WP/20/175

IMF Working Paper

Intelligent Export Diversification: An Export Recommendation System with Machine Learning

by Natasha Che

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

WP/20/175

IMF Working Paper

Intelligent Export Diversification: An Export Recommendation System with Machine Learning

by Natasha Che

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Western Hemisphere Department

Intelligent Export Diversification: An Export Recommendation System with Machine Learning

Prepared by Natasha Che

Authorized for distribution by Bas Bakker

August 2020

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

This paper presents a set of collaborative filtering algorithms that produce product recommendations to diversify and optimize a country's export structure in support of sustainable long-term growth. The recommendation system is able to accurately predict the historical trends in export content and structure for high-growth countries, such as China, India, Poland, and Chile, over 20-year spans. As a contemporary case study, the system is applied to Paraguay, to create recommendations for the country's export diversification strategy.

JEL Classification Numbers: F1, F4, O1, O4

Keywords: export diversification, comparative advantage, machine learning, collaborative filtering, economic growth, international trade

Author's E-Mail Address: nche@imf.org

1 Introduction

There is an ongoing debate in the economic development field about the use of proactive industrial policies to diversify the economy and facilitate growth-enhancing structural changes (e.g. Aiginger and Rodrik, 2020). This is partly due to the meteoric rise of China and other East Asian emerging market economies, many of which have governments that are deliberate and active in guiding export diversification and sectoral structural change. These are seen, justified or not, by the rest of the developing world as fresh examples to borrow from. On the other hand, the classical free-market approach to development has yielded, at best, mixed results for many developing countries in the past decades. Although a major commodity price boom in the 2000s lifted growth performance of many commodity exporters, the majority of which are developing economies in Latin America and Africa, in the wake of the subsequent price drop, many countries began to realize that they were facing a serious sustainability issue in their growth model.

As a result, there is an increasing demand for proactive policy actions to guide a country's export diversification strategy. But the tools and frameworks for designing such strategies remain scarce. In particular, although the benefit of export diversification has strong empirical support (see Section 2), the questions of what industries to diversify into and how to diversify have no easy answers from a policy standpoint.

For the first question, classical trade theory suggests that countries should export what they are relatively good at producing, i.e. following comparative advantages. But that advice yields little practical value in informing policy decisions about which specific sectors and industries a country should try to foster. After all, how does one ascertain comparative advantages in a product category which a country has not, up to now, invested much in or exported much of? Similarly, trade theories predict that developing countries tend to have comparative advantages in labor-intensive exports and should stay away from capital-intensive heavy industries. But in reality, comparative advantages contain far more dimensions than capital and labor. Some of them are linearly quantifiable, others are not. The

matter becomes even more complicated when we take into account the fact that comparative advantages evolve as a country grows, but there is no accurate measurement to comprehensively detect such evolution, at least in the short term. It may be easy to see that exporting rocket engines is not a comparative advantage of Guatemala (not today, anyway). But how about sweaters and plastic containers? Which would be a better product category for the country to diversify into? General theories often don't go very far in providing realistic insights.

On a micro level, diversification into new export categories happens serendipitously, is often introduced by new foreign investments, and frequently involves knowledge transfer and capacity spillover from existing, adjacent export categories. In some cases, the *potential*, or *latent comparative advantage* in an export category may not materialize, until one foreign company, scouting the globe for new production bases, decides to invest in the country. Such decisions involve a multitude of considerations, and the country's comparative advantages in the pure economic sense are only part of them. In other cases, domestic exporters may manage to branch into a new export category related to their existing exports, once they've accumulated sufficient knowledge, e.g. in production scaling and logistics, and established distribution channels in importing countries, through learning by doing from their past export experiences. In other words, a country may have a latent comparative advantage in product X. But it is not observed until actual exports of product X happen with at least some success. The occurrence and timing of the such events are highly unpredictable. So it's reasonable to assume that at any given time, the set of products that a country has a latent comparative advantage in is almost always bigger than the set of actual product categories that the country has demonstrated competence in exporting.

Therefore, the question of what exports a country should diversify into is really a question of identifying the product categories that the country may have a latent comparative advantage in. The goal of the present paper is to answer this question by algorithmically exploring the similarities among different product groups and among different countries.

The inspiration of the paper’s approach comes from two key observations. First of all, as mentioned above, products that involve similar production resources and knowledge tend to show up in the export basket together. A country that has successfully exported beef can branch into, with some effort, exporting dairy. A country that has mastered the trade of exporting desktop computer hardware is in a better position to produce and export cellphones, than otherwise. Therefore, the products in a country’s existing export basket contains valuable information regarding what other products the country can get good at producing. Secondly, countries with similar comparative advantages tend to export similar products. Bangladesh and Vietnam are both successful in exporting garments because of the countries’ shared abundance in low cost labor. New Zealand and Uruguay both specialize in cattle exports partly because of the high availability of pasture land. In other words, the export baskets of *similar countries* contain information about the comparative advantages the countries share.

The paper quantifies these insights to characterize a country’s latent comparative advantages and produce export diversification recommendations, using machine learning algorithms that implement *collaborative filtering*, an approach used widely by online commercial applications for their recommender systems. A recommender system based on collaborative filtering uses the revealed preferences of a group of users to make predictions about the preferences of a user similar to the group. There are numerous applications of this approach in the e-commerce space. For example, Amazon.com recommends new products to a customer by looking at the customer’s purchase history and the purchase records of other customers of a similar taste. After a user watches one movie, Netflix recommends to the user similar movies, using a ”people who have watched this movie also watched...” approach.

The paper is the first economic research to apply the methodology of collaborative filtering in analyzing export diversification. The literature closest to the present paper is the studies on the so-called *product space* and its implication for diversification and growth (e.g. Hausmann & Klinger 2007, Hidalgo & Hausmann 2009). Like the current paper, this strand

of research seeks to understand a country's diversification potentials by looking at the *relatedness* among products. But there are two key differences. The first is a technical one. The product-space literature uses a probability formula to represent the relatedness, or *proximity* between two products.¹ While this approach is easy to understand and makes the subsequent analysis computationally simple, it is at the cost of not fully using the information contained in the data matrix of world exports. In contrast, the collaborative filtering algorithms of the present paper make better use of the data to detect the unique blend of characteristics of countries and products, to make potentially better recommendations. Of course, it is at the cost of requiring more computational resources and being more difficult to explain with linear logic, which is a common drawback of many machine learning algorithms. The second difference is one of perspectives. The product-space literature makes specific value judgments about the worthiness of different products for diversification. A product's diversification value is seen to roughly depend on 1) how "complex" it is, meaning, how much sophisticated knowledge is required to produce the product, and 2) how closely related the product is to other more complex products. Again, specific formulas are used to calculate a complexity score for each product and for its proximity to other products. The rationale behind such judgments is a reasonable one—more complex products contain higher value added, use more human capital, face less global competition, and products that are "bridges" to more complex products make it easier for a country to move up the international value chain. And some empirical evidence shows that diversifying into these products is supposed to be better for growth. However, several issues emerge when this model is used for recommending export products to specific countries. First, it may be inconsistent with the framework of comparative advantage, which the product-space analysis is built on. By essentially assigning each product a score of merit, the model has a tendency to recommend products it deems universally worthy to countries of different circumstances (e.g. industrial products are good,

¹Specifically, the proximity between product A and product B is defined as the probability that a country exports product A given that it exports product B, or vice versa. For example, suppose that 17 countries export wine, 24 export grapes and 11 export both, all with revealed comparative advantage. Then, the proximity between the wine and the grapes is $11/24 = 0.46$.

commodities are bad). In the extreme— though improbably— scenario where all countries internalize the same merit rankings for developing their export structure, there would be no comparative advantages to speak of. Secondly, to come up with tractable, universally applied scores for concepts such as "proximity" and "complexity", strong assumptions need to be made to drastically reduce the feature dimensions of reality and throw away valuable country- and product- specific information, which may reduce the model's usefulness in generating realistic export recommendations for individual countries. In contrast, the present paper does not impose any opinion regarding the diversification value of any product. Instead, it seeks to fully exploit the information contained in the product-level export data to make realistic recommendations on specific countries' diversification possibilities.

It is important to note that the paper's recommendation methodology not only produces suggestions of individual product categories that may be good candidates for diversification, it can also yield important insights regarding the overall structure of a country's export basket and its optimal path of evolution that may benefit long-term growth. It's also worth noting that although the paper's algorithms and the product recommendations they produce can potentially be very useful to policy makers evaluating their industrial policy options and private investors entering new markets, they are no substitutes for detailed analyses of the viability of certain industries in a country from multidimensional perspectives. In addition, it goes without saying that unveiling the industries with latent comparative advantages does not automatically translate into specific policy recommendations.

The paper is organized as follows. Section 2 presents the literature and stylized facts regarding the relationship between export diversification and growth. Section 3 describes the empirical data used and related variable definitions. Section 4 introduces the three collaborative filtering algorithms that are the main workhorses of the paper's recommendation system. Section 5 applies the recommendation system to the historical data of selected countries and evaluates the system's performance. As a contemporary country case study, Section 6 uses the algorithms to analyze the export structure of Paraguay and diversification

potential. Section 7 concludes.

2 The Case for Optimized Diversification

The relationship between export diversification and countries' economic performance has been extensively studied in the literature. Overall, existing research asserts that export diversification is a key element in the economic development process, particularly for developing and emerging market countries trying to catch up with their advanced peers. Various studies provide evidence of a positive association between export diversification and economic development (e.g. Imbs and Wacziarg, 2003; Klinger and Lederman 2004 and 2011; Cadot et al., 2011).²

Figure 1: Number of High-RCA Exports of Select High Growth Countries

Numerous country studies also supports the benefits of export diversification. For example, Feenstra and Kee (2008) use data from a large set of countries exporting to the US, to

²Cadot et al. marks USD 25,000 per capita GDP at purchasing-power-parity as the turning point of re-concentration mostly due to the adjustments of rich countries at the extensive margin.

show that a sustained increase in export diversification results in increases in productivity and a notable increase in the GDP of the exporters. IMF (2014) finds that diversification in exports and in domestic production has been conducive to faster economic growth in LICs. Al-Marhubi (2000) provides similar findings within a set of developing economies. Balaguer and Cantavella-Jorda (2004) find that export variety plays a key role in Spain's economic development. And Herzer and Danzinger (2006) report a positive impact of export diversification on economic growth of Chile. Research also points to a positive association between export diversification and macroeconomic stability (e.g. IMF, 2014). These empirical findings can be easily confirmed with country-specific data. Figure 1 plots the evolution of export diversity for several poster-child high growth countries in recent decades. Here the extent of export diversification is approximated by the number of export products with a *Revealed Comparative Advantage Score* (RCA) > 1 . (See Section 3 for more details on definition and measurement). All of these countries experienced significant export diversification along their path to income convergence. A common story among them: certain institutional and policy changes triggered broader acceptance to international trade and foreign investments, which in turn allowed the country to capitalize on its latent comparative advantages and start exporting a wider range of products. As a result, the number of products with relatively high RCA scores went up, while the RCA scores of those few products that used to dominate the country's export basket went down. In contrast, Figure 2 plots the export diversity trend for a few countries where income convergence has been relatively slow. In these countries, export diversification has either been stagnant or decreasing.

The relationship between export diversification and economic growth, as well as between diversification and the stability of growth, can be observed in cross-country data as well. Examining the export data by SIC 4-digit industry.³ reveals the following stylized facts. First of all, higher-income economies tend to have more diversified export structures. Regressing the number of export products in which a country has "revealed comparative advantages"

³See Section 3 for a more detailed description of the underlining data.

Figure 2: Number of High-RCA Exports of Select Low Growth Countries

(RCA)– from this point on, these will be called *high-RCA exports*– on its real GDP per capita relative to the US level indicates a strong positive relationship between the two, even after the size of the country is controlled for (Figure 3).

Consistent with this fact, regressing the average RCA score⁴ among a country's high-RCA exports on its income level relative to the US shows a negative relationship between the two, controlling for country size. A lower average RCA score tends to indicate more diversification, i.e. few products overly dominate the export basket. Thus this negative relationship indicates that richer countries tend to earn their export income from a wider variety of products.

In addition, the more diversified a country's exports are, the faster it grows subsequently. This is observed by regressing the average annual growth rate of countries between 2007 and 2017 on the number of high-RCA exports the countries had in 2007, controlling for the

⁴See Section 3 for the definition of RCA score.

Figure 3: Number of High-RCA Exports vs Income Level, Partial Regression Plot

initial GDP per capita level (Figure 4).⁵ The same pattern holds in a panel regression with country fixed effects.

Thirdly, countries with more diversified export structures have lower growth volatility. Regressing the standard deviation of annual real GDP growth from 2007 to 2017 on the number of high-RCA exports a country had in 2007, controlling for initial GDP level and country size, indicates a strong negative relationship between the two (Figure 5). A similar regression using panel data shows the same pattern.

However, not all types of diversification are created equal, and diversification for its own sake is hardly a recipe for sustainable growth. A foundational idea of the classical international trade theory is that under free trade, countries will tend to export what they are relatively good at producing, i.e. products they have a comparative advantage in. "Diversi-

⁵The regression is run both with country size control and without. The results are not significantly different.

Figure 4: Real GDP Growth vs Number of High-RCA Exports, Partial Regression Plot

Figure 5: Growth Volatility vs Number of High-RCA Exports, Partial Regression Plot

“fying” into industries that are misaligned with a country’s current endowment fundamentals, as the former Soviet-block nations did after World War II through industrial policies that aimed to accelerate industrialization, has negative growth consequences (see e.g. Lin, 2009). On the other end of the spectrum, delayed industrialization also leads to negative growth outcomes, as the experience of many resource-rich countries that are entrenched in their overdependence on commodity exports has shown (e.g. Frankel, 2010). Relatedly, Hausmann et al. (2007) finds that countries that export more sophisticated, or knowledge-intensive products, tend to grow faster, controlling for initial income levels.

Overall, the literature has shown that choosing the right diversification strategy is no easy matter for any country. Considerations of a country’s current comparative advantage need to be balanced with evolutionary visions of structural change in production and exports that could deliver growth for the long term. Racing towards building advanced industries with no respect for a country’s current endowments is as ill advised as being stuck in the status quo. The task of the present paper is to construct a diversification recommendation system that takes into account a country’s current fundamentals and future potentials, to create *optimized* diversification strategies that support long-term growth.

3 Data and Definitions

The source data of the paper comes from the UN Comtrade database that consists of annual export value data by the 786 SITC 4-digit product categories, for around 260 countries.⁶

An important concept used throughout the paper is *Revealed Comparative Advantage* (RCA). The RCA indicator, first introduced by Balassa & Noland (1965), is a popular measure to calculate the relative importance of a product in a country’s export basket.

⁶The data is cleaned by the Atlas of Economic Complexity. The number of countries available varies by year.

Formally, the RCA score of country i in product j can be calculated as:

$$RCA_{ij} = \frac{E_{ij}/E_i}{E_j/\sum_{i' \in I} E_{i'}}$$

where E_{ij} is the export value of product i from country j , E_i is the total export values of country i , E_j is the total exports of product j from all countries around the world, and $\sum_{i' \in I} E_{i'}$ is the total world exports.

Throughout the paper, a *high-RCA product* for country i is defined as a product with its $RCA_{ij} > 1$. Mathematically, it means that the product's share in the country's export portfolio is greater than its share in the total world exports, which can be seen as an indication that the country has a comparative advantage in the product. For example, vehicle exports were about 12 percent of total world exports in 2017, while they constituted 22 percent of total exports from Mexico. Therefore, $RCA_{ij} = 22/12 = 1.8$ for Mexico's vehicle exports in 2017. Since it is > 1 , according to our criteria, Mexico has a revealed comparative advantage in automobiles. Or to put it another way, automobiles is a high-RCA product for Mexico.

It's worth emphasizing again that in the context of the present paper, "diversification" does not mean merely increasing the total number of products that a country exports. As discussed in Section 2, the goal of diversification should be to increase the types of products that a country has *comparative advantages* in or specializes in, i.e. increasing the number of high-RCA products. If the goal of diversification were to simply increase the number of exports, then hypothetically, a country that exports every goods under the sun with their export shares exactly matching the shares of those products in total world exports would be a perfect diversifier. In reality, that would be neither a possible nor a recommendable development strategy. The challenge facing the typical developing nations is that they specialize in too few things, and the few specializations completely dominate the country's export basket— for example, it's not uncommon for the key commodity exports of some resource-rich countries to have a RCA score greater than 100 or even higher. The real

diversification challenge of the developing world is therefore to discover more products that a country is good at exporting and reduce the dominance of the few existing exports.

Before we dive into the algorithms, we should also take a minute to discuss the measurement issue of diversification that involves the classification of products. How different should the products be in order to be counted as distinct exports? The SITC classification system provides 5 levels of product categorization with increasing details. Which level, if applicable at all, should we use to adequately measure diversification? It's probably obvious that at the SITC 2-digit level, the products are different enough to be counted as distinct categories, e.g. exporting "live animals" (00) is certainly different from exporting "vegetables and fruits" (05). But things may get murkier the further we go up on the levels. Should expanding from "bovine animals" (0011) to "swine" (0013) be counted as diversification? How about an export expansion from "sheep" (00121) to "goats" (00122)?

To answer this question, we shall keep in mind the purpose and benefit of export diversification, which is to create new growth engines and to reduce economic volatility. So any activities to expand the export basket should have the potential to support at least one of these two goals, for it to be counted as diversification in the economic sense. Examining the SITC product lists leads the present paper to conclude that the SITC 4-digit classification is by and large the appropriate categorization level to measure export diversification. Take one of the examples above. Although "bovine animals" (0011) and "swine" (0013) are both animal products, from the exporters' point of view, the market dynamics of these two are very different, including demand growth prospect and geography of demand. For example, for Paraguay, an economy heavily relying on agricultural exports, diversifying from exporting cows to pigs has opened up new markets, which faced less competition from neighboring countries that traditionally export similar products. In other words, it appears to fulfill the purpose of diversification. Let's take another two 4-digit product classes that may look even more similar: "bovine meat, fresh or chilled" (0111) and "bovine meat, frozen" (0112). To the casual onlookers, these may seem basically the same product. However, these two ex-

ports actually involve very different technologies in production, storage and transportation, which has important implications for where these exports go— 0112 can reach destinations that are much farther away compared to 0111. And expanding the geographical reach of a country’s exports is generally beneficial for boosting growth and reducing growth volatility, i.e. it meets the goal of diversification.

To be sure, the same reasoning may not apply to every product class in the 4-digit list, and it may not apply to every country equally. The SITC classification is a statistical tool, not a tool designed specifically for diversification analysis. According to the manual for the provisional Central Product Classification, the product grouping was done according to the products’ differences in: 1) the raw or basic material, 2) the degree of processing, 3) the use or function, and 4) economic activities. A product’s diversification value is likely to stem from one or more of these factors, but there is no absolute match between the two. To further complicate the issue, the size of the underlining economic entity also matters. Expanding production from whole milk to non-fat milk may yield ”diversification” benefit for a dairy company, but it may hardly make any economic difference for the export structure of a country. Similarly, the same product classification level may not be equally suitable for analyzing diversification needs for countries of drastically different sizes. But these caveats aside, the SITC 4-digit product grouping appears to work well overall for the analysis of the present paper, as Section 5 will demonstrate.

4 Empirical Methodology

The paper constructs a recommendation system for export categories based on three algorithms widely used in online collaborative filtering recommender systems: product-based and country-based K-nearest neighbors (KNN), and Singular Value Decomposition (SVD). All methods used in the paper produce the so-called ”top-N recommendations”— the goal of the exercise is to generate a list of N product categories that a country should export the

most of. The algorithms produce the list by predicting the RCA scores of different products for the underling country, using the training dataset of export values by country and SITC 4-digit product, and recommending the N products with the highest predicted RCA scores.

The underlining data used in the recommendation system can be represented as a $m \times n$ matrix R , where m is the number of countries in the database, and n is the total number of SITC 4-digit products. The content of R , i.e. r_{ij} , is country i 's RCA score in product j . R is a sparse matrix due to the fact that each country only exports a subset of the products in the SITC universe. In the case that country i does not export any product j , $r_{ij} = 0$. If an implementation uses multiple years of export data, then each country-year is a row in R , i.e. $m = c \times y$, where c = the number of countries in the dataset, and y = the number of years included. In most versions of implementations discussed below, $y = 1$, i.e. if the task is to generate export recommendations for country i in 2017, only the cross-country export data for 2017 is included in the training set.⁷

4.1 Neighborhood-Based Algorithms

4.1.1 Product-Based KNN

KNN is one of the most frequently used methods in solving classification and pattern recognition problems, and is a popular approach in constructing recommender systems. The basic idea of KNN is learning by analogy—classifying the test sample by comparing it to the set of training samples most similar to it. Different KNN implementations vary in terms of their choices of how the similarity between input vectors is calculated. In the present paper, the cosine similarity score is used as the similarity measure.

The intuition behind the current paper's product-based KNN implementation is simple—first look at what products a country already has a revealed comparative advantage in, and then recommend other products that are similar to those products. To explain the approach

⁷The paper experimented with including multiple years of data in the training set, but found no significant improvement in the evaluation metrics, while the model took longer to compute as the size of m increases.

in more details, let's first rewrite the RCA score matrix R as:

$$R = \begin{bmatrix} \mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n \end{bmatrix}$$

where $\mathbf{p}_j$ is a vector of length m that represents the RCA scores of product j for all the m countries in the sample:

$$\mathbf{p}_j = \begin{bmatrix} r_{1j} \\ r_{2j} \\ \vdots \\ r_{mj} \end{bmatrix}$$

In machine learning terminology, each product in the sample has m features. The cosine similarity between products i and j is equal to $\frac{\mathbf{p}_j \cdot \mathbf{p}_i}{\|\mathbf{p}_j\| \|\mathbf{p}_i\|}$, which ranges from -1, when the two vectors are the exact opposite, to 1, when the two are exactly the same. The intuition behind this is that by comparing the two sets of countries that export i and j , and how important the products are in the countries' export baskets, information can be inferred regarding how closely related the two products are.

The implementation of the product-based KNN recommender for country i involves the following steps:

1. Represent each product in the SITC 4-digit product space as a vector of RCA scores, $\mathbf{p}_j$.
2. Select the set of K products in which country i has a revealed comparative advantage, i.e. $r_{ij} > 1$. Let's call it the high-RCA product set of country i .
3. For each $j \in [1, n]$, calculate the predicted value of r_{ij} as the weighted average RCA score of the high-RCA product set, weighted by the cosine similarity between product j and the products in the country's high-RCA set.

4. The recommended products for country i are the N products with the highest predicted r_{ij} values.

4.1.2 Country-Based KNN

Instead of recommending product categories related to a countries' existing export products, the problem at hand can also be thought of as "recommending" other countries similar to the country in question, in other words, finding a group of countries that are similar to country i . And because these countries have similar comparative advantages, the products they export, beyond the ones country i is already exporting, can be good candidates for diversification for country i .

More specifically, the RCA score matrix R can be represented as:

$$R = \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{q}_2 \\ \vdots \\ \mathbf{q}_m \end{bmatrix}$$

where $\mathbf{q}_i$ is a vector of length n that represent country i 's RCA scores for the n product categories in the SITC 4-digit product space. The execution of the country-based KNN algorithm for country i can then be broadly described as follows:

1. Calculate the cosine similarity score between $\mathbf{q}_i$ and $\mathbf{q}_j$, where $1 \leq j \leq m$, and $j \neq i$.
2. Select a set of K countries with the highest similarity scores to country i .
3. For each $l \in [1, n]$, calculate the predicted value of r_{il} as the weighted average RCA score of product l across the K countries, weighted by the similarity score between each country and country i .

4. The recommended products for country i are the N products with the highest predicted r_{ij} values.

It is important to note that although the product-based and country-based KNN recommenders apply similar algorithmic logic, the differences in the perspectives of the two methods lead to different recommendation results, as will be demonstrated in the next section. Generally speaking, since in the data sample $n > m$,⁸ it may be easier to identify the relatedness between products with more accuracy than to identify similar countries, which make the product-based KNN a superior approach. The results presented in the next section seem to confirm this hypothesis.

4.2 Matrix Factorization Algorithm

The KNN algorithms, though intuitive and easy to implement, suffer from some significant drawbacks. First of all, these algorithms have limited scalability. As the sizes of m and n increase, the amount of computation required to calculate the similarity scores increases at $O(n)$ time, reducing the performance of the algorithm on larger data sets. Another disadvantage of the KNN algorithms is their problem with sparse data. Since the KNN algorithms require explicitly calculating similarities among vectors, the calculation becomes increasingly inaccurate when there's a lot of missing data in matrix R . This problem is exacerbated by the fact that the algorithm essentially treats each row of the product vector (or the country vector) as independent features of equal importance, which is not the most efficient use of information in the data, and also makes missing rows generally more damaging, compared to algorithms that impose some discretion on the relative importance of different data points. For the current use case, the first drawback is not a big concern, as the m and n of the country-product space are relatively small, especially when we do not include multiple years in the calculation. The second drawback is more problematic, as it implies that the

⁸There are close to 800 product categories in the SITC 4-digit product space, while there are just over 250 countries in the sample.

KNN algorithms would perform worse on countries that are significantly under-diversified, i.e. lots of missing entries in R for these countries. This would potentially defeat the purpose of the exercise, as under-diversified countries are arguably the ones that are most in need of diversification recommendations.

The Singular Value Decomposition (SVD) algorithm provides a possible remedy to the problem. SVD is a matrix factorization technique widely used in dimensionality reduction and principal component analysis. The basic idea is that matrix R can be decomposed into three matrices:

$$R = USV'$$

where U and V are two orthogonal matrices of size $m \times r$ and $n \times r$ respectively. r is the rank of R . And S is a $r \times r$ diagonal matrix, with the singular values of R as its diagonal elements, sorted in the order of decreasing magnitude.

The main purpose of the decomposition is to represent the products and countries as combinations of the *latent factors* in the data, which are implicit, orthogonal features that can be used to characterize the entire country-product space. U represents the relationship between countries and the latent factors, while V' represents the similarity between products and latent factors. The diagonal elements of S can be thought of as the relative scaling values assigned to various latent factors.

To give an extremely simplified example, suppose the matrix R can be summarized by three independent latent factors: labor, land, and knowledge. Row i of matrix U represents the comparative advantage of country i as a combination of the latent factors. $\mathbf{u}_i = \begin{bmatrix} .55, .4, .05 \end{bmatrix}$ would mean that country i 's profile can be described as 50% labor, 40% land, and 5% knowledge— a resource-rich, developing country. Column j of matrix V' represents the characteristics of product j as a combination of the latent factors. Thus $\mathbf{v}_j = \begin{bmatrix} .15, .05, .80 \end{bmatrix}'$ means that the production of product j can be characterized as 15% labor, 5% land, and 80% knowledge— a technology-intensive product that requires mostly intangible inputs. $r_{ij} = \mathbf{u}_i \cdot \mathbf{v}_j$, scaled by the appropriate diagonal element in S . It's not

difficult to see that r_{ij} would be relatively small, i.e. country i does not have a comparative advantage in producing product j . This is, of course, a very hypothetical example. In practice, the latent factors computed by the optimization algorithm are not human-interpretable, and only serve as features that more efficiently characterize the data.

The goal of the SVD algorithm is essentially to find the best estimations of U and V' , and then produce recommendations based on estimated $\hat{r}_{ij} = \hat{\mathbf{u}}_i \cdot \hat{\mathbf{v}}_j$. In practice, because R is already sparse, observing the orthogonality constraints for U and V' becomes computationally untenable. The execution of the algorithm thus centers on solving the following optimization problem:

$$\min_{\mathbf{u}_i, \mathbf{v}_j} \sum_{r_{ij} \in R} (r_{ij} - \mathbf{u}_i \cdot \mathbf{v}_j)^2 + \lambda(\|\mathbf{u}_i\|^2 + \|\mathbf{v}_j\|^2)$$

where λ is a regularization factor. The minimization is performed with stochastic gradient descent, using the python Surprise library for building recommender systems. The recommended products for country i are the products with the highest predicted $\hat{r}_{ij}$ value.

5 Results

One common obstacle in testing recommender systems is the difficulty to evaluate the quality of the recommendations. The common evaluation metrics for machine learning models, such as accuracy score and mean squared error, tend to be backward looking, essentially assessing the model performance by looking at how good the model fits the test data. Although that's a useful part of the evaluation, the goal of a recommendation system is not just correctly fitting existing data, but to make good recommendations for what have not yet manifested in the data. However, the definition of a good recommendation varies according to the user case, and is by no means undebatable.

In e-commerce, one of the primary use cases of recommender systems, the performance of the system may be evaluated by the click-through rate or purchase rate of the recom-

mendations made to customers. But simple metrics as such are actually less benign than what they may seem. For example, what should those rates be compared to? If they are compared against the same metrics for other versions of the same system, at best it probably only leads to local optimization.

In other use cases, evaluation is trickier. For example, the recommender systems of online news sites typically recommend to readers new articles that align with the reader's past reading history. But this assumes that preferences are fixed, and that the user wants to read the same type of things they read in the past. Though useful in the short term, this approach to recommendations defeats an important purpose of reading— to expand one's knowledge base and increase one's understanding of the wider world. In this case, metrics like click-through rates are even less capable of capturing the overall usefulness of the recommendations. And the proliferation of this type of recommender systems has contributed to the growing "echo chamber" phenomenon in online media that increases social political polarization.

There is a similar challenge in evaluating recommender systems for export diversification. On the one hand, a successful system should be able to correctly predict a country's current "taste", i.e. the profile of the country's current comparative advantages. On the other hand, it needs to be forward looking enough to guide the development choices for the future. In other words, a recommender system that accurately describes the country's current export structure and simply recommends more of the same things is perhaps not very useful. Ultimately, a good recommender system should be both practical and aspirational— it should be able to recommend both product categories that fit right in with a country's current comparative advantages, and also products that are "stretch goals", i.e. those that would support diversified and sustainable growth for the future.

To test how well the different recommendation algorithms fulfill these goals, the following section applies the algorithms described in section 4 to the historical data of several high-growth countries that have successfully diversified their export base while achieving growth

acceleration. In line with the considerations above, the paper focuses on two evaluation criteria:

1. how well the recommendations capture the country's export structure at present, and
2. how well the recommendations predict the long-term structural changes of the country's export basket.

For the second criteria, the paper will compare the recommendations generated in a given year with the country's actual export composition in 20 years down the road. The implicit assumption here is that since these countries have achieved significant growth and income convergence during the period of increasing diversification, that gives legitimacy to the direction that their diversification had taken, i.e. it was a real-life approximation of an *optimized* diversification strategy. This is a reasonable assumption, given that economic development is not a repeatable scientific experiment and there is no real data on what an optimally growth-enhancing diversification path for a country should look like. This is also why the paper picked the high growth countries that have diversified their exports as test cases. Comparing the model recommendations with how these countries had actually diversified can give us useful insights on the algorithms' performance.

5.1 Test Cases

5.1.1 China

China had 242 high-RCA exports in 1995. By 2015, the number increased to 302. The recommendation algorithms are run using the SITC-4-digit data for all countries in 1995. Alternatively, the paper also tried including several years of data before 1995, but the results did not seem to diverge much. So a single year's data was used in the presentation of results.

Table 1 shows the number of recommendations generated by each model, their *hit rate*—the number of recommendations that were actual high-RCA exports of China in 1995, divided

by the total number of recommendations, and the *top-N hit rate*— the number of top-N recommendations that were actual high-RCA exports⁹

Two things should be noted. First, the number of recommendations vary widely by model. This is partly due to the fact that in all three models, the final output is an array of scores for each item in the SITC 4-digit space. The score is a monotonically increasing function of the predicted RCA score of the products, but not exactly the predicted RCA scores themselves. This would be no problem for the typical use case of these algorithms, i.e. to produce "top-N" recommendations, since only the ranking of the scores matters. In producing the metric for the number of high-RCA recommendations, though, a cutoff score has to be chosen. In the reporting, the paper picked a cutoff score of 1, across all models. So a relatively higher (lower) number of items will be included in the recommended list, if the predicted scores from the underlining model is a scaled up (scaled down) function of the RCA-score space. This would affect the calculation of hit rate as well, as the hit rate is likely to be lower, the larger the number of products included in the recommended list. The top-N hit rate circumvents this bias, as the hit rate is calculated based on the 100 recommendations that have the highest scores.

The product-based KNN appears to be the worst performer when it comes to correctly predicting China's existing export profile in 1995, with a top-N hit rate of only 17 percent. On the other end of the spectrum, the recommendations from the SVD algorithm appear to overlap almost perfectly with China's high-RCA export basket in 1995. Yet only 2 percent of the recommendations are *new*, i.e. not already among the existing high-RCA exports, which seems to defeat the purpose of a recommender system.

⁹In the baseline version of the models, $N = 100$.

	Product-based KNN	Country-based KNN	SVD
Number of recommendations	322	195	111
Hit rate (in percent)	47.2	67.7	98.2
Top-N Hit rate (in percent)	17.1	41.0	90.1

Table 1: China: Stats of Recommendation Models

But as discussed in the previous section, the performance of a recommender system cannot be solely judged by how well it fits a country's existing export structure. Arguably, it's more important for the system to generate insights on the future direction of export development. To assess how well the algorithms may guide the evolution of export diversification, the paper compares the recommended export categories against the actual exports of China, 20 years after the recommendations.

Figures 6 to 8 present the actual exports of 1995, the recommended exports from each algorithm, and the actual exports of 2015, by SITC 1-digit categories. Although China was well on its way towards industrialization in 1995, the export structure continued shifting in the subsequent two decades as the country diversified its exports. As can be observed from the figures, the general direction of the change is towards less primary material exports, and increased sophistication in manufacturing.

The product-based KNN emulates these trends very well. In fact, export-share-wise, the recommended export basket from this algorithm is closer to China's export basket in 2015, for 9 out of 10 SITC 1-digit categories, than to the export basket in 1995. And the only category where the recommendations diverged from the actual data in 2015 is category 4, which China exports relatively little of. In other words, the algorithm is able to correctly predicts the direction and structural change of export diversification on a macro level, for a majority of the product categories.

Figure 6: China: Result from Product-based KNN

In contrast, the country-based KNN and SVD algorithms, though both had higher hit rates than the product-based KNN, appear much less forward looking. Both are closer aligned with China's export basket in 1995 than with the one in 2015. The country-based KNN correctly predicted the direction of export structural change for 5 out of the 10 1-digit categories. And the SVD only correctly predicted 4 categories.

Figure 7: China: Result from Country-based KNN

Figure 8: China: Result from SVD

5.1.2 India

The number of high-RCA exports of India increased from 160 in 1985 to 256 in 2005. The three recommendation algorithms were applied to the SITC 4-digit data in 1985 to produce export recommendations for India. The results are reported in Table 2 and Figures 9 to 11.

Similar to the results for China, Table 2 shows that the product-based KNN has the lowest hit rate, despite giving the largest number of recommendations among the three algorithms. The SVD algorithm is the best performing in terms of recommending products closely aligned with the country's existing export portfolio. The country-based KNN falls somewhere in between.

	Product-based KNN	Country-based KNN	SVD
Number of recommendations	363	133	102
Hit rate (in percent)	33.6	61.7	78.4
Top-N Hit rate (in percent)	9.1	51.9	78.4

Table 2: India: Stats of Recommendation Models

Despite a low hit rate, the product-based KNN again shows impressive foresight, correctly predicting the structural changes of the export basket for 9 out of the 10 SITC 1-digit

industries. The SVD and the country-based KNN again seem to fall short in providing evolutionary guidance for export diversification, with their recommendations aligning with the actual direction of export structural change in only 5 and 4 categories, respectively.

Figure 9: India: Result from Product-based KNN

Figure 10: India: Result from Country-based KNN

Figure 11: India: Result from SVD

5.1.3 Chile

Chile's economic takeoff since the mid 1980s is a much lauded success story rare in Latin America. From the mid 80s to 2000, per capita real GDP growth averaged over 4 percent, compared to less than 1 percent of Latin American average over the period. The economy also significantly diversified during this time, with the number of high-RCA exports increasing by 45 percent from 1980 to 2000.

Table 3 and Figures 12 to 14 present the summary results from the three recommendation models using data from 1980. The first thing to note is that the hit rate is lower across the board compared to the previous two test cases, especially for the product-based KNN and the SVD. Notably, both algorithms significantly under-recommend exports in the commodity categories, while recommending more industrial exports. In contrast, the recommendations from country-based KNN are closer aligned with Chile's existing export portfolio in 1980. As in the case with previous two test countries, the product-based KNN produced recommendations that are close to Chile's export structure in 2000 than to that in 1980.

Why are the hit rates lower than in the other test cases? One reason can be that the existing number of high-RCA exports in 1980 was small, 75 in total, which was quite

diversified compared to other Latin American countries, but not so much when compared to other countries around the world at a similar development stage. Therefore the algorithms had less *source materials* to work from. But as will be demonstrated in the next section's case study, the algorithms are perfectly able to produce a higher hit rate with even less diversified country data. A more likely possibility is that the algorithms indeed think Chile should be much less reliant on commodity exports, given the country's fundamentals, which is in line with many human experts' recommendations to commodity exporting countries in Latin America¹⁰.

	Product-based KNN	Country-based KNN	SVD
Number of recommendations	139	88	534
Hit rate (in percent)	10.1	42.0	8.4
Top-N Hit rate (in percent)	5.8	42.0	5.1

Table 3: Chile: Stats of Recommendation Models

Figure 12: Chile: Result from Product-based KNN

¹⁰See, for example, Bacha and Fishlow (2011).

Figure 13: Chile: Result from Country-based KNN

Figure 14: Chile: Result from SVD

5.1.4 Poland

For Poland, the recommender algorithms are applied to data in 1985, a few years before the economic liberalization of the country. Poland already had a relatively diversified export base in 1985. Still, from 1985 to 2005, the number of high-RCA exports increased by 50 percent, from 187 to 287. As in most other test cases, the SVD algorithm yielded recommendations that are most closely aligned with the country's existing export profile, as shown in Table

4. On the other hand, the product-based KNN had the lowest hit rate, but offered fairly forward-looking recommendations, correctly predicting the directional change in Poland's export portfolio, for 7 out of 10 SITC 1-digit product categories.

	Product-based KNN	Country-based KNN	SVD
Number of recommendations	261	200	125
Hit rate (in percent)	45.2	59.0	67.2
Top-N Hit rate (in percent)	21.1	34.0	61.6

Table 4: Poland: Stats of Recommendation Models

Figure 15: Poland: Result from Product-based KNN

Figure 16: Poland: Result from Country-based KNN

Figure 17: Poland: Result from SVD

5.2 Discussion

One of the frequently voiced complaints from economists and policy makers regarding the use of machine learning algorithms in empirical studies is the seemingly opaque nature of the algorithms. The human cognitive system can differentiate a picture of a dog from that of a cat easily. But how that discernment is made is such a complex process, that it's difficult to explain with any linearly sequenced, logical narrative. It's something more akin to intuition

than to logic. Many machine learning algorithms share the very same characteristics. Predictions and judgments are made. But a clean, cause-and-effect explanation of the predictions is hard to come by, which tends to throw economic researchers and policy makers out of our comfort zone. Although we rarely ask why a picture of a dog is not that of a cat (we take that as self explanatory, though in truth it's anything but), we don't feel comfortable with a policy recommendation until we're presented with a set of logical and convincing arguments as to why the recommendation is made. Yet answering why is not a strong suit of machine learning algorithms. Even the simplest models such as the KNNs give results without easy explanations of how they come about.

For example, why is the product-based KNN able to predict the evolution of export structure in the test cases much better than the other two methods? The algorithms themselves do not provide a clear answer on this. And we can only speculate by looking at the nature of the data. The country-based KNN– and to an extent, the SVD– looks for contemporaneous similarities among countries when making recommendations. This strategy runs the risk of pigeonholing a country with other countries at a similar income level or in the same region, which, though providing a good approximation for the country's current fundamentals, offers little insight regarding future directions. For instance, in 1995, China looked similar to Haiti in terms of how rich they were and what they exported. But being itself a developing country (with lower GDP per capita), it's questionable how useful the Haiti data was to inform a forward-looking diversification strategy for China. In contrast, the product-based KNN looks for products similar to those that a country is already good at exporting. This leaves more room for identifying promising products both up and down the value chain, and for discovering hidden relations among products that are not defined by geography or development stage. Therefore the strategy is better at connecting the dots between today's fundamental and tomorrow's potential, to produce recommendations that are evolutionary. Another possibility is that there are more products than countries. So an algorithm that primarily focuses on similarity/relatedness among products is able to capture more nuanced

information than an algorithm that looks for similarities among countries.

And as far as the test cases are concerned, although the product-based KNN appears to be good at producing forward-looking recommendations about how a country's export structure should evolve, the recommendations do not come with their economic rationales clearly laid out. In the case of Chile, for example, both product-based KNN and SVD appear to *dislike* the country's heavy reliance on primary material exports, and recommendations are concentrated in other sectors. But there is no explanation as to why. As Section 6 will show, this is likely not because the algorithms are somehow systemically biased against commodities and primary materials. And although industrialization does seem to be a common rite of passage of almost all countries that have achieved fast income convergence in modern history, the recommendations do appear to be highly country-specific. They just don't come with neat economic theories attached, which, again, can be a major obstacle towards wider adoptions of machine learning generated policy recommendations.

The rationales behind the recommendations aside, a more practical question is how these models should be best used in analyzing a country's diversification strategy. The test cases show that the product-based KNN is the apparent best among the three algorithms in providing guidance on the structural change of an export portfolio. But useful as it is, structural change is a long-term process. Diversifying into higher-value added categories that support long-term growth does not have to be in conflict with diversifying into the categories that fit more into the country's current comparative advantages. The latter is often easier in the short-to-medium term, and can still provide benefits in, for example, reducing growth volatility. So in analyzing the recommendation results, the paper suggests to give more weights to the product-based KNN recommendations, while still referencing the results from other algorithms.

With these notes in mind, the next section will use Paraguay as a case study to demonstrate the analysis of recommendation results in more details.

6 A Contemporary Case Study: Paraguay

Paraguay experienced a major economic boom in the first decade and a half of the 21st century. Real GDP per capita grew at an annual rate of 3 percent on average, compared to the 0.8 percent for the Latin American & Caribbean region. The agricultural commodity price boom during the 2000s was a major cause of the growth acceleration in Paraguay. The price boom not only boosted income growth in the short term, but also stimulated agricultural investment and increased production concentration in select crops that had seen large price increases, such as soybeans.

As a result, export diversity has declined during the boom period. From 2000 to 2017, the number of high RCA exports dropped by over 35 percent, from 83 to 53. In 2017, close to 35 percent of goods exports from Paraguay are soybean related. And another 25 percent of total exports consists of electrical energy primarily produced by the two large hydroelectric dams, respectively co-owned with Brazil and Argentina. The electricity exports, though a major source of fiscal revenue, has little connection with the rest of domestic production sectors. As Figure 19 shows, compared to countries of a similar size around the world¹¹, Paraguay is one of the least diversified economies.

¹¹Defined as countries with a population size in the range of +/- 10 percent of Paraguay's population in 2017.

Figure 18: Paraguay: Evolution of High-RCA Exports

Figure 19: Number of High-RCA Exports by Country

The paper ran the recommendation algorithms on the SITC 4-digit export data for 2016 and 2017 to generate diversification recommendations for Paraguay. Table 5 shows the

number of recommended high RCA exports and the hit rates from each algorithm.

	Product-based KNN	Country-based KNN	SVD
Number of recommendations	105	164	214
Hit rate (in percent)	34.3	22.6	18.2
Top-N Hit rate (in percent)	34.3	20.1	7.0

Table 5: Paraguay: Stats of Recommendation Models

Table 6 and Figure 20 compare the existing sectoral structure of high RCA exports in 2017 with the recommended sectoral structure of high RCA exports from the models, by SITC 1-digit category. Table 6 shows the numbers of 4-digit high RCA exports contained in each 1-digit category, for actual exports and for recommended exports. The shares are the numbers of 4-digit products in each category as percent of the total count of high RCA exports. The orange column in Figure 20 are weighted averages across the three models.¹²

Consistent with the observations at the beginning of the section, Columns 2 and 3 of Table 6 shows that Paraguay's exports are heavily concentrated in agricultural commodities (row 1) and primary materials also related to agriculture (row 3). The industrial exports are mostly concentrated in chemical products (row 6) and basic manufacturing mostly concerning the processing of raw materials (row 7).

The recommender algorithms, somewhat surprisingly, concur with the current structure overall. For example, the product-based KNN, the most forward-looking algorithm among the three according to the results of the test cases, suggests that the dominance of food and crude materials in the export basket should be maintained, with diversification focusing on expanding into more products within these categories. At the same time, the algorithm suggests moderate expansion in the machinery & transport equipment category, which would reduce the relative importance of some other sectors, such as energy, in the export basket.

The other two algorithms agree in the ballpark with the recommendations from the product-based KNN. The weighted average recommendations from the three algorithms can

¹²The weights are .6, .2, and .2, for product-based KNN, country-based KNN, and SVD respectively.

be broadly described as follows: an export diversification strategy for Paraguay should focus on expanding varieties within the categories of agricultural commodities and primary materials, which the country is already very strong in. Meanwhile, the country should create conditions to moderately expand on light manufacturing exports that involve the processing of primary materials, and more aggressively expand on the manufacturing of machinery and equipment. The growth of the latter categories will serve to reduce the country's reliance on electricity exports.

Categories	Actual high-RCA exports 2017		Product-based KNN		Country-based KNN		SVD	
	Number	Share (in %)	Number	Share (in %)	Number	Share (in %)	Number	Share (in %)
Food & live animals	16	30.2	32	30.5	56	34.1	28	13.1
Beverages & tobacco	3	5.7	1	1.0	8	4.9	2	0.9
Crude materials, except fuels	11	20.8	24	22.9	32	19.5	56	26.2
Mineral fuels & related	1	1.9	1	1.0	1	0.6	6	2.8
Animal and vegetable oils, fats & waxes	4	7.5	8	7.6	9	5.5	12	5.6
Chemical products	7	13.2	14	13.3	13	7.9	11	5.1
Manufactured goods	8	15.1	16	15.2	28	17.1	52	24.3
Machinery & transport equipment	1	1.9	5	4.8	8	4.9	29	13.6
Miscellaneous manufactured articles	2	3.8	3	2.9	9	5.5	18	8.4
Commodities and transactions, n.e.s.		0.0	1	1.0		0.0		0.0
Total	53		105		164		214	

Table 6: Paraguay: Actual vs Recommended Export Structure

Figure 20: Paraguay: Actual vs Recommended Export Shares by Industry

Tables 7 and 8 compare the recommended export structure from the product-based KNN and the actual exports, and list the top 10 product categories with the largest increases and decreases in overall export shares, by SITC 2-digit category. Examining Table 7, it is interesting to note that the top recommendations cover multiple broader sectors. Paraguay currently exports little in most of these categories, even though some of them belong to the commodity and primary material sectors that Paraguay specializes in. On the other hand, of a couple categories in this list Paraguay already exports a lot, such as 01: meat and preparation, but the recommender system suggests there is still significant room to increase the export varieties in these categories.

For the top categories with reduced export shares in the recommendations (Table 8), it's important to note that with a couple of exceptions, the system is not suggesting that exports from these categories should be reduced. But rather, more diversification should happen in other categories that lead to a reduction of export shares for the categories in the list.

Product code	Product name	Share (actual)	Share (rec.)	Difference (rec. - actual)
72	Machinery specialized for <u>particular industries</u>	0.0	3.8	3.8
02	Dairy products and birds' eggs	0.0	2.9	2.9
21	Hides	0.0	2.9	2.9
26	Textile fibers and their wastes	0.0	2.9	2.9
58	Artificial resins and plastic materials	0.0	2.9	2.9
01	Meat and preparations	3.8	5.7	1.9
28	Metalliferous ores and metal scrap	3.8	5.7	1.9
09	Miscellaneous edible products	0.0	1.9	1.9
53	Dyeing	0.0	1.9	1.9
66	Non-metallic mineral manufactures	0.0	1.9	1.9

Table 7: Paraguay: Top 10 Categories with Increased Product Shares

Product code	Product name	Share (actual)	Share (rec.)	Difference (rec. - actual)
06	Sugar	5.7	1.0	-4.7
12	Tobacco and tobacco manufactures	5.7	1.0	-4.7
61	Leather	5.7	2.9	-2.8
59	Chemical materials and products	7.5	4.8	-2.8
27	Crude fertilizer and crude minerals	1.9	0.0	-1.9
77	Electric machinery	1.9	0.0	-1.9
29	Crude animal and vegetable materials	3.8	1.9	-1.9
55	Oils and perfume materials	3.8	1.9	-1.9
08	Feeding stuff for animals	5.7	3.8	-1.9
22	Oil seeds and oleaginous fruit	7.5	5.7	-1.8

Table 8: Paraguay: Top 10 Categories with Reduced Product Shares

7 Conclusion

A strategic vision for export diversification that aligns with a country's comparative advantages and development aspirations can usefully inform the design of growth-enabling policies. Yet there is no overarching empirical framework in economic research for creating such a vi-

sion that takes into account the diverse and often complicated fundamentals of individual countries. The paper aims to fill this gap by designing a machine-learning based recommendation system for export diversification, taking inspiration from the recommender systems based on collaborative filtering, popularized by e-commerce and online media platforms.

The basic idea of the system is to recommend new export categories to a country by looking at other countries with similar export profiles and products related to the country's existing exports. The implantation of the system is done via two neighborhood algorithms (country and product based) and a matrix factorization algorithm (singular vector decomposition). The algorithms are run on the historical data of several test-case countries that have achieved significant income convergence in the past few decades. And the recommendations are compared with the actual evolution of the export basket of these countries in subsequent years. Overall, the product-based KNN is able to recommend exports that are consistent with the subsequent structural changes in a country's export portfolio for a majority of the test cases. The algorithms are then applied to the current export data of Paraguay to demonstrate an analysis of export diversification strategy for the country.

Building the framework of the present paper, future research extensions could apply the algorithms to a wider set of countries and time periods, to explore the commonalities and differences of recommendations for countries that may share certain characteristics. In addition, the merit of the recommender system could be further tested, by applying it to a wide range of countries, identifying the divergence between recommended and actual export structures, and correlating the divergence with manifested economic outcomes.

Wider adoption of machine learning methodologies in international trade studies likely faces two challenges, which also apply to the adoption of machine learning tools in empirical economic research in general. First, improving the performance of the algorithms requires large quantities of valid *learning data*. In the context of export diversification, that means historical data on diversification episodes that have worked (failed), i.e. generated positive (negative) economic outcomes. Such data is scarce in both quantity and quality. Secondly,

unlike most conventional tools in empirical economic research, results from machine learning algorithms often seem *magical*— they are generated without tractable economic logic backing them up, which creates an obstacle for the evaluation and policy adoption of the results, especially given the first challenge. Further applications of machine learning in international trade studies, as well as in other economic research topics, will need to address these challenges.

References

- Adeniyi, David Adedayo, Zhaoqiang Wei, and Y Yongquan. “Automated web usage data mining and recommendation system using K-Nearest Neighbor (KNN) classification method”. In: *Applied Computing and Informatics* 12.1 (2016), pp. 90–108.
- Aiginger, Karl and Dani Rodrik. “Rebirth of Industrial Policy and an Agenda for the Twenty-First Century”. In: *Journal of Industry, Competition and Trade* (), pp. 1–19.
- Bacha, Edmar L and Albert Fishlow. “The recent commodity price boom and Latin American growth: More than new bottles for an old wine?”. In: *The Oxford Handbook of Latin American Economics*. 2011.
- Balaguer, Jacint and Manuel Cantavella-Jorda. “Structural change in exports and economic growth: cointegration and causality analysis for Spain (1961–2000)”. In: *Applied Economics* 36.5 (2004), pp. 473–477.
- Cadot, Olivier, Céline Carrère, and Vanessa Strauss-Kahn. “Export diversification: what’s behind the hump?”. In: *Review of Economics and Statistics* 93.2 (2011), pp. 590–605.
- Clickstreams, Multi-faceted Web. “Workshop Notes”. In: (2005).
- Feenstra, Robert and Hiau Looi Kee. “Export variety and country productivity: Estimating the monopolistic competition model with endogenous productivity”. In: *Journal of international Economics* 74.2 (2008), pp. 500–518.
- Frankel, Jeffrey A. *The natural resource curse: a survey*. Tech. rep. National Bureau of Economic Research, 2010.
- Giri, Rahul, Mr Saad N Quayyum, and Rujun Yin. *Understanding Export Diversification: Key Drivers and Policy Implications*. International Monetary Fund, 2019.
- Hausmann, Ricardo, Jason Hwang, and Dani Rodrik. “What you export matters”. In: *Journal of economic growth* 12.1 (2007), pp. 1–25.
- Hausmann, Ricardo and Bailey Klinger. “The structure of the product space and the evolution of comparative advantage”. In: *CID Working Paper Series* (2007).
- Herzer, Dierk and Felicitas Nowak-Lehmann D. “What does export diversification do for growth? An econometric analysis”. In: *Applied economics* 38.15 (2006), pp. 1825–1838.
- Hidalgo, César A and Ricardo Hausmann. “The building blocks of economic complexity”. In: *Proceedings of the national academy of sciences* 106.26 (2009), pp. 10570–10575.
- Imbs, Jean and Romain Wacziarg. “Stages of diversification”. In: *American Economic Review* 93.1 (2003), pp. 63–86.
- IMF. “Sustaining Long-Run Growth and Macroeconomic Stability in Low-Income Countries—The Role of Structural Transformation and Diversification”. In: *IMF Policy Paper* (2014).
- Klinger, Bailey and Daniel Lederman. *Discovery and development: An empirical exploration of “new” products*. The World Bank, 2004.
- “Export discoveries, diversification and barriers to entry”. In: *Economic Systems* 35.1 (2011), pp. 64–83.
- Koren, Yehuda, Robert Bell, and Chris Volinsky. “Matrix factorization techniques for recommender systems”. In: *Computer* 42.8 (2009), pp. 30–37.
- Lathia, Neal, Stephen Hailes, and Licia Capra. “kNN CF: a temporal social network”. In: *Proceedings of the 2008 ACM conference on Recommender systems*. 2008, pp. 227–234.
- Lin, Justin Yifu and Feiyue Li. *Development strategy, viability, and economic distortions in developing countries*. The World Bank, 2009.

- Al-Marhubi, Fahim. “Export diversification and growth: an empirical investigation”. In: *Applied economics letters* 7.9 (2000), pp. 559–562.
- Paterek, Arkadiusz. “Improving regularized singular value decomposition for collaborative filtering”. In: *Proceedings of KDD cup and workshop*. Vol. 2007. 2007, pp. 5–8.
- Prebisch, Raul. “The economic development of Latin America and its principal problems”. In: *Economic Bulletin for Latin America* (1962).
- Sarwar, Badrul et al. *Application of dimensionality reduction in recommender system-a case study*. Tech. rep. Minnesota Univ Minneapolis Dept of Computer Science, 2000.
- “Incremental singular value decomposition algorithms for highly scalable recommender systems”. In: Citeseer.
- Singer, Hans W. “The distribution of gains between investing and borrowing countries”. In: *The Strategy of International Development*. Springer, 1975, pp. 43–57.
- Singer, HW. “The Distribution of Gains between Investing and Borrowing Countries”. In: *The American Economic Review* 40.2 (1950), pp. 473–485.