical complexity score

Dataset	TCS	Attributes	Labels	Labelsets
emotions	9.364	72	6	27
scene	10.183	294	6	15
yeast	12.562	103	14	198
genbase	13.840	1 186	27	32
cal500	15.597	68	174	502
medical	15.629	1 449	45	94
enron	17.503	1 001	53	753
reuters	17.548	500	103	811
mediamill	18.191	120	101	6 555
corel16k001	19.722	500	153	4 803
corel5k	20.200	499	374	3 175
stackex-cs	20.532	635	274	4 749
bibtex	20.541	1 836	159	2 856
tmc2007	21.093	49 060	22	1 341
eurlex-sm	21.646	5 000	201	2 504
eurlex-dc	21.925	5 000	412	1 615
rcv1subset1	22.313	47 236	101	1 028
delicious	22.773	500	983	15 806
bookmarks	22.848	2 150	208	18 716
eurlex-ev	26.519	5 000	3 993	16 467

be seen emotions and scene, two of the most popular MLDs, are the simplest ones. The two MLDs from genetics/proteins field, genbase and yeast, are more complex. Multimedia datasets, such as mediamill and corel5k, are located at the middle of the table. Some of the MLDs coming from text media, such as delicious, bookmarks, EURLex, etc., appear as the most complex ones.

## Sampling Multilabel Datasets

Almost all studies and proposals in the multilabel field imply some classification experimentation. Hold out, 2×5 and 10 folds cross validation are among the most common schemes, always with the same strategy to chose the patterns included in train and test partitions, random sampling. Despite the fact that some other sampling strategies [STV11] have been described for some time, the random approach is still the most used option.

Random sampling does a good work in selecting training and test patterns when most labels have enough representation in the MLD. However, sometimes it could be a risky strategy. That some labels have only one or two patterns representing them in the MLD is quite usual. Random sampling can place all of them either in the training or the test partition. To avoid this problem a stratified sampling approach can be used.

## Stratified Sampling of MLDs

Stratified sampling is a usual technique in cross validation [RTL09] for traditional classification. Since only one class is assigned to each instance, it is possible to compute the distribution of each class in the whole dataset and then draw the equivalent proportion of samples for training and testing. On the contrary, samples in an MLD are associated to several labels at once. If one instance is chosen for the train partition because it holds a certain label, it must be taken into account that some other labels are also included in the operation since they jointly appear with the selected one.

In [STV11] a stratified iterative method for sampling MLDs is proposed. It goes label by label through the MLDs, choosing individual samples and updating a set of counters. Due to its iterative nature it is a slow method when compared with random sampling, specially with MLDs having thousands of labels. Nonetheless, the authors stated that it was able to improve the classifier performance while dealing with some MLDs.

## Stratified Random Sampling Method

The method outlined in Algorithm 1 is a new proposal to partition MLDs. It follows a stratified random sampling approach, but unlike the one in [STV11] it is not iterative by label.

---

**Algorithm 1** Partitioning method based on stratified random sampling

---

```

1: function STRATIFIED.KFOLDS(MLD  $D$ , Integer  $nfolds$ )
2:   for each instance  $i$  in  $D$  do
3:      $D_{i_w} \leftarrow \prod freq(l \in D_i)$  ▷ Weight for each instance
4:   end for
5:    $D \leftarrow \text{SortBy}(D_w)$  ▷ Sort instances according to their weight
6:   ▷ Group samples with similar weight in separate strata
7:   for  $i = 1$  to  $nfolds$  do
8:      $strata_i \leftarrow D_{|D|/nfolds*(i-1)} - D_{|D|/nfolds*i}$ 
9:   end for
10:  for  $i = 1$  to  $nfolds$  do ▷ Generate nfolds folds
11:    for  $j = 1$  to  $nfolds$  do ▷ Taking part of the samples in each stratum
12:       $trainfold_i \Leftarrow \text{drawRandomly}(strata_j, |D|/nfolds \times (nfolds - 1))$ 
13:       $testfold_i \Leftarrow strata_j - trainfold_j$  ▷ Remainder samples in stratum
14:    end for
15:  end for
16:  return ( $trainfold$ ,  $testfold$ )
17: end function

```

---

In line 3 a weight is computed for each instance in the MLD. It is obtained as the product of the relative frequencies of active labels in the data sample. If one or more rare labels appear in it, the score will be very low. On the contrary, the occurrence of one or more common labels will produce a higher value. The number of active labels also influences this score. The larger is the set of labels in the instance the lower will be the score. The goal is to group instances with a similar label distribution relying in a simple procedure.

Once the instances have been ordered according to their weight (line 5), they are divided into as many strata as folds have been requested. Each training partition gets a portion of each stratum proportional to

the number of samples in  $D$  and the number of folds. The remainder samples in the stratum are given to the test partition. The samples in each stratum are randomly picked.

## Experimentation

Aiming to validate the usefulness of the two proposals made in the previous sections, five MLDs with diverse TCS values have been selected from Table 1. Those are emotions, yeast, enron, stackex-cs and delicious. All of them can be downloaded from the R Ultimate Multilabel Dataset Repository [CCR<sup>+</sup>16], and they can be partitioned randomly or using the stratified strategy described in Section by means of the `mldr.datasets`<sup>2</sup> R package.

The datasets were partitioned using 10 fcv, once randomly and once with stratified random sampling. These partitions were given as input to tree multilabel classifiers, one based on binarization (BR [GS04]), one based on label combinations (LP [BLSB04]), and one on lazy learning adapted to multilabel data (ML-kNN [ZZ07]). From the results produced by the classifiers three general performance metrics, Accuracy (2), Precision (3) and Recall (4) have been computed to analyze the meaningfulness of the TCS metric. Another two more specific metrics, MacroPrecision and MacroRecall, have been obtained to compare the two sampling strategies. The macro-averaging strategy (5) allows the calculation of any standard performance metric label by label, then averaging to obtain the final measure. In these equations  $n$  is the number of instances in in the MLD,  $Y_i$  the real labelset associated to  $i$ -th instance,  $Z_i$  the predicted one,  $k$  the number of labels in  $\mathcal{L}$ , and  $TP$ ,  $FP$ ,  $TN$  and  $FN$  stand for *True Positives*, *False Positives*, *True Negatives* and *False Negatives*, respectively.

$$Accuracy = \frac{1}{n} \sum_{i=1}^n \frac{|Y_i \cap Z_i|}{|Y_i \cup Z_i|} \quad (2)$$

$$Precision = \frac{1}{n} \sum_{i=1}^n \frac{|Y_i \cap Z_i|}{|Z_i|} \quad (3)$$

$$Recall = \frac{1}{n} \sum_{i=1}^n \frac{|Y_i \cap Z_i|}{|Y_i|} \quad (4)$$

$$MacroMet = \frac{1}{k} \sum_{l \in \mathcal{L}} EvalMet(TP_l, FP_l, TN_l, FN_l) \quad (5)$$

## Influence of Complexity in Classifier Performance

To analyze how the intrinsic complexity of each MLD influences the classifiers performance, Figures 1 to 3 show for each classifier the Accuracy, Precision and Recall values along with TCS. The x-axis corresponds to the five MLDs ordered according to their complexity.

From these plots observation that higher TCS values are correlated to worse performances can be easily deduced. For the LP algorithm the three evaluation metrics show a similar behavior, whereas for BR and ML-kNN Precision seems to be less affected than Recall and Accuracy. To formally analyze this relationship, a Pearson correlation test was applied over the TCS and performance values for each metric and algorithm. The obtained results are the shown in Table 2. As can be seen, with the exception of Precision and BR, all

<sup>2</sup><https://cran.r-project.org/web/packages/mldr.datasets/index.html>




Figure 1: Performance measures with respect to TCS values for the BR algorithm.

the results are above 0.9 in absolute value, meaning that a strong correlation exists. The negative values imply an inverse relation, thus the higher is TCS the worse would be the result.

Table 2: Pearson correlation test results

Algorithm	Accuracy	Precision	Recall
BR	-0,91	-0,67	-0,93
LP	-0,95	-0,94	-0,93
ML-kNN	-0,96	-0,99	-0,95

## Influence of Sampling Strategy in Classifier Performance

Once the partitions for each MLD using the two sampling strategies are generated, it would be useful to know how potentially problematic cases affect each method. As mentioned above, some MLDs contain labels that could be considered as rare, since they only appear once or twice in the whole dataset. Using a 10 fcv scheme, it is easy that these singular cases fall down in the train partition almost always, leaving the test set with poor or null representation of these labels. The own sampling method guarantees that at least once they will appear in the test partition.

The plots in Fig. 4 show the amount of labels with one or none occurrences in the test set produced by each strategy. Here "Random" refers to the classical random approach and "Stratified" to the proposed stratified random sampling. The emotions MLD does not have any case, thus his plot would be empty.




Figure 2: Performance measures with respect to TCS values for the LP algorithm.

Looking at delicious, for instance, it can be verified that for folds 1, 3, 8 and 9 the random sampling clearly produces more problematic cases than the stratified approach. Only for fold 10 the result is definitely worse for the stratified strategy. With stackex-cs the differences appear to be smaller due to the y-axis scale. In more than half of the folds the stratified strategy worked better than the random one. The yeast MLD is not affected by the described problem as much as the other ones. There are two folds in which the random approach produced one problematic case, against only one for the proposed stratified method. Lastly, enron has the most mixed situation, with large and small differences in both ways.

The results produced by the classifiers were, in general, better for the stratified strategy in those partitions where it produced less problematic cases. The same was applicable for the random approach. Since the results obtained from cross validation are always average values, these differences tend to compensate among them. These final evaluation measures are the shown in Table 3. Best values are highlighted in bold.

Overall there is a tie between the two strategies. Although there are MLDs working better with the stratified one, such as yeast, and others with the random alternative, such as stackex-cs, the remainder MLDs reflect a mixed behavior. Even though there are some noticeable differences between the results produced by the two strategies, most of them are in the order of a few thousandths.

## Conclusions

The performance of a multilabel classifier is influenced by a plethora of circumstances, starting with the own model goodness, the learning process and the traits (imbalance, missing values, outliers, label concurrence,




Figure 3: Performance measures with respect to TCS values for the ML-kNN algorithm.

etc.) of the data used to train it. We hypothesized that two key aspects could be the inherent complexity of the data and the strategy used to partition the MLDs, and described two useful tools to face them.

With the proposed TCS metric the theoretical complexity of any MLD can be quickly and easily computed. As has been demonstrated with experimental results, a clear correlation between the TCS level and the performance of the tested MLC algorithms can be established. Therefore, this metric could be used to know in advance if an MLD would obtain better or worse classification results than others depending on their TCS values.

Regarding the sampling strategies to partition the datasets, the most used approach in MLC is the random way. It can produce some problems with certain MLDs, as has been explained, that could be solved with an stratified strategy. Such a method has been proposed, and its behavior has been compared with the standard random sampling. Although it clearly improved the balanced presence of rare labels among folds in some cases, the classifiers performance did not show fair overall differences. A further more extensive analysis, including additional MLDs, algorithms and sampling strategies, will be needed to determine which could be the best way for MLD partitioning.

**Acknowledgments:** This work was partially supported by the Spanish Ministry of Science and Technology under projects TIN2014-57251-P and TIN2012-33856, and the Andalusian regional projects P10-TIC-06858 and P11-TIC-7765.




Figure 4: Occurrences of problematic cases depending on the sampling strategy.

Table 3: Performance measures for each MLD, algorithm and sampling strategy

Dataset	Algorithm	Macro Precision		Macro Recall	
		Random	Stratified	Random	Stratified
delicious	BR	0.4882	<b>0.4919</b>	0.0696	<b>0.0705</b>
	LP	0.1113	<b>0.1114</b>	<b>0.1101</b>	0.1095
	ML-kNN	<b>0.6385</b>	0.6134	<b>0.0433</b>	0.0420
emotions	BR	0.5961	<b>0.5997</b>	0.5578	<b>0.5832</b>
	LP	<b>0.5761</b>	0.5573	0.5565	<b>0.5631</b>
	ML-kNN	<b>0.7439</b>	0.7122	<b>0.6051</b>	0.5837
enron	BR	0.4556	<b>0.4780</b>	<b>0.1684</b>	0.1667
	LP	<b>0.1855</b>	0.1696	<b>0.1657</b>	0.1536
	ML-kNN	0.5942	<b>0.6266</b>	0.0880	<b>0.0899</b>
stackex-cs	BR	<b>0.4026</b>	0.3864	<b>0.1160</b>	0.1156
	LP	<b>0.0964</b>	0.0937	<b>0.0911</b>	0.0847
	ML-kNN	<b>0.6242</b>	0.5876	<b>0.0200</b>	0.0184
yeast	BR	0.4425	<b>0.4576</b>	0.3817	<b>0.3971</b>
	LP	0.3764	<b>0.3784</b>	0.3762	<b>0.3814</b>
	ML-kNN	0.6783	<b>0.6803</b>	0.3503	<b>0.3515</b>

# Bibliography

- [Bel56] Richard Bellman. Dynamic programming and lagrange multipliers. *Proceedings of the National Academy of Sciences of the United States of America*, 42(10):767, 1956.
- [BLSB04] M. Boutell, J. Luo, X. Shen, and C. Brown. Learning multi-label scene classification. *Pattern Recognit.*, 37(9):1757–1771, 2004.
- [CCR<sup>+</sup>16] Francisco Charte, David Charte, Antonio Rivera, Mara Jos del Jesus, and Francisco Herrera. R ultimate multilabel dataset repository. In *Hybrid Artificial Intelligence Systems*, LNCS. Springer International Publishing, 2016.
- [CRdJH14] Francisco Charte, Antonio Rivera, Mara Jos del Jesus, and Francisco Herrera. Concurrence among imbalanced labels and its influence on multilabel resampling algorithms. In *Hybrid Artificial Intelligence Systems*, volume 8480 of *LNCS*, pages 110–121. Springer International Publishing, 2014.
- [CRdJH15a] Francisco Charte, Antonio J. Rivera, María J. del Jesus, and Francisco Herrera. Addressing imbalance in multilabel classification: Measures and random resampling algorithms. *Neurocomputing*, 163(0):3–16, 2015.
- [CRdJH15b] Francisco Charte, Antonio J. Rivera, Maria J. del Jesus, and Francisco Herrera. QUINTA: A question tagging assistant to improve the answering ratio in electronic forums. In *EUROCON 2015 - International Conference on Computer as a Tool (EUROCON)*, *IEEE*, pages 1–6, Sept 2015.
- [DBdFF02] P. Duygulu, K. Barnard, J. de Freitas, and D. Forsyth. Object Recognition as Machine Translation: Learning a Lexicon for a Fixed Image Vocabulary. In *Proc. 7th European Conf. on Computer Vision-Part IV, Copenhagen, Denmark, ECCV'02*, pages 97–112, 2002.
- [FHLMB08] Johannes Fürnkranz, Eyke Hüllermeier, Eneldo Loza Mencía, and Klaus Brinker. Multilabel classification via calibrated label ranking. *Mach. Learn.*, 73:133–153, 2008.
- [GS04] Shantanu Godbole and Sunita Sarawagi. Discriminative Methods for Multi-Labeled Classification. In *Advances in Knowl. Discovery and Data Mining*, volume 3056, pages 22–30, 2004.
- [GV14] E. Gibaja and S. Ventura. Multi-label learning: a review of the state of the art and ongoing research. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 4(6):411–444, 2014.

- [GV15] Eva Gibaja and Sebastián Ventura. A tutorial on multilabel learning. *ACM Computing Surveys*, 47(3):52:1–52:38, April 2015.
- [HB02] Tin Kamo Ho and Mitra Basu. Complexity measures of supervised classification problems. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(3):289–300, 2002.
- [HFCB08] Eyke Hüllermeier, Johannes Fürnkranz, Weiwei Cheng, and Klaus Brinker. Label ranking by learning pairwise preferences. *Artificial Intelligence*, 172(16):1897–1916, 2008.
- [KY04] Bryan Klimt and Yiming Yang. The Enron Corpus: A New Dataset for Email Classification Research. In *Proc. ECML’04, Pisa, Italy*, pages 217–226. 2004.
- [LFGH11] Julián Luengo, Alberto Fernández, Salvador García, and Francisco Herrera. Addressing data complexity for imbalanced data sets: analysis of smote-based oversampling and evolutionary undersampling. *Soft Computing*, 15(10):1909–1936, 2011.
- [Rea08] J. Read. A pruned problem transformation method for multi-label classification. In *Proc. 2008 New Zealand Computer Science Research Student Conference (NZCSRS 2008)*, pages 143–150, 2008.
- [RPHF11] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Mach. Learn.*, 85:333–359, 2011.
- [RTL09] Payam Refaeilzadeh, Lei Tang, and Huan Liu. Cross-validation. In *Encyclopedia of Database Systems*, pages 532–538. Springer US, 2009.
- [SLH13] José Antonio Sáez, Julián Luengo, and Francisco Herrera. Predicting noise filtering efficacy with data complexity measures for nearest neighbor classification. *Pattern Recognition*, 46(1):355 – 364, 2013.
- [STV11] Konstantinos Sechidis, Grigorios Tsoumakas, and Ioannis Vlahavas. On the stratification of multi-label data. In *Machine Learning and Knowledge Discovery in Databases*, pages 145–158. Springer, 2011.
- [TKV08] G. Tsoumakas, I. Katakis, and I. Vlahavas. Effective and Efficient Multilabel Classification in Domains with Large Number of Labels. In *Proc. ECML/PKDD Workshop on Mining Multidimensional Data, Antwerp, Belgium, MMD’08*, pages 30–44, 2008.
- [TV07] Grigorios Tsoumakas and Ioannis Vlahavas. Random k-labelsets: An ensemble method for multilabel classification. In *Proc. 18th European Conf. on Machine Learning, Warsaw, Poland, ECML’07*, volume 4701, pages 406–417, 2007.
- [ZZ07] M. Zhang and Z. Zhou. ML-KNN: A lazy learning approach to multi-label learning. *Pattern Recognit.*, 40(7):2038–2048, 2007.