



## Research Article

# Machine Learning to Assess Relatedness: The Advantage of Using Firm-Level Data

Giambattista Alborna <sup>1,2</sup> and Andrea Zaccaria <sup>1,3</sup>

<sup>1</sup>Enrico Fermi Center for Study and Research, Rome, Italy

<sup>2</sup>Sapienza University, Department of Physics, Rome, Italy

<sup>3</sup>Institute for Complex Systems-CNR, UOS Sapienza, Rome, Italy

Correspondence should be addressed to Giambattista Alborna; [albornagiambattista@gmail.com](mailto:albornagiambattista@gmail.com)

Received 9 April 2022; Accepted 18 June 2022; Published 31 July 2022

Academic Editor: Pierluigi Contucci

Copyright © 2022 Giambattista Alborna and Andrea Zaccaria. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The relatedness between a country or a firm and a product is a measure of the feasibility of that economic activity. As such, it is a driver for investments at a private and institutional level. Traditionally, relatedness is measured using networks derived by country-level co-occurrences of product pairs, that is counting how many countries export both. In this work, we compare networks and machine learning algorithms trained not only on country-level data, but also on firms, which is something not much studied due to the low availability of firm-level data. We quantitatively compare the different measures of relatedness, by using them to forecast the exports at the country and firm level, assuming that more related products have a higher likelihood to be exported in the future. Our results show that relatedness is scale dependent: the best assessments are obtained by using machine learning on the same typology of data one wants to predict. Moreover, we found that while relatedness measures based on country data are not suitable for firms, firm-level data are very informative also for the development of countries. In this sense, models built on firm data provide a better assessment of relatedness. We also discuss the effect of using parameter optimization and community detection algorithms to identify clusters of related companies and products, finding that a partition into a higher number of blocks decreases the computational time while maintaining a prediction performance well above the network-based benchmarks.

## 1. Introduction

Relatedness [1], a key tool of the economic complexity framework [2], refers either to the similarity between two economic activities or between an activity and an economic actor. As such, it is also known as coherence [3] in the standard economic literature. This concept can be easily applied to different sets of such activities, such as the export basket of countries [4, 5], the technology portfolios of companies [6–8], or regional diversification patterns [9]. In these cases, relatedness is a measure of the feasibility of an activity (e.g., exporting a product) with respect to what an economic actor already does. This tool is, at present, widely adopted by policymakers and institutions such as the World Bank Group [10, 11] and the European Commission [12, 13] to inform governments and the private sector with respect to

industrial and innovation policy, both at a country and regional level.

Relatedness being a general concept, the precise way to assess the similarity between two activities or the feasibility of an activity for an economic actor is, a priori, not determined. As a consequence, various formulations coexist in the literature; most of them are however related to the so-called co-occurrences, that is counting how many countries are exporting a couple of products (the more the counting, the more related the two products will be). This is equivalent to projecting the input data (a bipartite economic actor-activity network, typically a country-product network) into one of the two layers [14], usually the economic activities, for instance the exported products. For example, the projection of the bipartite country-product network into the layer of products gives rise to a monopartite network of products.

Among the various possibilities, Teece et al. [3] proposed to use the  $t$ -statistics of co-occurrences in industries with respect to a randomized diversification of firms. Hidalgo et al. [4] introduced the product space, in which the co-occurrences of exported products are normalized with respect to their ubiquity. Zaccaria et al. [5] normalized the co-occurrences with respect to both ubiquity and diversification, considering the nested structure of the country-product network, the idea being that highly diversified countries carry a relatively small amount of information. In the cases above, the result is an almost fully connected network, whose pictorial representation is not very informative and, most importantly, no effort is present to remove any possible noise. So various attempts to filter such projections are present in the literature. Saracco et al. [15] proposed to statistically validate the single links with respect to a null configuration model [16]. Cimini et al. [17] however showed that the adoption of different null models leads to different filtered networks; see also Dosi et al. [18] and Bottazzi and Pirino [19] for a critical discussion of the consequences of using different null models while computing relatedness.

This situation calls for a framework to systematically compare and validate the different relatedness measures. Our proposal is to use an out-of-sample prediction task for this purpose. Tacchella et al. [20] and Straccamore et al. [21] have shown that standard co-occurrence methods perform worse than autocorrelation benchmarks, and that tree-based machine learning algorithms such as random forest [22, 23] provide the present state-of-the-art with respect to the assessment of relatedness. Alborn et al. [24] described this approach in detail, providing a comparison between different machine learning algorithms.

Having established that the relatedness between a country and a product is better assessed by means of machine learning, a natural question arises, namely whether the country data provide an optimal assessment, given the fact that companies, and not countries, are actually producing the exported products. Moreover, often recommendations are given to the private sector, so one could expect that algorithms should be trained on companies, and not countries. In this article, we provide a systematic and quantitative comparison of machine learning and network-based approaches to forecasting new products both at the country and firm level. In particular, we will leverage a database of more than 70000 Italian firms and compare it with country-level data, providing a cross-database analysis for both training and testing. Our results provide quantitative evidence about which algorithm and which database should be used to optimally assess relatedness; moreover, we are able to economically motivate these results by investigating the different structures of the two databases and how the algorithms extract the relevant information.

## 2. Materials and Methods

In this section, we discuss our database, the metrics to compute relatedness, and the testing procedure.

**2.1. Firm-Level Data.** The Italian National Institute of Statistics (<http://www.istat.it>) provided data about the export of all Italian firms. After a preliminary cleaning procedure, we have 71826 Italian firms exporting at least two products in the period between 1996 and 2017 and at least one product every year from the one on which they exported their first product till 2017. The exported products are classified according to the UN-COMTRADE ([comtrade.un.org](http://comtrade.un.org)) Harmonized System, 1992 edition. This is a hierarchical classification encoded by a number of digits corresponding to different levels of aggregation. For our investigation we use 4 digits, corresponding to 1233 codes, defining as many different products. These data can be organized as a set of temporal bipartite networks, one for each year, linking Italian firms with their exported products; at first, the weight of each link is the export volume. This is equivalent to defining 22 matrices  $E(y)$  ( $y=1996..2017$ ) of size  $71826 \times 1233$ , where each row represents a firm and each column a product. The element  $E_{fp}(y)$  is the volume of product  $p$  (expressed in euros) that firm  $f$  exported during the year  $y$ .

**2.2. Country-Level Data.** Country-level data come from UN-COMTRADE database ([comtrade.un.org](http://comtrade.un.org)) and consist in the exports of 169 countries in the period between 1996 and 2017. All the considerations made above still apply. In order to match these data with the firm level, we use only the 1233 products that are present also in the Italian firms' data. We point out that, Italian economy being highly diversified, this corresponds to discarding less than 1% of products (0.8%). So, at the country level, we have 22 matrices (one per year) of size  $169 \times 1233$ .

**2.3. Data Preprocessing.** Since export volumes strongly depend on the size of both the economic actor (country or firm) and the specific product, the direct use of this quantity would introduce a strong bias. The usual solution in the economic complexity literature [4, 25] is to compute the RCA values (revealed comparative advantage) introduced by Balassa [26] defined as:

$$RCA_{fp}(y) = \frac{E_{fp}(y)}{\sum_{p'} E_{fp'}(y)} \frac{\sum_{f'} \sum_{p'} E_{f'p'}(y)}{\sum_{f'} E_{f'p}(y)}. \quad (1)$$

In this way the export is normalized with respect to both the total export of the firm and the product and, using a physics jargon, we go from the extensive variable  $E$  to an intensive one. In order to have a binary variable, we say that a product  $p$  is (competitively) exported by a firm  $f$  if its RCA is greater than 1 and with this threshold we define the binary matrix  $M$

$$M_{fp}(y) = \begin{cases} 1 & \text{if } RCA_{fp}(y) \geq 1 \\ 0 & \text{if } RCA_{fp}(y) < 1. \end{cases} \quad (2)$$

**2.4. Relatedness Measures.** The first aim of our analysis is to compare different approaches to measure the relatedness

between firms and products, that is how much a firm is close to being able to export a product. This is something largely studied when the economic actors are not firms but countries; as discussed in the Introduction, two types of approaches exist: complex networks [3–5] and supervised machine learning algorithms [20, 21, 24, 27]. Here we compare these methods in an out-of-sample forecast exercise both at the country and firm level, in which we assume that exporting more related products is easier. The output of both the network-based and the machine learning approach is an  $S$  matrix in which the element  $S_{fp}$  is the relatedness between firm  $f$  and product  $p$ .

**2.4.1. Network Models.** Traditionally, in order to measure the relatedness between a country and a product, one starts with a measure of the similarity or proximity between products, that can be visualized as a network of products. The next step is the computation of the density [4] or coherence [8]: the average similarity between the target product and the ones already exported by the target country. This is what we will call the relatedness between a country and a product.

To compute the proximity between two products, one counts how many countries export both, that is the number of co-occurrences. The weight of the link of the resulting network of products is this quantity possibly divided by a normalization factor. According to the latter, we can define different types of networks. In the product space [4] we divide the number of co-occurrences with the maximum ubiquity between the two products (i.e., how many countries export that product). In formula:

$$B_{pp'}^{PS} = \frac{1}{\max(u_p, u_{p'})} \sum_c M_{cp} M_{cp'}, \quad (3)$$

where  $u_p = \sum_c M_{cp}$ . However, the co-occurrence of products in a country that exports almost all the products is not so relevant like the one in a country that exports few products. This is a relevant problem, given the nested structure of the matrix  $M$  [28]. An improvement that considers this factor is the taxonomy network [5] in which each co-occurrence is also normalized with respect to the diversification of the country  $d_c = \sum_p M_{cp}$ :

$$B_{pp'}^{TN} = \frac{1}{\max(u_p, u_{p'})} \sum_c \frac{M_{cp} M_{cp'}}{d_c}. \quad (4)$$

Once we have the network  $B$ , we define the relatedness  $S$  between a country and a product by using the density [4]:

$$S_{cp} = \frac{\sum_{p'} M_{cp'} B_{pp'}}{\sum_{p'} B_{pp'}}. \quad (5)$$

In practice, we sum the export matrices  $E$  from 1996 to 2012 to obtain the total export either of the countries or of the firms, we compute the RCA and  $M$  values from the resulting matrix and we estimate all the weights of the links  $B$ . Then we apply the formula above using the  $M$  matrix of

2012 to compute the relatedness  $S$  either for the countries or for the firms that belong to the test set defined below.

**2.4.2. Random Forest.** The measure of relatedness given by the use of machine learning algorithms based on decision trees has been shown to provide a better assessment of the probability for the future exports of countries than network-based approaches [20]. In particular, it has been shown that random forest [23] and XGBoost [29] are the most performing algorithms for the task of assessing relatedness [24]. In this article, we decided to adopt the random forest (RF) since, even if with country-level data XGBoost gets slightly superior results [24], the computational time required to train a random forest is much lower, so it allows us to make a more complete analysis with a tuning of the hyperparameters and, as we will see, the use of community detection algorithms.

Before talking about the training procedure, we need to talk about how we split the data. When we are working with country-level data, we use all the countries both for the training and the testing of the algorithms, but when we work with firm-level data, we split the firms into three datasets: 20000 firms are used to train the algorithm, another 20000 are used for the validation procedure to make the tuning of the hyperparameters, and the remaining 31826 firms are used to do the test. The reason why we do not use all the firms in the training as we do with country-level data is that firms are much more than countries and using all of them would increase the computational time. The firms in each of the three datasets are chosen randomly.

For each product  $p$ , we build a random forest that has the task to predict if firms (or countries) will export  $p$  after 5 years. In the case of firm-level data, during the training procedure, the features are given by the concatenation of the 12 RCA matrices  $20000 \times 1233$  between 1996 and 2007 in order to have a single matrix  $240000 \times 1233$ , in which each row contains the RCA values in a year  $y$  of a firm that belongs to the training set. The labels of the training are given by the concatenation of the column  $p$  of the  $M$  matrices from 2001 to 2012. So, during the training, the model learns if with a certain configuration of the RCA values in year  $y$  (i.e., the export basket of a firm), the firm can start to export the product  $p$  after 5 years.

There are some hyperparameters that can be optimized to improve the performance of the random forest and to avoid overfitting. Here we take into consideration *max depth* and *min sample leaf* [30]. The value of the first hyperparameter regulates the maximum depth of the tree, if one of the trees of the random forest reaches this value during its construction, its training stops even if not all the training samples are perfectly classified. So, if the trees are very deep and the too many splits bring to overfitting, to avoid it we can lower the max depth. The value of the second hyperparameter regulates the minimum value of training samples that a leaf node must have. If during the training the algorithm finds a split that creates a node with less samples than the value of *min sample leaf*, this split is discarded. So, a high value of *min sample leaf* prevents the random forest

from creating nodes with few training samples that are the ones that bring to overfitting. In order to optimize the hyperparameters, we train different random forests varying the value of one of the hyperparameters at once, then we perform a test using the firms of a validation set and we choose the value of the hyperparameter that brings to the higher best  $F1$  score.

Once found the optimal values of *max depth* and *min sample leaf*, we train a random forest with these values and we do the test with the firms in the test set. Giving the RCA values of 2012 as input to the random forest related to a product  $p$ , it returns a vector with the column  $p$  of the  $S$  matrix, so with all the 1233 random forests we build the whole  $S$  matrix for the firms needed in the test set. The same procedure can be repeated for the 169 countries.

In Results section, we are going to compare the models trained on countries and the ones trained on firms to make predictions about countries or firms. This cross-test may change depending on the peculiar value used as an input. In particular, one may use directly RCA instead of the binary  $M$ . Our idea is to always use the most informative input. When we train the random forest on firm-level data and we also do the test on firm-level data, we use RCA; however, if we train the model on firm-level data to make predictions on countries, we use  $M$  values as input variables. The reason is that countries and firms have very different RCA values since they are very different objects. Indeed, the average nonzero RCA value for firms is about 500, while for countries it is 2, so for a firm  $RCA = 4$  is a low value, while for a country it is well above average. So, in what follows, when we show the results of a model that has been trained on firms to make predictions on countries the input variable is the  $M$  values, instead, if we make predictions on firms, the input variable is RCA. The same goes for the models trained on country-level data, if the predictions are on firms we use  $M$ , while if they are on countries, we use RCA.

**2.5. Testing Procedure.** In order to test the goodness of the relatedness assessment, we assume that firms (or countries) will export in the future products with the higher  $S$  (relatedness) values. In particular, we build the models discussed above by using data from 1996 to 2012, from which we compute the  $S$  matrix. The comparison of these relatedness measures with the  $M$  (2017) matrix can be seen as evaluating the output of a binary classifier. In particular, the hypothesis is that higher the  $S_{fp}$  the more likely will be that firm  $f$  will start to export product  $p$ . This is analogous to common machine learning classification exercises [30, 31], so in order to compare the goodness of the different relatedness metrics, we can use the performance indicators we introduce in Section 2.6. However, given the strong self-correlation of the export matrices, what interests us is not predicting if firms will export products that they already export, but if they will export new products. For this reason, when we do the comparison of the  $S$  matrix with  $M$  (2017), we consider only the activations of new products, or in other words, we remove the elements  $(f, p)$  that do not satisfy this requirement:

$$RCA_{fp}(2012) < 0.25. \quad (6)$$

In this way, we look at how good the model is in predicting the activation of new products by firms. The value 0.25 of the threshold follows [20, 24]. The idea is that using a threshold  $RCA < 1$  would increase the noise in the test set given by products whose RCA value for a firm fluctuates around 1. Moreover, predicting that a firm can be competitive in the export of a product on which its RCA value is already close to 1 is less interesting with respect to a firm genuinely becoming competitive. In order to check the robustness of our findings, we repeated the forecast exercise using different values of the threshold finding similar results. As already said, to build the model we use a set of 20000 firms (training set) and to make the comparison with  $M$  (2017) we use a separate set of 31826 firms (test set), in this way the test is out of sample because during the construction of the model no year between 2013 and 2017 is used and the firms from which the model learns are not the ones on which we make predictions.

When we work with countries, there are some differences, since the countries are only 169 we use all of them both to train the model and to do the test. The other difference is that with country-level data we use a stronger definition of activation in order to align the results with the ones published in [24]:

$$RCA_{cp}(y) < 0.25 \forall y \in [1996, 2012]. \quad (7)$$

The main reason we chose to use two different definitions of activations is that, while countries always export at least one product in the years from 1996 to 2012, a firm may have been created in a certain year and, before that year, it does not appear in the export data.

**2.6. Performance Indicators.** In this section, we describe the indicators which quantify the goodness of the forecast. When evaluating a binary classification, the choice of the performance indicator depends on both the research purpose and the database structure [32, 33]. In our case, the fraction of ones in the  $M$  matrices of firms is only the 0.4% of the total elements, the remaining being equal to zero. So, we have to deal with a very high class imbalance and for this reason we have to carefully choose our performance indicators. For instance, if we would use indicators that involve the true negatives like accuracy, they would get very high values because even if we do not guess any true positive, the number of true negatives will likely be huge. Here we make a quick description of the indicators we use:

- (i) Precision [32]: It is the number of true positives divided by all the positives, respectively, how many products we guess to be competitively exported by firms after 5 years and how many products we expected to be exported considering also the wrong predictions.
- (ii) P@K: The precision@K corresponds to the fraction of the top  $K$  positives that are correctly predicted

or, in other words, the fraction of elements that the model guesses if we ask it to tell us the  $K$  most probable positives.

- (iii)  $mP@K$  The mean precision@ $K$  is computed considering only the first  $K$  predicted products separately for each firm, then we look at how much of them are correct and finally we average on the firms. By using  $mP@K$ , we quantify the correctness of our possible recommendations of  $K$  products, on average, for a firm. In this sense,  $mP@K$  is a local measure of performance, while  $P@K$  considers the whole matrix.
- (iv) Recall [32]: It is the number of true positives divided by the sum of true positives and false negatives.
- (v)  $F1$  Score [34, 35]: Precision and recall vary by changing the scores' binarization threshold (the value we use to define positives and negatives). Usually the higher it is the precision, the lower it is the recall and vice versa.  $F1$  score is a harmonic mean of these two metrics and its value is high only if both of them are relatively high. We define as best  $F1$  score the  $F1$  score computed by finding the threshold that maximizes it.
- (vi) ROC-AUC [36, 37]: It is computed by ranking all the scores and computing for each possible threshold the true positive rate (TPR) and the false positive rate (FPR). In this way, we draw a curve in the TPR/FPR plane and the ROC-AUC corresponds to the area under this curve. It can be seen as the probability that, if we randomly select a positive and a negative element, the first will receive a higher score [38]. With highly imbalanced data, due to the high number of true negatives, the ROC-AUC tends to give too optimistic results [39, 40]. For a random classifier, ROC-AUC = 0.5.
- (vii) AUC-PR: The area under the precision-recall curve is the area under the curve that, in the plane defined by precision and recall, is obtained by varying the scores' binarization threshold. Since true negatives are not considered, its value is not misled by the class imbalance [39].
- (viii) MCC [41]: Matthews' correlation is computed using the scores' binarization threshold that maximizes the  $F1$  score. It is a metric that considers all the four classes of the confusion matrix and also the class imbalance issue [42, 43].

Precision, recall,  $F1$  score, and MCC require a threshold to define if the value of the score should be associated with a positive or negative prediction. In these cases, we chose the threshold that maximizes the  $F1$  score.

### 3. Results

*3.1. Random Forest on Firms vs Product Space on Countries: Worked Examples.* In this section, we present data-driven examples to compare the different forecasts on future exports

of the same Italian firms given by (i) a product space (PS) approach built on country-level data and (ii) a random forest (RF) trained on firm-level data. What we want to highlight is not only that the RF has more predictive power than the PS, as we will quantify better later, but also that the choice of the data with which we build the model is fundamental.

The first reason why a model that is trained on country-level data produces worse forecasts on firms' future exports is that what is similar from the point of view of a country is oftentimes not similar from the point of view of a firm too. To clarify this point we take an example from the data: in 2012 a historic firm (firm A) that deals with jewelry, in particular with corals, exported the products (Table 1).

The RF trained on firm-level data correctly predicts that this firm in 2017 will export the product with HS code 7113 that is *Jewellery articles of precious metal or of metal clad with precious metal*. However, PS wrongly predicts the product 0307 that is *Molluscs, whether in shell or not, live, fresh, chilled, frozen*. So, while the RF understood that firm A deals with jewelry, the PS recommended molluscs. The reason is that the firm exports corals and in a country where there are corals there are also molluscs. If you are a country and you have the sea you will have both firms that treat corals to export jewelry and firms that treat and export molluscs and other seafood products, but if you are a single firm either you treat jewelry or you treat seafood products. So, the PS found a relation between corals and molluscs that is relevant only if you are a country.

Another reason why the PS built on country data is not suitable for forecasting future exports of firms lies in the specialized nature of firms. A big difference between a country and a firm is that the first tends to diversify and export as many products as possible, while the second is specialized in a category of products [44]. When one builds the PS using country-level data one observes and counts the co-occurrences between different products following the nested pattern of the export data (see Figure 1(a)). In particular, simple products are exported by almost all countries and so the PS counts a large number of co-occurrences with the complex products. On the contrary, complex products are exported by only those countries that export most of the products, including the less sophisticated ones. So, one can expect that, by using such a model, when one tries to forecast the future exports of a firm specialized in complex products one can wrongly predict that firm to export random simple products.

In Figure 1 we show a real example to clarify this point.

On the left plot, we show the adjacency matrix of the bipartite firm-product network. Companies are on the vertical axis while products are on the horizontal axis. The ordering of rows and columns is given by the BRIM community detection algorithm [45]. The orange points highlight the products that firms exported in 2012. The evident modular structure of the firm-product matrix reflects the specialized nature of firms. We selected as a target firm an important company specialized in the design and production of kitchens (firm B) that in the last 20 years exported products for more than 500 million euros. The firm exported ten products in 2012, eight of which belong to the

TABLE 1: Example of exported products by firm A. Clearly, this firm deals with jewelry and corals and in the last 20 years it exported these products for about 10 million euros.

Code	Description	HS section
0508	Coral and similar materials	1
7103	Precious (excluding diamond) and semi-precious stone	14
7117	Imitation jewelry	14
9601	Ivory, tortoise-shell, horn, coral, and other animal carving material and articles	20

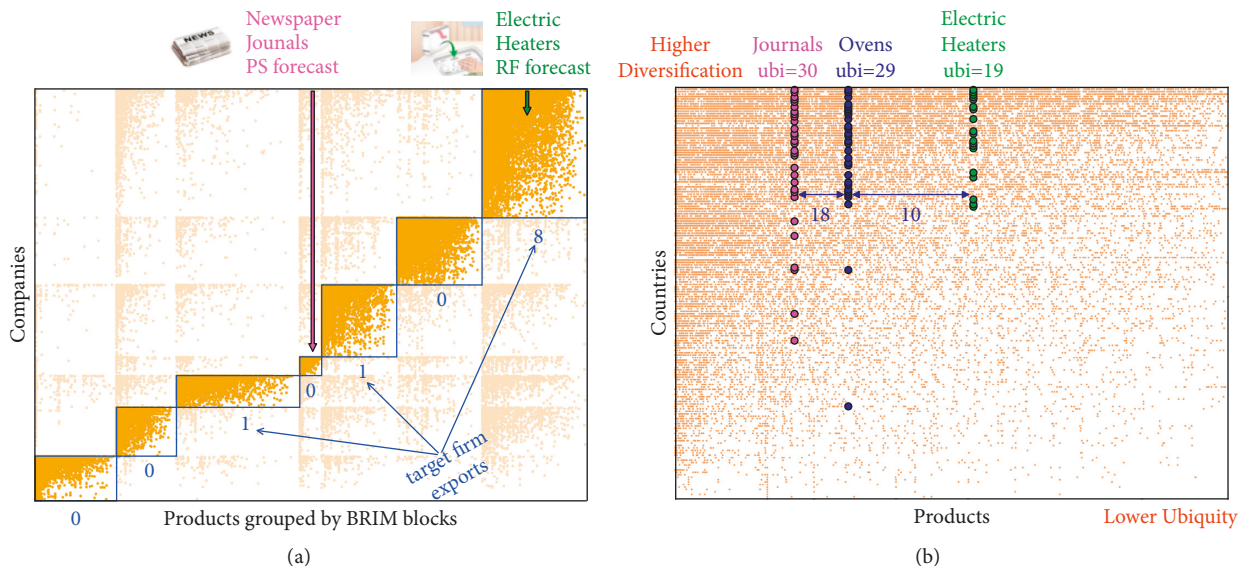


FIGURE 1: Visual representation of the predictions given by product space built on countries and random forest built on firms. (a) Matrix representation of the firm-product network. The products are grouped in blocks detected by the BRIM community detection algorithm. Under each block we report the number of products that belong to the exports of the target firm, firm B (selling large kitchens). The random forest prediction (electric heaters, green arrow) falls in the block in which the target firm has 8 products; the magenta arrow points to the most probable future export (newspapers) according to a product space model built on country-level data and falls in a block in which the firm has no products. (b) Country-product network. The products are sorted in decreasing order of ubiquity while the countries are sorted in increasing order of diversification. We highlighted in blue a product exported by the target firm—ovens. As expected, the product space forecast for the target firm is a more ubiquitous (simpler) product that is exported by many countries and so has many co-occurrences (18) with the ovens. The random forest forecast is a less ubiquitous (more complex) product and for this reason it has fewer co-occurrences (10) with the ovens.

same block. According to the PS model built on country-level data, the most related product, and so its best guess for a future export of the target firm, is *newspapers, journals, and periodicals* (magenta arrow). The un-relatedness of this product with respect to the target firm is self-evident. Moreover, newspapers belong to a different block than the one on which the firm is active, in particular, no products of this block are exported by the firm, and obviously, if we check these predictions, in 2017 the firm will not export journals. When we ask the same question to a RF built on firm-level data, the output is *electric water, space, and soils heaters*, which belongs to the same block where the target firm already exports 8 products. Moreover, in 2017 the firm will start exporting also electric heaters. It is evident that, while the RF understood what category of products the firm deals with, the PS did not. The fact is that, being built on country-level data, the PS did not learn the specialized nature of firms; what it knows is that among the countries there are many co-occurrences between newspapers and the products of the target firm, but this is only a consequence of

the fact that many countries export newspapers since it is a simple product. We can see this from the right plot. On the horizontal axis we report the products in decreasing order of ubiquity (that is highly anti-correlated with the complexity) and on the vertical axis we report the countries in increasing order of diversification. Newspapers in 2012 are exported by 30 countries; they are less sophisticated than electric heaters, which are exported by 18 countries. We highlighted with blue points one of the ten products exported in 2012 by the target firm, that is *Furnaces and ovens; industrial or laboratory* and it is exported by 29 countries. Journals and ovens have 18 co-occurrences, while electric heaters and ovens have only 10 co-occurrences, so it is evident that PS built on these data learns that ovens are more similar to journals than to electric heaters because of the nested structure of the matrix.

3.2. *Firm-Level Relatedness Outperforms.* In the previous section, we show that the modular structure of the firms' database allows for a better quantification of relatedness

with respect to the nested structure of the countries database. Now we compare a RF trained on country-level data with a RF built on firm-level data, and we compare the performance of the two models on the prediction of both the future exports of firms and the future exports of countries. The assumption is that a higher relatedness implies, on average, a higher probability that the country or firm will export the target product. While in the previous section we presented specific examples, here we provide a general, quantitative evaluation of the model's predictions, always using PS as a benchmark.

In Figure 2(a) we show the performance of different models when we try to predict the future exports of firms. The blue bars refer to a RF trained on country-level data, the orange ones refer to a RF trained on firm-level data and the green ones refer to a PS built on firm-level data. We show different indicators to show the robustness of our results; we rescaled their values to allow a visual comparison. As expected from the qualitative discussion of the previous section, RF trained on firm-level data is by far the best choice; in particular, the model trained on country-level data is totally unsuitable to make predictions about the future exports of firms. Figure 2(b), the performance refers to a prediction of the future exports of countries and, accordingly, now the PS is built on country-level data. This time the RF trained on country-level performs better and the reason is that there are relationships between products that are not present at the firm level, like the ones discussed above: a model built on country-level data is trained using observations like marine countries that export both corals and molluscs, while a model built on firm-level data is trained using samples that, if exporting corals, will never export molluscs. So, the former will use the co-occurrences between molluscs and corals while the latter will not learn them, because at the firm level this relationship does not exist. However, the difference between country-level and firm-level RF is smaller with respect to the plot on the left. Strikingly, the RF trained on firm-level data performs better than the PS trained on country-level data even if we are making predictions on countries. We can conclude that machine learning models trained at the firm level are able to extract a measure of relatedness that is relevant also at the country level, and in particular, more relevant than the information that the PS is able to extract even at the country level. In conclusion, firm-level data provide a relatedness measure that is objectively better than the one given by country-level data.

**3.3. Model Comparison.** In this section we compare different models to predict the exports of firms. In particular, we compare RF with network models like PS and taxonomy network (TN), and we will also show the results obtained by using a quasi-trivial benchmark, RCA itself. Indeed, one may think to assume an autocorrelation model and consider as prediction score the RCA value in year  $y = 2012$  of a firm  $f$  on a product  $p$ . The resulting matrix  $S$  is the relatedness between firms and products and it is treated exactly in the

way we show in Sections 2.5 and 2.6. The higher the RCA, the higher the likelihood that  $f$  will start to export  $p$ .

In Figure 3 we show a radar plot in which the performance of the PS (green line) built on firms is used to normalize the other scores.

Each of the vertices in the radar plot refers to a different metric and the area of each polygon is a proxy of the total performance of the corresponding model. The brown and purple lines refer to the PS and the RF built on country-level data, respectively. They not only perform worse than all models trained on firm-level data, but they underperform the RCA predictions too (red line). The orange line is the TN built on firm-level data, which is slightly better than the PS. The RF built on firms vastly outperforms all the other models.

In Table 2 we compare the results of the models using all the performance metrics described in the Methods section. All the models are trained on firm-level data. The large majority of the classified elements are true negatives because of the high class imbalance of the problem; so metrics that involve the true negatives like accuracy should be avoided since they would give very high results only because it is very easy to predict a true negative. For instance, the ROC-AUC is very high for all the models except the RCA one and provides misleading results.

**3.4. Random Forest Optimization: Leveraging Modular Structure and Hyperparameters.** As we have shown in the previous sections, the firm-product export matrix has a modular structure. In this section we investigate if such a structure can be exploited to improve the prediction performance of machine learning; in particular, we perform a community detection on the bipartite graph and we train each RF by giving as input only the RCA values of the products that belong to the same block of the target product. Community detection is characterized by a number of different algorithms [46], so it is natural to consider various possibilities to build the partitions. Two natural block decompositions can be derived from the hierarchical structure of the harmonized system classification: the 1233 4-digits products can be organized in 21 sections or 96 chapters (the latter corresponding to a 2-digit aggregation level). Moreover, using community detection algorithms we can find other partitions; in this paper, we use BRIM [45], BILOUVAIN [47], and IBN [48].

To which extent the performance of RF can be improved by tuning its hyperparameters is still debated [49]. Here we discuss the effect of changing two of these parameters in assessing relatedness: *max depth* and *min sample leaf*. The default values are *max depth* =  $\infty$  and *min sample leaf* = 1 [30]. However, this choice can lead to overfitting, because each decision tree is expanded up to a perfect classification of the training sample.

In Figure 4, we compare the performance of different RF models trained using the partitions given by different community detection models and with different choices of the two hyperparameters. In particular, we compare both the

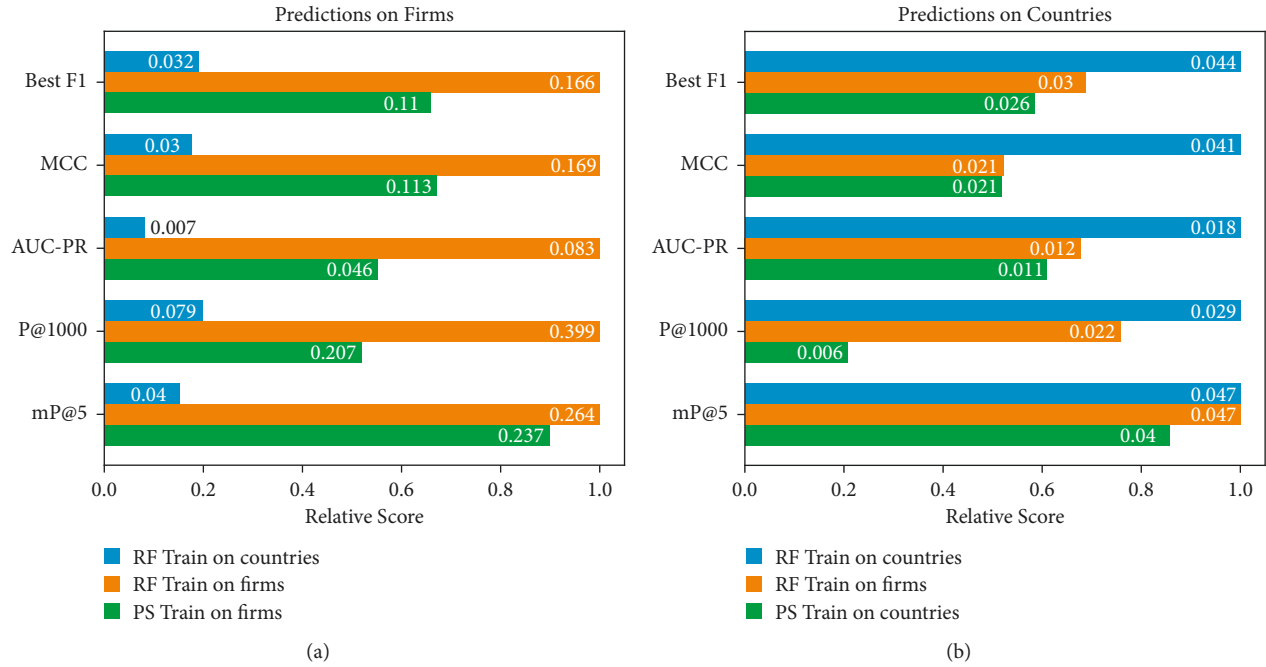


FIGURE 2: Comparison using different metrics between random forest built on firms and built on countries when we predict future exports of both firms (a) and countries (b). The orange bars refer to a random forest that has been trained on firm-level data, the blue bar refers to a random forest trained on country-level data, and the green bar refers to a product space built on firms (a) or country (b) data. The best assessment of relatedness at the country and at firm level is given by the relative random forest. However, the relatedness computed at the firm level performs better than the product space also at the country level.

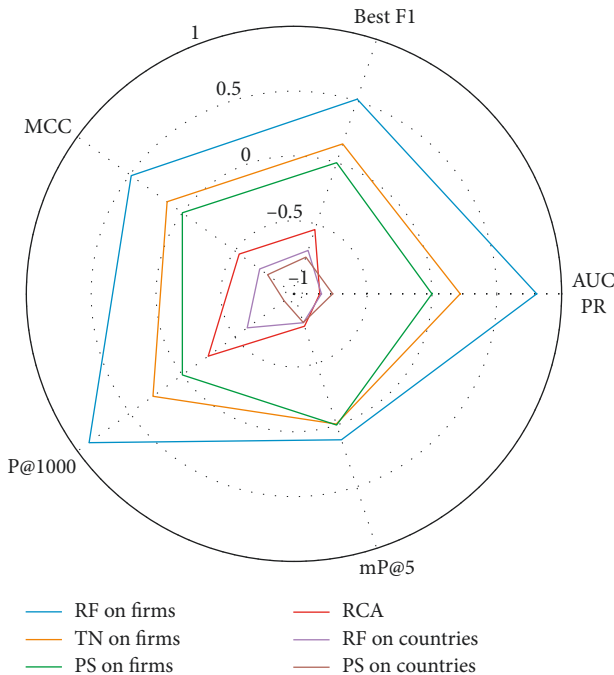


FIGURE 3: Model comparison in assessing firms' relatedness. Each vertex refers to a performance indicator normalized with respect to a product space built on firms. Models built on countries perform worse than the RCA benchmark. Taxonomy network performs better than the product space. The random forest trained on firms results to be the best model overall.

prediction performance (quantified by the best  $F1$  score) and computational time on a standard desktop computer.

On the right plot, each point represents one RF model with a *max depth* (circles) or *min sample leaf* (triangles) optimization. Close to each point we report the partition criteria that defines the blocks seen by the RFs: since we train one model for each product, each product is predicted by using the block it belongs to. Here 1-block means that we use all available data (so all products see all products) and *BRIM*<sup>2</sup> means that we applied the BRIM algorithm twice. On the horizontal axis, we report the training time and on the vertical axis the best  $F1$  score. The color of the points represents the number of blocks resulting by the corresponding partition. On the left plot, we compare the range of the prediction performance spanned by the different RFs with other prediction approaches: it is evident that all variations of RF outperform the other models.

The results of this analysis are:

- (1) *min sample leaf* tuning brings to better predictions, but *max depth* optimization speeds up the training time;
- (2) Prediction performance is higher when using a low number of blocks (for instance, using BRIM with 8 blocks or no partition at all). We can deduce that the RF is, in a sense, able to recognize the blocks on its own. However, by using the BRIM blocks we can reach the same prediction power in less time;



TABLE 2: Prediction performance of firm-based models. Our problem is characterized by a huge number of true negatives, and this precludes the use of metrics such as accuracy and ROC-AUC. In any case, the random forest outperforms all other models.

	RCA	Random forest	Product space	Taxonomy network
Best F1	0.050	0.166	0.110	0.126
AUC-PR	0.006	0.083	0.046	0.056
ROC-AUC	0.519	0.936	0.919	0.928
Precision@1000	0.156	0.399	0.207	0.265
mP@5	0.047	0.264	0.237	0.235
MCC	0.052	0.169	0.113	0.130
TP	2703	15584	11482	12706
FP	33387	100292	126310	116773
FN	69037	56156	60258	59034
TN	38952640	38885735	38859717	38869254




FIGURE 4: Performance of the random forest by varying the hyperparameters and the partitions that define the feature vector given as an input for each product. The results are optimized by tuning of *max depth* (circles) or *min sample leaf* (triangles). We also write the adopted partition and the color of the points represent the resulting number of blocks. On the horizontal axis, we report training time and on the vertical axis a performance indicator. Using more data (larger blocks for each product) provides better performance but takes longer time for training; smaller blocks lead to faster but less precise results. In any case, as evinced by the zoom given by the red lines on the plot on the left, random forest always outperforms the other models.

- (3) The more blocks we define with the community detection models, the more the RF is trained quickly; however, the performance tends to decrease;
- (4) The 2-digit aggregation represents a bad definition of relatedness, since the RF trained with these blocks is the one that performs worse; for instance, IBN defines almost the same number of blocks and requires about the same computational time, but performs significantly better.
- (5) Also the 21 HS sections do not represent a good definition of relatedness. This can be seen by comparing the performance with using BRIM<sup>2</sup>, which provides the same performance with more blocks (42) and so less computational time;
- (6) Even if we consider the worst model, that is the 2-digit blocks and the optimization on *max depth*, the best F1 score is significantly better than the one of the network models.

Now we motivate these results by investigating how these choices influence the training of the RF.

Result 1: *max depth* speeds up the training time because it represents a more drastic constraint than *min sample leaf*. Indeed, the average depth of a tree without any constraint is about 60, while the optimal value of *max depth* we find is usually less than 10 (it is 15 only if we do not use blocks), and shallower trees are trained faster. On the other hand, *min sample leaf* is a less drastic cut: what changes is at the level of the leaf nodes, so it is targeted to the removal only of the splits that bring to overfitting, and for this reason the performance is better. We can imagine that we have a real tree with some sick leaves and we have to remove the sick leaves knowing that the probability to have a sick leaf is proportional to the length of the branch. What we can do is either cut all the branches longer than a certain threshold or remove the sick leaves one by one. The first option corresponds to a *max depth* tuning, it is faster, but it reduces the quality of the tree. The second option corresponds to a *min sample leaf* tuning, it requires more time, but the quality of the resulting tree is better.

Results 2 and 3: since the RF is able to recognize the blocks on its own, if we provide a good partition of the

products what we can obtain is, at most, that we do not decrease its predictive power; however, if the selected partition contains too many blocks, we reduce the information the RF can learn, since each model has fewer products to see, and in this way we have a decrease in performance. However, with a higher number of blocks the training is faster for two reasons: the first is that the input has less features and the second is that, without a lot of products that have nothing to do with what we want to predict, the decision trees need less cuts and a lower depth.

Results 4 and 5: let us consider the jewelry firm we discussed in Section 3.1. As we can see from Table 1, its products are spread into different HS sections and chapters. So, the HS does not provide good partitions for relatedness analyses at the firm level.

Result 6: this result has practical consequences. If one wants to speed up the training of the algorithm through the use of (possibly good) partitions and a *max depth* optimization, in any case, the RF will outperform all network models. Note, however, that realistic applications do not usually require real-time investigations.

#### 4. Conclusions

The concept of relatedness, or coherence, is usually applied to quantify the closeness between an economic actor such as a country or a firm and an activity such as competitively exporting a given product. The possible practical applications of relatedness assessments are widespread; for instance, policymakers and institutions may want to quantify how much a developing country is far from entering into a given market (given its present export diversification) before deciding on an investment strategy, or if a new product is feasible given the present export basket of a firm. In both the mainstream and in the economic complexity literature, relatedness is measured in two steps. First of all, one builds a network of economic activities, typically products, in which the weights of the links are given by the so-called co-occurrences: the more countries export both products, the more the two will be similar. Different ways of building such networks coexist in the literature. The second step consists in computing the relatedness as the average similarity between the exports of a given country and the target product. In this article, we discuss and investigate two radical improvements: the use of supervised machine learning and firm-level, instead of country-level, data. In order to quantitatively compare the resulting different measures of relatedness, we test them against a forecast task, the assumption being that, on average, an economic actor will likely diversify in products that are relatively more related to. By means of both specific examples and general statistical assessments, we are able to show that: (i) machine learning, and in particular random forest, outperforms network-based methods regardless of the data typology; (ii) firm-level data provide a better assessment of relatedness, in the sense that while a model built on country-level data is totally unsuitable to predict future exports of firms, a model built on firm-level data is still able to accurately predict future exports of countries. This is due to the relative specialization of firms

that accurately tracks the similarity between products. On the contrary, successful countries are highly diversified, providing misleading co-occurrences; (iii) community detection algorithms provide partitions in the subsets of products which reduce the computational effort needed to train the algorithms, and (iv) regardless of the method used to build the relatedness measure, the optimal strategy is to train the forecast model using data of the same typology one wants to predict (in particular, firm-level data to build relatedness measures to be used at the firm level).

In summary, in order to compute the feasibility of a product for a firm, one should use machine learning algorithms trained on firm-level data, since the widespread use of co-occurrences computed at the country level leads to poor assessments of the relatedness.

These results open up a number of consequent investigations. First, the very same exercise should be replicated with different kinds of human activities, for instance, patents. Indeed, the relatedness between technological sectors is usually measured by counting country-level co-occurrences, and this assessment very likely suffers from the same issues which we have exposed, and solved, here. Second, relatedness enters in a number of derived quantities which are used to characterize the diversification strategies of countries, firms, and regions. The robustness of these quantities should be checked in light of the findings hereby reported. Finally, the validation strategy we propose to quantitatively compare the different relatedness measures—a rigorous out-of-sample forecast exercise—could be applied, more in general, to the various concepts used in economic complexity, in order to scientifically validate or falsify the different approaches, an issue of general relevance in the physics of complex systems.

#### Data Availability

The data generated and analyzed during the study are not publicly available for legal reasons, but are available from the corresponding author upon reasonable request.

#### Disclosure

A preprint has previously been published [50].

#### Conflicts of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

#### Acknowledgments

This work was supported by the Centro Ricerche Enrico Fermi Research Project “Complessità in Economia.” The authors thank ISTAT for providing the Italian firms data, in particular Dr. Stefano Menghinello and Dr. Cristina Lanzi. The authors thank Francis Farrelly for data encryption and pre-processing, and Luciano Pietronero for useful discussions.

## References

- [1] C. A. Hidalgo, P. A. Balland, R. Boschma, M. Delgado, M. Feldman, and K. Frenken, "The principle of relatedness," in *Proceedings of the International conference on complex systems*, pp. 451–457, Springer, July 2018.
- [2] C. A. Hidalgo, "Economic complexity theory and applications," *Nature Reviews Physics*, vol. 3, no. 2, pp. 92–113, 2021.
- [3] D. J. Teece, R. Rumelt, G. Dosi, and S. Winter, "Understanding corporate coherence," *Journal of Economic Behavior & Organization*, vol. 23, no. 1, pp. 1–30, 1994.
- [4] C. A. Hidalgo, B. Klinger, A. L. Barabási, and R. Hausmann, "The product space conditions the development of nations," *Science*, vol. 317, no. 5837, pp. 482–487, 2007.
- [5] A. Zaccaria, M. Cristelli, A. Tacchella, and L. Pietronero, "How the taxonomy of products drives the economic development of countries," *PLoS One*, vol. 9, no. 12, p. e113770, Article ID e113770, 2014.
- [6] S. Breschi, F. Lissoni, and F. Malerba, "Knowledge-relatedness in firm technological diversification," *Research Policy*, vol. 32, no. 1, pp. 69–87, 2003.
- [7] L. Nesta and P. P. Saviotti, "Coherence of the knowledge base and the firm's innovative performance: evidence from the U.S. Pharmaceutical industry," *The Journal of Industrial Economics*, vol. 53, no. 1, pp. 123–142, 2005.
- [8] E. Pugliese, L. Napolitano, A. Zaccaria, and L. Pietronero, "Coherent diversification in corporate technological portfolios," *PLoS One*, vol. 14, no. 10, p. e0223403, 2019.
- [9] F. Neffke, M. Henning, and R. Boschma, "How do regions diversify over time? Industry relatedness and the development of new growth paths in regions," *Economic Geography*, vol. 87, no. 3, pp. 237–265, 2011.
- [10] J. Lin, M. Cader, and L. Pietronero, *What African Industrial Development Can Learn from East Asian Successes*, The Role of Complexity and Economic Fitness, 2020.
- [11] A. Zaccaria, S. Mishra, M. Z. Cader, and L. Pietronero, *Integrating Services in the Economic Fitness Approach*, p. 8485, World Bank Policy Research Working Paper, 2018.
- [12] E. Pugliese and A. Tacchella, *Economic Complexity for Competitiveness and Innovation: A Novel Bottom-Up Strategy Linking Global and Regional Capacities*, Joint Research Centre (Seville site), 2020.
- [13] E. Pugliese and A. Tacchella, *Economic Complexity Analytics: Country Factsheets*, Joint Research Centre (Seville site), 2021.
- [14] T. Zhou, J. Ren, M. Medo, and Y. C. Zhang, "Bipartite network projection and personal recommendation," *Physical Review*, vol. 76, no. 4, Article ID 046115, 2007.
- [15] F. Saracco, M. J. Straka, R. D. Clemente, A. Gabrielli, G. Caldarelli, and T. Squartini, "Inferring monopartite projections of bipartite networks: an entropy-based approach," *New Journal of Physics*, vol. 19, no. 5, p. 053022, Article ID 053022, 2017.
- [16] T. Squartini and D. Garlaschelli, "Analytical maximum-likelihood method to detect patterns in real networks," *New Journal of Physics*, vol. 13, no. 8, p. 083001, Article ID 083001, 2011.
- [17] G. Cimini, A. Carra, L. Didomenicantonio, and A. Zaccaria, *Meta-validation of Bipartite Network Projections*, arXiv preprint arXiv:210503391, 2021.
- [18] G. Dosi, M. Grazzi, and D. Moschella, "What do firms know? What do they produce? A new look at the relationship between patenting profiles and patterns of product diversification," *Small Business Economics*, vol. 48, no. 2, pp. 413–429, 2017.
- [19] G. Bottazzi and D. Pirino, *Measuring Industry Relatedness and Corporate Coherence*, LEM Working Paper Series, 2010.
- [20] A. Tacchella, A. Zaccaria, M. Miccheli, and L. Pietronero, "Relatedness in the era of machine learning," 2021, <http://arxiv.org/abs/210306017>.
- [21] M. Straccamore, L. Pietronero, and A. Zaccaria, "Which will be your firm's next technology? Comparison between machine learning and network-based algorithms," 2021, <http://arxiv.org/abs/211002004>.
- [22] L. Breiman, "Bagging predictors," *Machine Learning*, vol. 24, no. 2, pp. 123–140, 1996.
- [23] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [24] G. Albora, L. Pietronero, A. Tacchella, and A. Zaccaria, "Product Progression: a machine learning approach to forecasting industrial upgrading," 2021, <http://arxiv.org/abs/210515018>.
- [25] A. Tacchella, M. Cristelli, G. Caldarelli, A. Gabrielli, and L. Pietronero, "A new metrics for countries' fitness and products' complexity," *Scientific Reports*, vol. 2, no. 1, p. 723, 2012.
- [26] B. Balassa, "Trade liberalisation and "revealed" comparative advantage," *The Manchester School*, vol. 33, no. 2, pp. 99–123, 1965.
- [27] N. X. Che and B. B. Bakker, "Intelligent export diversification," *IMF Working Papers*, vol. 20, no. 175, p. 2020, 2020.
- [28] M. S. Mariani, Z. M. Ren, J. Bascompte, and C. J. Tessone, "Nestedness in complex networks: observation, emergence, and implications," *Physics Reports*, vol. 813, pp. 1–90, 2019.
- [29] T. Chen and C. Guestrin, "Xgboost: a scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, ACM, CA, USA, August 2016.
- [30] G. Aurélien, *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly Media, 2018.
- [31] S. B. Kotsiantis, I. Zaharakis, P. Pintelas et al., "Supervised machine learning: a review of classification techniques," *Emerging artificial intelligence applications in computer engineering*, vol. 160, no. 1, pp. 3–24, 2007.
- [32] D. M. Powers, "Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation," *Journal of Machine Learning Technologies*, vol. 2, no. 1, pp. 37–63, 2011.
- [33] R. Caruana and A. Niculescu-Mizil, "An empirical comparison of supervised learning algorithms," in *Proceedings of the 23rd international conference on Machine learning*, pp. 161–168, Pittsburgh, PA, 2006.
- [34] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.
- [35] C. J. Van Rijsbergen, "Foundation of evaluation," *Journal of Documentation*, vol. 30, no. 4, pp. 365–373, 1974.
- [36] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.
- [37] K. Hajian-Tilaki, "Receiver operating characteristic (ROC) curve analysis for medical diagnostic test evaluation," *Caspian journal of internal medicine*, vol. 4, no. 2, pp. 627–635, 2013.
- [38] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.
- [39] T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets," *PLoS One*, vol. 10, no. 3, p. e0118432, Article ID e0118432, 2015.

- [40] A. Fernández, S. García, M. Galar, R. C. Prati, B. Krawczyk, and F. Herrera, *Learning from Imbalanced Data Sets*, Springer, 2018.
- [41] B. W. Matthews, “Comparison of the predicted and observed secondary structure of T4 phage lysozyme,” *Biochimica et Biophysica Acta (BBA) - Protein Structure*, vol. 405, no. 2, pp. 442–451, 1975.
- [42] D. Chicco and G. Jurman, “The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation,” *BMC Genomics*, vol. 21, no. 1, p. 6, 2020.
- [43] S. Boughorbel, F. Jarray, and M. El-Anbari, “Optimal classifier for imbalanced data using Matthews Correlation Coefficient metric,” *PLoS One*, vol. 12, no. 6, p. e0177678, Article ID e0177678, 2017.
- [44] M. J. Bruno, F. Saracco, T. Squartini, and M. Dueñas, “Colombian export capabilities: building the firms-products network,” *Entropy*, vol. 20, p. 785, 2018.
- [45] M. J. Barber, “Modularity and community detection in bipartite networks,” *Physical Review*, vol. 76, no. 6, Article ID 066102, 2007.
- [46] S. Fortunato, “Community detection in graphs,” *Physics Reports*, vol. 486, no. 3-5, pp. 75–174, 2010.
- [47] V. A. Traag and Y. Van DoorenNesterov, “Narrow scope for resolution-limit-free community detection,” *Physical Review*, vol. 84, no. 1, Article ID 016114, 2011.
- [48] A. Solé-Ribalta, C. J. Tessone, M. S. Mariani, and J. Borge-Holthoefer, “Revealing in-block nestedness: detection and benchmarking,” *Physical Review*, vol. 97, no. 6, Article ID 062302, 2018.
- [49] P. Probst, M. N. Wright, and A. L. Boulesteix, “Hyperparameters and tuning strategies for random forest,” *WIREs Data Mining and Knowledge Discovery*, vol. 9, no. 3, Article ID e1301, 2019.
- [50] G. Albora and A. Zaccaria, “Machine learning to assess relatedness: the advantage of using firm-level data,” 2022, <http://arxiv.org/abs/220200458>.