

ACCEPTED MANUSCRIPT • OPEN ACCESS

Which will be your firm's next technology? Comparison between machine learning and network-based algorithms

To cite this article before publication: MATTEO STRACCAMORE *et al* 2022 *J. Phys. Complex.* in press <https://doi.org/10.1088/2632-072X/ac768d>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2022 The Author(s). Published by IOP Publishing Ltd.

As the Version of Record of this article is going to be / has been published on a gold open access basis under a CC BY 3.0 licence, this Accepted Manuscript is available for reuse under a CC BY 3.0 licence immediately.

Everyone is permitted to use all or part of the original content in this article, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected and is not published on a gold open access basis under a CC BY licence, unless that is specifically stated in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Which will be your firm's next technology? Comparison between machine learning and network-based algorithms

Matteo Straccamore

MATTEO.STRACCAMORE@CREF.IT

Dipartimento di Fisica Università "Sapienza"

P.le A. Moro, 2, 00185 Rome, Italy

Centro Ricerche Enrico Fermi

Piazza del Viminale, 1, 00184 Rome, Italy

Luciano Pietronero

Centro Ricerche Enrico Fermi

Piazza del Viminale, 1, 00184 Rome, Italy

Andrea Zaccaria

Istituto dei Sistemi Complessi (ISC) - CNR

UoS Sapienza, P.le A. Moro, 2, 00185 Rome, Italy

Centro Ricerche Enrico Fermi

Piazza del Viminale, 1, 00184 Rome, Italy

Editor: //

Abstract

We reconstruct the innovation dynamics of about two hundred thousand companies by following their patenting activity for about ten years. We define the technological portfolios of these companies as the set of the technological sectors present in the patents they submit. By assuming that companies move more frequently towards related sectors, we leverage on their past activity to build network-based and machine learning algorithms to forecast the future submissions of patents in new sectors. We compare different prediction methodologies using suitable evaluation metrics, showing that tree-based machine learning algorithms outperform the standard methods based on networks of co-occurrences. This methodology can be applied by firms and policymakers to disentangle, given the present innovation activity, the feasible technological sectors from those that are out of reach.

Keywords: Economic Complexity, Technological Innovation, Predictions, Patenting firms

Introduction

In this work, we quantify the relatedness between a firm and a technology sector in different ways, namely using standard methods based on co-occurrences networks and supervised machine learning algorithms (Tacchella et al. (2021); Albora et al. (2021)). In order to compare such assessments, we develop an out-of-sample prediction framework based on the assumption that, on average, the next technology sector in which a firm will patent will be among the ones that are more related with its present patenting portfolio. In this way, we can build and study the *technological adjacent possible* of innovative firms, this concept being originally introduced by Kauffman (1996) and subsequently mathematically formalised in Tria et al. (2014); Loreto et al. (2016). We find that machine learning algorithms not

only show better prediction performances but allow for a two-dimensional representation of technology sectors that we call Continuous Technology Space (CTS). The CTS can be used to visualize the patenting portfolio of companies and to design strategic investments and acquisitions.

The question regarding the nature of the link between the performance of firms and their internal allocation of resources (Penrose, 1959) and capabilities (Teece et al., 1994) has fueled the interest of economics and management scholars for a long time, since opening the black box of corporate strategy would be key to gain insight into the determinants of corporate heterogeneity and hence a better understanding of markets and their evolution. To the best of our knowledge, these analyses are all aiming at finding explanatory variables for the present performance and not at forecasting future activity. On the contrary, the approach known as Economic Fitness and Complexity Tacchella et al. (2012); Sbardella et al. (2018), widely applied at both country and regional level, naturally focuses on forecasting, which represent a natural and scientifically sound framework to validate and falsify the different approaches (Tacchella et al., 2018; Albora et al., 2021; Tacchella et al., 2021). The aim of the present paper is to apply the EFC forecasting methods at firm level, and in particular to the bipartite network of firms and the technology sectors in which they show patenting activity.

One of the main problems for the economic literature is to empirically track the capabilities and the strategic choices of companies. Unfortunately, these elements are generally intangible, so that the empirical literature often struggles to find instruments to keep up with the theoretical richness of the debate. One of the more easily measurable footprints left behind by the strategic decision making of firms is *diversification*, i.e. the scope of activities (both at technological and productive level) to which internal resources are devoted. This has been recognized early by scholars, who have often focused their efforts in this direction to reconcile theory with empirical evidence (Penrose, 1960; Gort, 1962a; Berry, 1971b). Though diversification is interesting in and of itself, perhaps the more interesting question regards the degree of complementarity (or relatedness) between the various elements included in the portfolio of activities in which businesses engage. Notable early efforts to address this aspect have been proposed by Rumelt (1974) and Rumelt (1982). Both studies examine diversified manufacturing firms and focus on the link between profitability and the degree of correlation between the business units of the same firms. From this, they test the hypothesis that greater profitability correlates with expansion mainly in areas that share a competence or basic resource. Teece et al. (1994) have built on the above intuition by employing plant-level data classifying establishments according to the standard 4-digit SIC industrial codes relative to the industrial sectors in which they operate and measuring the relatedness between sectors through the frequency of their co-occurrence within the same productive plant, that is two sectors are related if many plants produce both. The hypothesis underlying this approach is the so-called *survivor principle* (Teece et al., 1994), i.e. the assumption that economic competition eventually drives inefficient organizational forms out of the market, thus promoting the co-occurrence of activities that are well integrated with one another because of complementarities in the *technological capabilities* they require. In virtue of the survivor principle, efficient combinations of activities should occur with a significantly higher frequency than one would expect if activities were paired randomly. Indeed, the authors find that internal coherence matters, as firms that diversify tend to add activities that are related to at least a part of their existing portfolio. More recent analyses confirmed this hypothesis (Rahmati et al., 2020; Buccellato, 2016; Lo Turco and Maggioni, 2016).

Production is not the only aspect of corporate strategy in which building a coherent portfolios of related activities has been shown to matter (for example in Gort (1962b); Rumelt (1974); Berry (1971a) the manufacturing sector is considered). Indeed, in the last twenty

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

years, the empirical analysis of the innovative output of firms as measured by patents has gained increasing popularity (Rycroft and Kash, 1999). It is worth noting that patent data have become in general a workhorse for the literature on technical change over the past few decades due to the growing availability of machine-readable patent documents and widespread access to sufficient computing power (Youn et al., 2015). All the above has played a pivotal role in fueling this trend spurring scholarly (e.g. Hall et al. (2001)), institutional (e.g. PATSTAT, REGPAT) and corporate (e.g. Google Patents) efforts aimed at constructing comprehensive collections of patent-related documents. Increasing data availability has in turn allowed researchers to inquire into the nature of patented inventions, their role in explaining technical change, their reciprocal connections, and their link to inventor - and applicant-specific characteristics (Strumsky et al., 2011, 2012; Youn et al., 2015). One of the characteristics of patent documents, which historically has lent itself more to economic analysis, is the presence of codes associated with the claims contained in the patent applications. These are used to mark the boundary of the commercial exclusion rights demanded by inventors. To allow evaluation by patent office examiners, claims are classified based on the technological areas they impact according to classifications (e.g. the IPC classification (Fall et al., 2003)), which consist of a hierarchy of 6-digit codes that associate progressively finer-grained definitions of technological areas to the codes lower in the hierarchy. Mapping claims to classification codes allows to localize patents and patent applications within the technology space. Taking advantage of the increasing availability of patent data, several studies (Jaffe et al., 2000; Leten et al., 2007; Joo and Kim, 2010; Rigby, 2015) have found significant empirical evidence suggesting that evidence that relatedness in the composition of R&D activities has implications for the ability of firms to innovate successfully.

Within this stream of literature, a well-known study (Breschi et al., 2003) has recovered the methodology proposed by Teece et al. (1994) and built upon it to investigate whether firms tend to diversify their innovative efforts in a coherent fashion by patenting in technological fields that share a common knowledge base with the technological fields in which they innovated in the past. In particular, the authors have analyzed the technological diversification of firms through the co-occurrences between technology codes.

In another well-known paper, Nesta and Saviotti (2006) have studied corporate knowledge coherence in the US pharmaceutical industry showing that both the scope and the coherence of the knowledge base “contribute positively and significantly to the firm’s innovative performance”, as measured by the number of patents it produces weighted by the number of citations received.

Some authors of the present paper introduced the concept of “coherent diversification” (Pugliese et al., 2019b), showing that firms that diversify (i.e., expand their technological portfolios by patenting in a relatively large number of technological sectors) in a coherent way (i.e., by preferring related sectors to unrelated ones) on average show a higher performance in terms of labor productivity.

In Yan and Luo (2017), the authors present a method for choosing an optimal compromise between explanatory power of on the diversification and the removal of the weak links in a network of technology codes.

Finally, we mention the work by Kim et al. (2021), whom have studied the relatedness between technology codes in Korean firms, finding that “firms are more likely to develop a new technology when they already have related technologies”.

Results

The data we will use in this study is the matrix representation of the temporal bipartite company-technology network. In particular, we will consider 643 technology sectors embedded in the patents submitted by 197944 firms in 12 years. In practice, we will use 12 \mathbf{V}^y matrices that link the layer of firms with that of technology codes, where y ranges from 2000 to 2011. In the following, we will interchangeably use the terms technological code, sector, or simply technology to express the same concept, since the codes written in the patents do represent technology sectors and so, in this sense, technologies.

The matrix element $V_{f,t}^y$ quantifies the patenting activity of firm f in the technology field t during year y . In particular, it is the number of patents submitted by the given firm in that sector. Note that this number can be fractional, since (usually) more than one code is present in each patent and (rarely) a single patent could be submitted by more than one applicant firm. In these cases, the unitary weight corresponding to one patent is split among the sectors and/or the applicants. Note that it may also happen that the same invention is linked to multiple patent application documents. In this case, each group of documents in the PATSTAT database is called "Patent Family" according to primary citations among them (Publishing et al., 2001). Referring to the same inventions, these families are associated to the same technology codes, and they are counted as single patents. In summary, each matrix element V_{ft}^y is obtained as follows: we assign to each patent (or family of patents), in a given year y , one unit of weight. This is then divided into equal shares between all the observed (firm f -technology t) pairs and, finally, the matrix \mathbf{V} is built by summing element-wise these contributions. The construction process is explained with more details and a numerical example in the supplementary information. The starting data is obtained by matching the *AMADEUS* database (<https://amadeus.bvdinfo.com>), that covers over 20 million firms with European registered offices, with the *Patstat* (www.epo.org/searching-for-patents/business/patstat) database about patents submissions. More details can be found in the Methods section and in Pugliese et al. (2019b).

The matrix element V_{ft}^y gives a quantification of the patenting activity of firm f in the technology t . However, in the EFC framework one usually deals with binary matrices; our choice is to use different thresholds T to define the 12 binary matrices \mathbf{M}^y , one for each year from 2000 to 2011, and to compare a posteriori the effect of using different values of T . In formulas, the binarizing procedure reads

$$\begin{aligned} \text{If } T = 0 &\rightarrow M_{ft}^y = \begin{cases} 1 & \text{if } V_{ft}^y > T \\ 0 & \text{if } V_{ft}^y = T \end{cases} \\ \text{If } T > 0 &\rightarrow M_{ft}^y = \begin{cases} 1 & \text{if } V_{ft}^y \geq T \\ 0 & \text{if } V_{ft}^y < T \end{cases} \end{aligned}$$

So the element $M_{f,t}^y$ is equal to 1 if a firm f submits more than T patents with technological code t in the year y , and 0 otherwise. We point out that in the Economic Complexity framework one usually binarizes the export (or, if patents are considered, the innovation activity) matrix using Balassa's Revealed Comparative Advantage (Balassa, 1965; Hidalgo et al., 2018; Pugliese et al., 2019a). Since in this case the \mathbf{V} matrix is very sparse, the effect of RCA is practically negligible so we preferred to use the V_{ft}^y elements for clearer interpretability.

These \mathbf{M} matrices can be used to train different algorithms to calculate our predictions about their temporal evolution. In order to have an out of sample forecast, we use data from 2000 to 2009 for the training phase and to obtain a score matrix \mathbf{S}^{2011} which will represent the relatedness between companies and technologies; in other words, we expect that a higher value of the matrix elements $S_{f,t}^{2011}$ is connected to a higher probability for

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

firm f to patent in technology code t in year 2011.

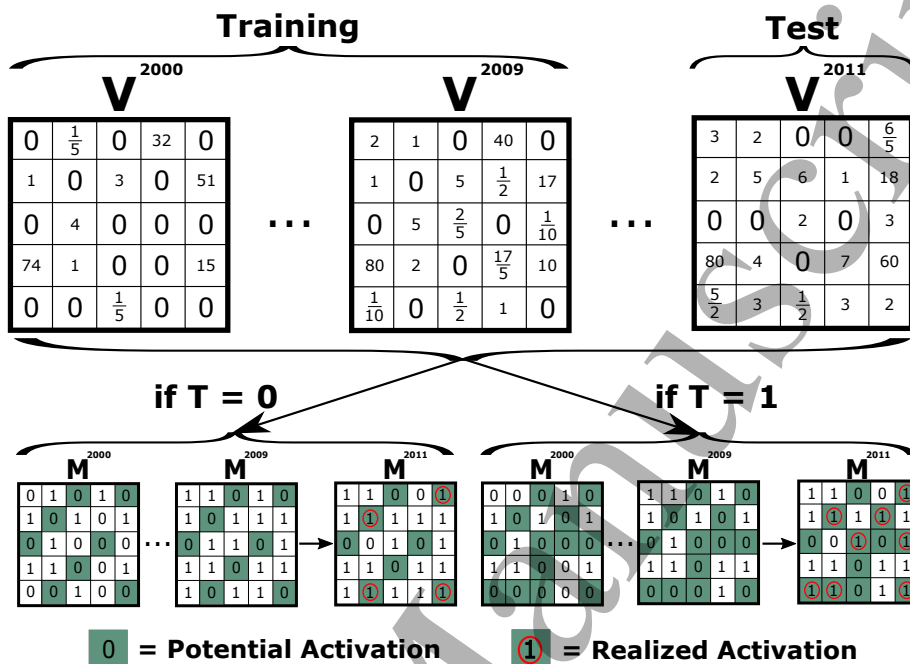


Figure 1: Schematic representation of the data processing. The \mathbf{V} matrices represent the yearly bipartite networks that link firms and technologies; each element $V_{f,t}$ represents how much a technology t is present in the patenting activity of firm f . The elements $M_{f,t}$ of \mathbf{M} matrices instead, gives us the binary information about *whether* a technology t is made by a firm f or not, that is if $V_{f,t}$ exceeds or not a threshold T . The dark green elements are the potential activations and the red circles in 2011 matrices are the realized activations (i.e. those elements that were 0 in all training years and then, in 2011, become 1: a new technology for that firm). Also the distinction between training and test set is shown.

We point out that both the matrices \mathbf{V} and \mathbf{M} are highly *autocorrelated* in time: if a firm does submit patents with a given technological code in a year y , it is likely that it will also in the year $y + \delta$, and viceversa. As a consequence, we focus our attention on those matrix elements that we call *potential activations*: the elements of \mathbf{M} that are 0 in all training years (from 2000 to 2009). Then, we will check whether in the test year (2011) this element remains equal to 0 or becomes 1. We will call this last case a *realized activation*: a firm enters (that is, starts patenting) in a technological sector which is new to this firm. In Figure 1 we represent a numerical example clarifying how we managed the \mathbf{V} and the \mathbf{M} matrices, the division of the data in training and test set, and the definitions of both potential and realized activations.

Our forecast exercise permits to compare different prediction algorithms using the test year 2011. So we will compute one score matrix \mathbf{S}^{2011} for each algorithm and we will compare it with \mathbf{M}^{2011} (obtained by binarizing the empirical \mathbf{V}^{2011}), and quantifying the prediction performance as in usual supervised classification tasks (Kotsiantis et al., 2007).

In order to obtain the prediction scores, we use different algorithms to evaluate the relatedness (Hidalgo et al., 2018) between a firm and a technology. In the case of co-occurrences based networks, an intermediate step is to assess the *similarity* between technology codes. Here we list the tested algorithms by category, leaving a more detailed discussion for the Methods section.

- **Benchmarks:** We use a quasi-trivial Random and Autocorrelation-based predictions as benchmarks. The first is a random model where we fix the diversification of the firms $d_f = \sum_t M_{f,t}$, i.e. the number of the technology codes in its patents. The second is a benchmark model that takes into account the temporal autocorrelation of the \mathbf{M} matrices: the scores \mathbf{S} are equal to the mean of $V_{f_t}^y$ in all the training years (i.e. with $y \in [2000, 2009]$).
- **Networks:** The standard Economic Complexity approach usually starts from the evaluation of normalized co-occurrences; in the simplest case

$$B_{t,t'}^y = \sum_f M_{f,t}^y M_{f,t'}^y$$

that is, t and t' are similar if many firms patent in both sectors. Different normalizations lead to the Product Space, or in this case, the Technology Space (Hidalgo et al., 2007), the Taxonomy Network (Zaccaria et al., 2014), and the Micro-Partial network, based on the paper of Teece et al. (1994). In all these cases, the network \mathbf{B} represents a projection of the bipartite network \mathbf{M} into the space of technology codes, and each element $B_{t,t'}$ represent the proximity/similarity between the two technology codes. In order to obtain a measure of the relatedness between a firm f and a target technology t , to be used as a prediction score, one then computes the coherence (Pugliese et al., 2019b) using eq. 2. Other approaches, such as the density normalization introduced by Hidalgo et al. (2007), perform sensibly worse. More details are provided in the Methods section.

- **Machine Learning:** Since our prediction exercise can be expressed as a supervised classification exercise, we can use the Random Forest (RF) algorithm (Breiman, 2001; Albora et al., 2021), and what we call the Continuous Technology Space (CTS). The first is a popular machine learning algorithm based on decision trees, while the CTS is based on the studies of Tacchella et al. (2021), and it is a projection on the space of the technology codes obtained from the scores obtained with the Random Forest, which can be seen as a high dimensional representation of the codes themselves. This is done by using a Variational Auto Encoder (Kingma and Welling, 2013) followed by the t-SNE dimensionality reduction algorithm (Van der Maaten and Hinton, 2008). In this way, we are able to make the results of the Random Forest, in a sense, more interpretable. As better specified in the Methods, the RF scores can be seen as an high dimensional representation of the technology codes (one dimension for each firm). In order to visualize this space, the t-SNE dimensionality reduction algorithm is applied, which results in the CTS. Note that in order to produce prediction scores from the CTS one has to compute a coherence or density measure as in the network based approaches. The use of Random Forest somehow hides the reason why a company is close to a technology code, in other words, where the relatedness results comes from. However, by applying t-SNE to the prediction scores one can obtain a visual representation of the relative position of the codes in this new space we define. Now the motivation behind our relatedness assessment, and the consequent forecast, becomes (hopefully) more interpretable: the company is close to a given technology if it is already patenting in close technologies. This is visible as a diffusion process only if a low-dimensional representation is adopted. We point out that this is not an explanation of how the Random Forest works, which is beyond the scope of the paper, but a posteriori justification of our results that, being represented in a 2-dimensional plane, can provide insights about companies' innovation strategy.

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

Two types of Random Forest are used, the non-Cross Validated (RF) and the Cross Validated one (RF.CV). With the cross validation, we remove a portion of firms at a time from the training, and then we use them in the test. The starting rationale is that the algorithm produces its predictions by using two pieces of information: the similarity between technologies and its ability to recognize a firm. By cross-validating the RF we try to force the algorithm to use the former, and not the latter (Albora et al., 2021).

Prediction results

Here we compare the relatedness assessments of the co-occurrences based networks with the machine learning algorithms, showing how the latter are able to give better prediction results. The results are shown in Figure 2. In order to compare the various prediction

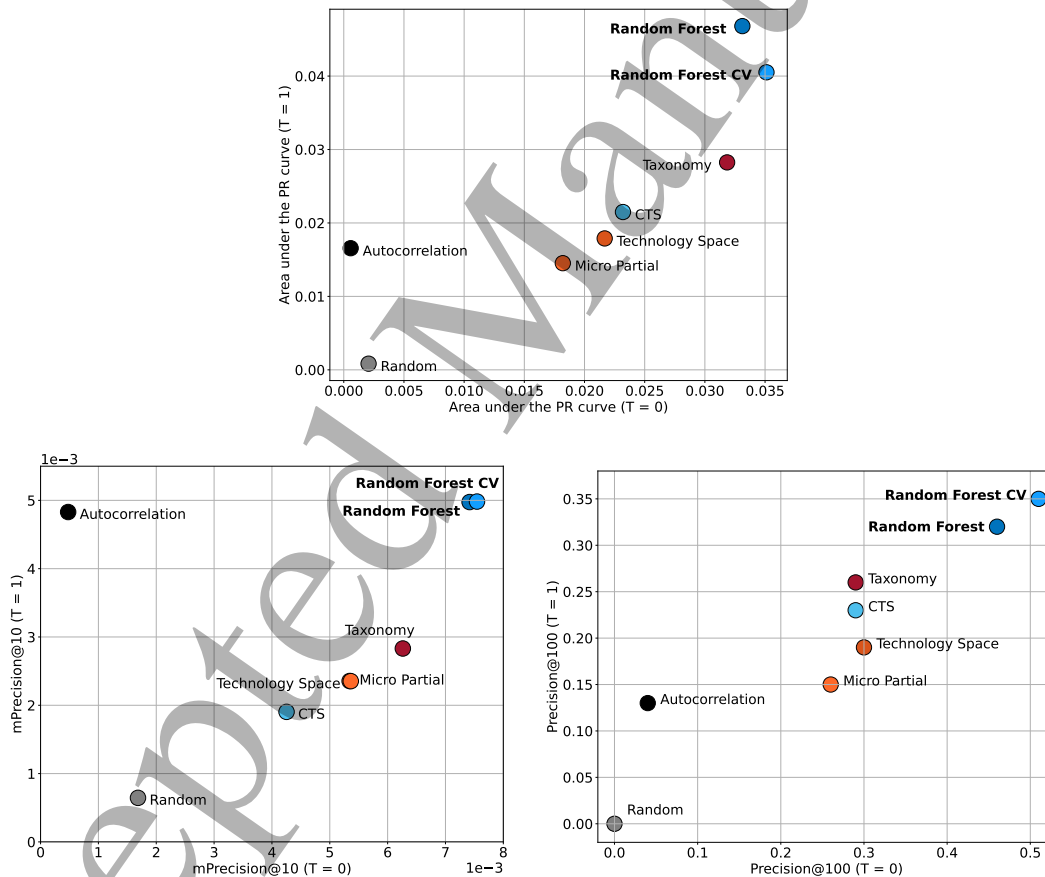


Figure 2: Comparison of the prediction results obtained by different approaches, using three different evaluation metrics and two binarization thresholds. Random Forest outperforms all other approaches in all perspectives. Regarding the Area under the PR curve, the second bests are the Taxonomy Network and the Continuous Technological Space.

methods from various viewpoints, we adopted different metrics to quantify the goodness of a prediction (these metrics are discussed in detail in Methods):

- **Area under the PR curve:** the area under the curve in the precision recall plane. The latter is obtained by varying the threshold that identifies the value above which the scores are associated to positive predictions;
- **Precision@100:** the fraction of the largest 100 elements of the score matrix S^{2011} that are actually activated;
- **mPrecision@10:** for each firm, we consider the largest 10 scores and we compute the fraction of realized activations; then we average over the firms.

In Fig.2 we report the scores of the previous metrics for different values of the threshold parameter T ; the results are consistent even if one varies such threshold (or uses the RCA to binarize).

We start noticing that the random benchmark is surpassed by all the different approaches, showing that all are able to compute a measure of similarity that is able to grasp links between the technology codes.

Also the autocorrelation benchmark is outperformed, but in the mPrecision@10 case. In particular, it performs better when T increases, because the number of zeros in both the training and the test matrices increases (that is, the number of potential activations, the green elements in Figure 1, that are not realized).

In the Area under the PR curve and Precision@100, the only network-based algorithm that manages to overcome the CTS is the Taxonomy. In particular, it is interesting to observe how this network exceeds the Technology Space. We can argue that for the technology codes, a network based on the taxonomy principle, i.e. how firms move from low-complexity to high-complexity technologies only after developing the necessary skills (Zaccaria et al. (2014)) shows a better prediction performance than a proximity-based one, i.e. a network where two technologies have a high link if they need the same capabilities (Hidalgo et al., 2007).

The Micro-Partial approach does not show a competitive performance despite being quite popular in both academic and corporate applications (Smith and Linden, 2017).

In any case, the superiority of the RF_CV and of the RF with respect to both benchmarks and density-based approaches (networks and CTS) is evident. Although the other algorithms are able to give prediction scores able to overcome the benchmark models (especially for $T = 0$), clearly these are not able to fully highlight the non-linear relationships among the technological portfolios of firms and the technological sectors they will move to.

In Figure 3 we compare the frequency distributions of the scores of both the realized and the not realized activations for the Random Forest (y-axis) and Technology Space (x-axis).

In order to make them comparable, both scores are rescaled using the respective maxima and minima. The red line is the bisector, showed for further reference. From the left figure it emerges that the Random Forest assigns, on average, higher scores to those potential activations which will be actually realized in two years. On the contrary, the possible but not realized activations show similar distributions; this is due to the much greater number of true negatives which is present in both approaches. Note that, as expected, the scores given to the Not Realized Activations are lower than the Realized ones.

Finally, it is also important to point out that in the present study, true positives are more significant than true negatives. This has a twofold rationale:

- The high class imbalance implies that a performance measure such as accuracy is not adequate for the problem. The majority of the elements of our matrices is equal to zero and therefore an accuracy measure would consider only the overwhelming number of true negatives. Even if we made a prediction in which we assume that all the elements of the matrix will be zeros

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

we would get an accuracy higher than 99%. More in detail, the percentage ratio between elements equal to 1 and 0 is about 0.2% for $T = 0$ and 0.08% for $T = 1$.

- For a firm, it is more interesting to know which technologies are close to it than those it already active into (i.e. which technologies it can successfully activate in the future), rather than knowing which ones it will not do in the future (which is often a trivial information, due to a totally different scope, for instance).

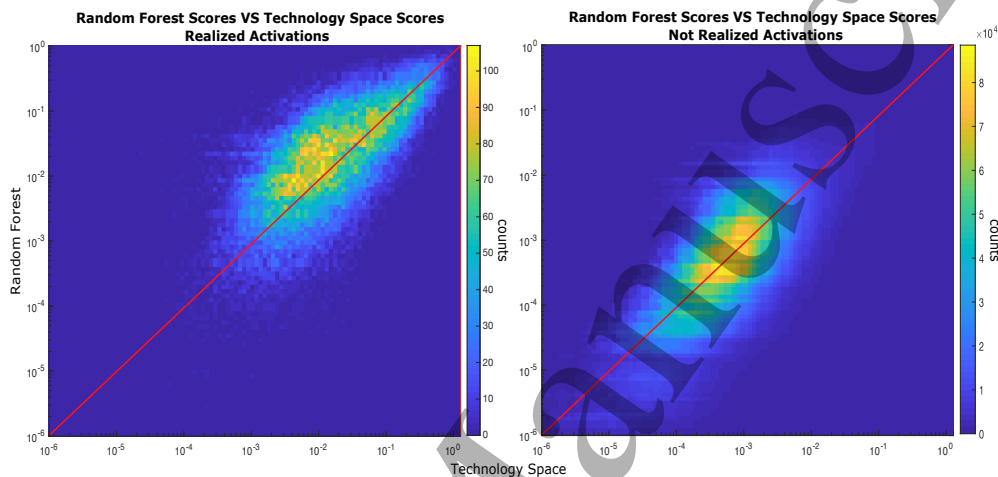


Figure 3: Comparison of the Random Forest and Technology Space scores with a 2D histogram plot. In each Figure the bins' color represents the number of data points within each bin. The left figure is referred to the Realized Activations, i.e. those elements that are always 0 during the training years (from 2000 to 2009) and then become 1 in the test year (2011). The red line is the bisector. The Realized Activations' scores obtained with the Random Forest are, on average, higher than those obtained with the Technology Space. The right figure is referred to the Not Realized Activations, i.e. elements that also in the 2011 are zero: here, the distributions are roughly similar.

Continuous Technology Space

Even if the prediction performance of the Random Forest algorithm vastly outperforms the other approaches, its practical feasibility in policy making could be limited by its low interpretability. From a policy perspective, indeed, it is not easy to justify a strategic decision such as to invest or not in a technological sector on the basis of a quasi-black box algorithm. In order to provide a visual tool to inform and justify strategic decisions, we introduced the Continuous Projection Space (Tacchella et al., 2021), that uses the scores obtained from the machine learning algorithms to build a two-dimensional and, as such, easy interpretable space to visualize and describe the temporal evolution of bipartite networks. The key idea is to interpret the scores matrix obtained with the Random Forest as a matrix of coordinates of technology codes in a high-dimensional space. These embeddings are then made representable by applying suitable dimensionality reduction techniques; in this case, t-SNE Van der Maaten and Hinton (2008). Note that here we are using the term "interpretable"

1
2
3
4
5 in a policy perspective: in this sense, machine learning algorithms are a black box in the
6 sense that, for a policymaker, the origin of our results is not clear a priori (for instance, a
7 company being close to a technology code). However, by applying t-SNE to the Random
8 Forest scores one can obtain a visual representation of the relative position of the codes in
9 this new space we define. Now the motivation behind our relatedness assessment, and the
10 consequent forecast, become (hopefully) clear: the company is close to a given technology
11 if it is patenting in the close technologies. This is visible as a diffusion process only in a
12 low-dimensional representation. We point out that this is not an explanation of how the
13 RF works, which is beyond the scope of the paper, it is just an a posteriori justification of
14 our results that, being representable in a 2-dimensional plane, can provide insights about
15 companies' innovation strategy (e.g., exploration vs exploitation).
16 Here we apply this methodology - which is fully described in the Methods section - to the
17 firm-technology network; the result is a plane in which each technology sector is a point,
18 and similar sectors are close. In Figure 4a) the different colors correspond to IPC macro-
19 categories, i.e. the first of the 4 digits that define the classification codes. We point out
20 that, differently from network-based representations, here the similarities are simply repre-
21 sented by the spatial proximity between technology codes. The use of euclidean distances
22 instead of topological ones permits to use a wider range of tools, for instance clustering
23 and anomaly detection algorithms. A visual inspection of the CTS permits to obtain a
24 number of insights: in the Figure 4a) one can observe that technology codes tend to cluster
25 following the macro categories; this is a first hint that the positions in the plane present a
26 certain degree of significance. However, also the departures from the classification reveal
27 meaningful relationships. In particular, on the upper left one can observe the presence of
28 a dense area where it is possible to find veterinary medicine close to farm. In the Motor
29 vehicles area on the left we find motor vehicle technology codes; in particular, there is a
30 red color technology code (A47C) that corresponds to chairs and seats specially adapted
31 for vehicles, black technology codes colors, corresponding to B60 (considering the first 3
32 digits), that represent vehicles, light blue technology codes corresponding to the first 3
33 digits F01 and F02, i.e. machines and engines, and combustion engineering, and one brown
34 technology code color (G05G), physics of command systems. Weapons area is associated
35 with weapons technologies: we find principally (considering the first 3 digits) codes B63
36 and B64, i.e. ships and aircrafts, C06 associated to explosive chemistry, and F41 and F42,
37 i.e. weapons and ammunition.
38

39 The example discussed above can be generalized by comparing two frequency distri-
40 butions of distances (see Figure 4b). The first distribution (cyan line) is relative to the
41 distances between the new activated codes and the patenting company (that is, the dis-
42 tance from the closest sector the company patented in); the second (orange line) is the
43 density distribution of all distances between the points in the CTS. We can see how the
44 first distribution has a lower mean, evidencing that, on average, firms tend to patent in
45 codes that are relatively close to what they already do. Obviously, also high distances are
46 present, indicating strategic choices which lead the firm far from its usual scope.
47 In order to show a concrete application of the CTS, we show in Figure 5 the portion of
48 this space relative to an American nanotechnology company, Nanotek Instruments Inc., as
49 an example. In 2002 Nanotek patented three inventions, two based on batteries ([https://
50 patentimages.storage.googleapis.com/f4/d8/3d/d663e43fe48e2b/US6773842.pdf](https://patentimages.storage.googleapis.com/f4/d8/3d/d663e43fe48e2b/US6773842.pdf) and
51 [https://patentimages.storage.googleapis.com/66/b3/7f/6fa873ae402fbf/US6864018.
52 pdf](https://patentimages.storage.googleapis.com/66/b3/7f/6fa873ae402fbf/US6864018.pdf)) and the third is the *Nano-scaled graphene plates* ([https://patentimages.storage.
53 googleapis.com/e5/3d/0d/1c25e5f68a77ab/US7071258.pdf](https://patentimages.storage.googleapis.com/e5/3d/0d/1c25e5f68a77ab/US7071258.pdf)). The first two are associ-
54 ated to the code *H01M*, while the third to the codes *C08K*, *C04B*, *C01B* and *C22C*, that
55 correspond to the gold points. The red points are the technology codes in which Nanotek
56
57
58
59
60

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

1
2
3
4
5 patented in 2000 and 2001, while, as mentioned, the gold ones are those activated in the
6 2002. The black arrows underline the non random position of the new technologies, that
7 are close to the ones already present in the patenting activity of the company. This is
8 because we find that technology codes that have a high similarity are represented close to
9 each other, and therefore a sort of "technological diffusion" is expected starting from the
10 codes that firms already have in their portfolios (as shown in Figure 4b).

11 12 13 Discussion

14
15 In this work we compare machine learning and network-based approaches to forecast which
16 will be the future patenting activity of firms; in particular, their next technological sector
17 of innovation. To the best of our knowledge, this is the first attempt to assess the related-
18 ness between a firm and a technology sector using machine learning. In order to compare
19 the various possible measures of relatedness we analyze a very large database consisting in
20 about two hundred thousand firms and 643 technology sectors and we develop a forecast
21 exercise using the assumption that, on average, firms will patent in sectors related to their
22 present technological activity. We find that supervised machine learning techniques (Ran-
23 dom Forest) clearly outperform the standard methodologies usually adopted in Economic
24 Complexity, that is, networks of co-occurrences. Our results are robust with respect to
25 different definitions of what a "new" technological sectors is, and if different metrics to
26 evaluate the prediction performance are adopted. Indeed, Random Forest assigns on av-
27 erage higher activation scores to those technologies which will be explored by firms with
28 respect to all network based approaches. Finally, we introduce the Continuous Technol-
29 ogy Space (CTS), that permits to visualize the dynamics of firms during their innovation
30 activity. The introduction of this approach opens a number of possible applications and
31 developments. First of all, our activation scores represent an assessment of the achievability
32 of a given jump to a new technology sector, a measure of how easy will be to produce in-
33 novations in that sector given the present activity of the firm. Moreover, the CTS allows a
34 compact visualization of the past, the present and the possible patenting activity of a firm.
35 Using these tools, and in the spirit of the "adjacent possible" approach Kauffman (1996);
36 Tria et al. (2014), it is now possible to quantify how much a firm is *exploring* the space of
37 technologies or *exploiting* what it already does. One can then compare different strategic
38 choices with various measures of performance, both in terms of profitability and further
39 innovation activity. Furthermore, these measures can be applied to investigate Mergers
40 and Acquisitions, and in particular to study whether acquirers prefer to target companies
41 which are "close" or "far" from their present patenting activity. Finally, following the work
42 of Brummitt et al. (2020) and Pugliese et al. (2017), a future research project could be
43 the prediction of some performance-related monetary variables of firms, such as revenue or
44 labor productivity, from knowledge of firms' patent activity.

45 46 47 Methods

48
49 In this Section we describe in more detail the database, algorithms, and evaluation metrics
50 used in the analysis.

51 52 Data

53
54 The bipartite firm-technology network is obtained by matching two databases: AMADEUS
55 for firms and PATSTAT for the technology codes.

FIRMS

AMADEUS (<https://amadeus.bvdinfo.com>) contains information about over 20 million companies, mainly concentrated in the European continent. This database is managed by Bureau van Dijk Electronic Publishing (BvD) which specializes in providing financial, administrative and budget information relating to companies. It is compatible with the PATSTAT database for patents as BvD includes the same patent identifiers as the European Patent Office (Pugliese et al., 2019b). We mention here one of the well-known problems with AMADEUS, namely that large companies are fully covered while those with less than 20 employees are under-represented (Ribeiro et al., 2010); however, this is not a severe issue for the present analysis.

TECHNOLOGY CODES

The dataset from which we take information about the patent and the technology codes is PATSTAT (www.epo.org/searching-for-patents/business/patstat). Globally, PATSTAT considers approximately 100 millions of patents registered in about 100 Patent Offices. This information spans from mid-19th century to three-four years before release of the database; this is evident from the quickly decreasing number of patents in the last available years. As a consequence, we decided to restrict our analysis in a conservative time interval (2000-2011). A key element is the presence of a set of alphanumeric codes in each patent submission; these codes can be assigned by the inventors or by the reviewers and represent the technological sector the patent belongs to. The WIPO (World International Patent Office) uses the IPC (International Patent Classification) (Fall et al., 2003) to assign these technology codes to each patent in such a way as to classify, and better manage, the inventions presented. The IPC codes define a hierarchical classification consisting of six levels: sections (that we call macro category), sub-sections, classes, sub-classes, groups, sub-groups. For example, code Axxxxx corresponds to the "Human Necessities" macro category and Hxxxxx to the "Electricity" macro category; considering the following digits we have, for example, with A01xxx the sector "Agriculture; Hunting", and with A43xxx the "Footwear" sector. It is important to note that we discard classes "99" and sub-classes "Z", as they represent other technologies not classified in other classes or sub-classes, and they are therefore not well defined.

It may happen that the same invention may be referred to for multiple patent application documents. In this case, each group of documents in PATSTAT is called "Patent Family" according to primary citations among them (Publishing et al., 2001), which is nothing more than the set of patents presented in different countries to protect the single invention. Patent Families can be built with different criteria (Martínez, 2011), but among these we choose the one related to the "Extended Family", also called IN-PADOC. This corresponds to the category considered in such a way as to associate the inventions with the widest possible technological spectrum. Once patents are assigned to firms, we can assign them the corresponding technology codes and build the firm-technology bipartite network, and its adjacency matrix \mathbf{V}^y , one for each year y . The matrix element represents the number of patent submitted by the firm in the technology sector. Note that this number may be fractional, since more than one code is usually present in the same patent: for instance, if a firm submits only one patent with three technology codes, the three nonzero elements of the corresponding row of \mathbf{V}^y will be equal to one third. The interested reader can find more details about this data in the results section and in Pugliese et al. (2019b).

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

Data processing

The starting database can be represented using the following structure: 12 matrices, one for each year from 2000 to 2011, that link 426983 firms f (rows) to 7456 (6-digits) technology codes t (columns). We chose to work at a higher aggregation level, and so to compress the technology codes from 6 to 4 digits, summing the columns corresponding to the 6 digit codes with the same first 4 digits. From the 6 to the 4 digit level the number of technologies goes from 7456 to 643. This operation leads to both better quantitative results and shorter computation times (from a qualitative point of view, instead, the results are unchanged). A key element of both the machine learning and the network based approaches is to provide an assessment of the similarity between technology codes; this information can be extracted from the co-occurrences of technological sectors in the same firms. So we consider only firms that, in years from 2000 to 2009, make at least 2 technology codes; these firms are 197944. This leads to the data mentioned to the main text: 12 \mathbf{V} yearly matrices that link 197944 firms and 643 technology codes.

In order to compute the relatedness measures, in the Economic Complexity literature (Hidalgo et al., 2007; Zaccaria et al., 2014; Pugliese et al., 2019a) one usually computes the Revealed Comparative Advantage or RCA (Balassa, 1965), and then these matrices are binarized using a threshold equal to 1. As far as exports are concerned this choice of threshold has a natural economic meaning, traceable to the works of Ricardo and Balassa himself: considering the bipartite country-product network, $RCA_{c,p} \geq 1$ means that country c is significantly competitive in the export of the product p (Hidalgo et al., 2007). So the country's share of that product in its market is equal to or greater than the product's share on the world market. However, the economic meaning of patents submission is different, so the choice of RCA is not straightforward. In this work, we binarize the matrices \mathbf{V} with different values of threshold T , without computing the RCA; in this way, the matrices \mathbf{V} are better interpretable as how much a technology code t is present in the patenting activity of a firm f . We have in any case checked the robustness of our results for different threshold values and the use of RCA.

Network-based approaches

In this and in the next sections we discuss how to obtain a prediction score matrix \mathbf{S} for 2011 from each method starting from the same training data \mathbf{V} and \mathbf{M} , relative to the years 2000-2009. The score matrix gives the model's estimation of the likelihood that a firm will patent in the given technology sector, and the comparison between the scores and the actual \mathbf{M}^{2011} using the performance metrics will give an assessment of the models' performance.

The basic idea of network-based approaches is to compute a similarity of technology codes from their co-occurrences in companies. Introduced by Teece et al. (1994), and popularized in the network/complexity community by Hidalgo et al. (2007), the basic quantity is the number of firms that have patented inventions relating to both codes:

$$B_{t,t'}^{CO} = \sum_f M_{f,t} M_{f,t'}. \quad (1)$$

The idea is that if many firms are active in two technology sectors t and t' at the same time, this means that the capabilities, the techniques and, in general, the necessary means to patent in these sectors, are roughly the same, and so these sectors are, in this sense, similar, or related.

Different scholars presented various ways to normalize the co-occurrences, on the basis of

different theoretical frameworks or interpretations. In general, we can write:

$$B_{t,t'} = \frac{1}{A} \sum_f \frac{M_{f,t} M_{f,t'}}{C}$$

and discuss the various options for the quantities A and C :

- Simple Co-Occurrences (Teece et al., 1994): for $A = 1$ and $C = 1$ one simply counts the number of companies that are active in both sectors. This case corresponds to $B_{t,t'}^{CO}$ of Equation 1;
- Technology Space (same normalization of the Product Space (Hidalgo et al., 2007)): $A = \max(u_t, u_{t'})$ and $C = 1$, where $u_t = \sum_f M_{f,t}$ is the *ubiquity* of technology code t , that is, the number of firms active in that technology sector. Using this type of normalization we give a lower connection weight to those technology codes done by many firms, that we can consider as basic.
- Taxonomy (Zaccaria et al., 2014): $A = \max(u_t, u_{t'})$ and $C = d_f$, where $d_f = \sum_t M_{f,t}$ is the *diversification* of firm f . The Technology Space, for how it is built, gives a higher score for high complexity technology codes (i.e. codes done by few firms) and, as a result, bias towards them. Consequently, it is not possible to justify the evolution of low-complexity technology codes towards high-complexity ones. Normalizing also for the diversification we avoid this problem as we penalize low ubiquity scores and low complexity technology codes are weighted more.
- Micro Partial (Teece et al., 1994): we compute

$$B_{t,t'}^{MP} = \frac{B_{t,t'}^{CO} - \mu_{tt'}}{\sigma_{tt'}}$$

with

$$\mu_{tt'} = \frac{u_t u_{t'}}{N},$$

and

$$\sigma_{tt'}^2 = \mu_{tt'} \frac{(N - u_t)(N - u_{t'})}{N(N - 1)},$$

where N is the number of companies. Here we use a null model in which the ubiquities of the technologies are kept fixed and everything else is randomized. This case can be analytically solved: the resulting distribution for the co-occurrences is hypergeometric with mean $\mu_{tt'}$ and variance $\sigma_{tt'}^2$. We call this network Micro Partial following the notation used by Cimini et al. (2022): this null model is microcanonical in the sense that the degree sequence is exactly fixed and partial because only one layer is constrained. So the idea is that, if the weight of the link between two technology codes t and t' exceeds the expected value $\mu_{tt'}$, this means that t and t' are highly related with respect to this random case. Furthermore, as a t-statistic, $B_{t,t'}^{MP}$ measures how much the observed link between the two technology codes exceeds what would be expected if the companies were randomly assigned.

For the latest formulas we obtain one matrix \mathbf{B}^{Net} for each network. In order to consider all years available in the training data, we using as \mathbf{M} matrix in the previous formulas a total matrix obtained by summing the \mathbf{V} matrices from years 2000 to 2009, using all the 197944 firms, and then binarizing this sum.

Based on the network used, we get a \mathbf{B}^{Net} which we use in the coherence equation from Pugliese et al. (2019b):

$$S_{f,t}^{2011} = \sum_{t'} M_{f,t'}^{2009} B_{t't}^{Net}, \quad (2)$$

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

where $M_{f,t}^{2009}$ is the \mathbf{M} matrix obtained by binarizing the \mathbf{V}^{2009} matrix. In practice, t is highly coherent with the patenting activity of firm f if f is active in many sectors highly connected with t . On the contrary, if a sector is far from what a firm actually does, we will assign to it a lower activation score. Note that this equation differs from the density equation of the Product Space (Hidalgo et al., 2007); we use coherence instead of density since we have found a better predictive performance.

Random Forest

Random Forest (Breiman, 2001) (RF) is a tree-based machine learning algorithm that we use to better capture the non-linear links between technology codes. In particular, we use this binary classification algorithm to determine whether or not a technology code will appear in the patenting portfolio of a particular company in the future starting from the knowledge of the technology codes in which the firms patented in the last training year.

In general, during the training of a supervised machine learning algorithm an input data \mathbf{X} matrix is passed. Because our problem is a supervised one, to each vector (row) of the matrix is associated a label presents in a different input \mathbf{y} . To give an example, \mathbf{X} can be the matrix where each row is a flattened handwritten digit, and each element of the row is the intensity of a pixel; in this case \mathbf{y} will be the label corresponding to the digit, and that must be associated, in order to be recognized, to all those present in \mathbf{X} . Once the model is trained, one gives new samples \mathbf{X}_{test} and the model is able to make associate a prediction \mathbf{y}_{test} (in this case, a digit), to each sample.

In our case, we train one RF for each technology code: we want the RF to learn to which typologies of portfolios is associated each code after two years. So, as samples matrix \mathbf{X} we use the matrix obtained by concatenating, or stacking vertically, the \mathbf{V} matrices from the year 2000 to 2007, and as \mathbf{y} we use one column at a time (and therefore one technology code at a time) of the matrix obtained by concatenating the matrices \mathbf{M} from the year 2002 to 2009. In this way, each row is a firm in a year from 2000 to 2007, that has 643 features. We associate this sample to the respective label in \mathbf{y} , that is, if after 2 years the technology code associated to the element in \mathbf{y} is active, or not. In such a way, we associate the codes of each portfolio to the possible presence of the target code in the future.

From a practical viewpoint, we use the "RandomForestClassifier" from the "sklearn.ensemble" python library (Pedregosa et al., 2011), called in this way:

$$\text{RandomForestClassifier.fit}(\mathbf{V}_{2000}^{2007}, \vec{\mathbf{M}}_{2002}^{2009}),$$

where \mathbf{V}_{2000}^{2007} are the vertically stacked matrices and with the vector symbol over \mathbf{M} we indicate that one column is used at a time, that is, we train one RF for each technology code. The delay of 2 years is used to insert a dependence on time, as we want to produce forecasts about the innovative development of firms. We optimized the RF parameters as described in the supplementary information; the results shown here refer to: number of trees = 50, min_samples_leaf = 4; max_depth = 40 and method = 'entropy'. The use of all available companies in the training is computationally demanding, so we used only the top 10000 most diversified firms (10KHD firms). If we use more firms for the training we get a saturation of the forecast performances (see the Supplementary information). The fact that firms with higher diversification should be used is due to the fact that these provide a better coverage of the possible technologies and the possible combinations among them. After fitting the data, that is, training the machine learning model, we obtain the \mathbf{S}^{2011} scores by using the \mathbf{V}^{2009} matrix as \mathbf{X}_{test} in a predicting phase. The command line reads

$$\mathbf{S}^{2011} = \text{RandomForestClassifier.predict_proba}(\mathbf{V}^{2009})$$

and this associates a probability to activate the target technology to each firm in 2009. In Figure 6 we schematically represent how \mathbf{S}^{2011} is obtained from the RF, the latter being represented by a set of differently trained decision trees.

In this paper we also compare the approach described above with a cross-validated RF, for which we use the same optimized parameters and, in order to provide a consistent comparison of the results, the same training and test sets. In this respect, note that the $\text{textbf{X}}_{test}$ should always produce a prediction for all the 197944 firms (including the 10KHD firms used in the training). In the cross-validation framework, we train $k = 4$ different RFs, using the technique called *k-fold Cross Validation*, that is we separate into $k=4$ groups both the 10KHDs used for the training and the 197944-10KHDs used for the test. Then, in training each RF we remove one of the 4 groups of 10KHD companies. The test is instead performed on the removed group from the 10KHD companies along with 1/4 of the 197944 - 10KHD firms used for the test. This is performed 4 times, each time removing a different set of firms. At the end, we will have prediction scores for all companies by merging the scores produced by the four RFs.

The idea behind the use of cross validation is the following. During the training the RF basically learns two pieces of information: to recognize the portfolio of a company and the similarity among technologies. Even if we are more interested to the latter, the learning of the two can not be avoided. However, we can try to force the algorithm to use the similarities in the test phase: if we give a new company in the \mathbf{X}_{test} , the RF can not recognize it and so it is forced to use the similarities to produce its predictions. This procedure, even if computationally more demanding, leads to better results, as shown in Figure 2.

Continuous Technology Space

Random Forest shares with most of the machine learning algorithms an intrinsic difficulty of interpretation, i.e. the rationale behind how the input is connected to the output is not evident. In this respect, network approaches (note: if made sparse by a suitable filter) are more clear, since the coherence or density based approach are clearly visualizable: a technology is coherent with a firm's portfolio if has a lot of heavy connections with what the firm already does. In order to restore the interpretability of networks and keep the predictive performance of machine learning, Tacchella et al. (2021) propose the Continuous Projection Space, that here we reformulate, with suitable modifications, as the Continuous Technology Space (CTS).

To compute the CTS we start from Random Forest CV method but, in \mathbf{X} , only the first 2K HD firms are used, because we have a saturation of the scores: using more firms does not change the scores and increases the computational time.

Another difference with the Random Forest CV is that the predictions are obtained using as \mathbf{X}_{test} the same 2K HD firms used for the training (i.e., in the CTS $\mathbf{X} = \mathbf{X}_{test}$; however, we use k-fold CV to avoid overfitting problems). At the end we obtain a scores matrix of shape $[N \times \text{years}] \times [\#t]$, where N is the number of companies ($N = 2000$), years = 10 and $\#t =$ number of technology codes = 643; in total this scores matrix has shape 20000×643 . Each column of the score matrix represents the likelihood that each company (rows) will patent in each technology code (columns). We can then argue than two sectors are *similar* if the RF predicts that the same companies will (or will not) produce patents in these sectors. In this sense, the columns of the score matrix can be seen as the coordinates in a high-dimensional space for each technology code, where the number of dimensions is given by the number of companies multiplied by the number of training years (in this case, 20000). In order to provide a better interpretability to the relatedness assessment, one should find a low dimensional visualization of this high dimensional representation. Obviously, it is

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

impossible to visualize this continuous space of technologies in such an high dimensionality; so we project these points in a lower dimension by combining a Variational - Autoencoder Neural Network (Kingma and Welling, 2013), to reduce the dimension from 20000 \rightarrow 150, and then t-SNE (Van der Maaten and Hinton, 2008), to reduce the embedding space from 150 \rightarrow 2 dimensions, finally obtaining the Continuous Technology Space (CTS), that we show in Figure 4a. Now the similarity between technology codes is simply given by the relative distance in this 2 - D space, and it is easy to understand and visualize how firms move from the codes already present in their portfolios to the ones that are immediately close, as shown in Figure 5.

Now we want to use the idea of a coherent diffusion in this low dimensional space to produce forecasts; in practice, to obtain a score matrix \mathbf{S}^{2011} to compare with the possible activations of 2011. We start by computing a similarity matrix for the CTS, that for the sake of simplicity we keep calling \mathbf{B} . We use the distances between technology codes on the CTS and gaussian kernels:

$$B_{i,j} = \frac{e^{-\|\mathbf{y}_i - \mathbf{y}_j\|^2 / 2\sigma_i^2}}{\sum_k e^{-\|\mathbf{y}_i - \mathbf{y}_k\|^2 / 2\sigma_i^2}},$$

where \mathbf{y}_i is the coordinate of the i -th technology code in the CTS (i.e. in the 2 - D space). The σ_i is the standard deviation of the Gaussian kernel related to the technology code i -th; this parameter can be set differently for each i code, through a binary search process in which the number of first neighbors is fixed. As we can see in Figure 4a, there are codes in dense areas and codes in less dense areas, so the idea is to assign a high sigma value to the codes in less dense areas and low sigma values in more dense areas in order to keep the interaction with the number of first neighbors constant. The binary search process is described in the supplementary information where we also show that the best optimal value of nearest neighbors is 75.

After the similarity matrix \mathbf{B} is obtained, one can compute the score matrix \mathbf{S}^{2011} from the coherence equation Eq. (2):

$$S_{f,t}^{2011} = \sum_{t'} M_{f,t'}^{2009} B_{t',t}.$$

The number of nearest neighbors is calculated out of sample using the 4-fold cross validation as in the case of the Random Forest CV: we use 3/4 of the companies to determinate the number of nearest neighbors that maximize the Area under the PR curve and then we calculate the scores using the Equation 2 for the remaining companies.

Note that the CTS is, like the network approaches, density-based: the more a firm surrounds a technology sector, more likely it will be part of its patenting activity in the near future.

Benchmark models

In order to understand the effective goodness of our forecast results, a comparison with some relatively trivial benchmark models is required. We used two benchmark models:

- The first consists in a simple randomization of the technology codes. In practice, we shuffle the columns of the \mathbf{M}^{2009} matrix in the calculation of Equation 2. The \mathbf{B} used is that calculated with Technology Space network starting from the not randomized \mathbf{M}^{2009} (using the other networks there is no significant change in the metric scores). In this way, the diversification of firms is preserved.
- The second benchmark model checks the hypothesis that the simple temporal autocorrelation of the bipartite networks can explain the observed dynamics. In this case, we use the mean

STRACCAMORE, PIETRONERO AND ZACCARIA

of the \mathbf{V} s matrices from 2000 to 2009, $\bar{\mathbf{V}}$ (all the years used in the training) of the test firms as score matrix \mathbf{S}^{2011} , that is, element-wise:

$$\mathbf{S}^{2011} = \bar{\mathbf{V}}.$$

In this way we check if the number of patents done in the past can forecast the number of patents done in the future by the same company in the same technology sector. As shown in Figure 2, this benchmark model can outperform some of the density-based approaches.

Prediction performance metrics

In order to compare the goodness of the predictions of the different approaches we use standard evaluation metrics, widely used for classification tasks in supervised machine learning (Hossin and Sulaiman, 2015). As different metrics capture different aspects of the prediction problem, only the comparison between various measure of performance can provide a global view of the effectiveness of a forecast approach.

The elements that we want to predict are the possible activations, that is, those elements of \mathbf{M}^{2011} that were always zero in 2000-2009. The 0s are called negatives, and the 1s are called positives. The elements equal to 1 that are correctly predicted are called true positives (TP); and similarly one can define the false positives (FP), the true negatives (TN) and the false negatives (FN) as, respectively, the 0s predicted as 1s, the correctly predicted 0s, and the 1s predicted as 0s.

To evaluate the predictions done with the different approaches, we have used three performance metrics:

- **Area under the PR curve:** This indicator is equal to the area in the Precision - Recall plane. Precision is defined as $\frac{TP(\tau)}{TP(\tau)+FP(\tau)}$, while recall is equal to $\frac{TP(\tau)}{TP(\tau)+FN(\tau)}$. These quantities are close to 1 if FP and FN are respectively minimized. Note that in order to compute precision and recall one has to specify the scores' binarization threshold τ , that is, the number above which the score is associated to a positive prediction (1). The PR curve is defined by varying the τ parameter, because for different values of τ we obtain different precision and recall values. The last step is the computation of the area under this curve, which is independent from the threshold τ .
- **Precision@100:** to compute this indicator we focus on the top 100 scores elements in \mathbf{S}^{2011} : if the model is correct, many of these possible activations should become realized activations. The Precision@100 is the ratio between the number of these 100 that are true positives (that is, correctly predicted realized activations), and 100, i.e. the number of elements that we are considering. This represents a global assessment, that considers the score matrix as a whole.
- **mPrecision@10:** While the Precision@100 provides a global measure of the precision of the approach, we would like to have a measure of our average predictive performance for each firm. To do this, we evaluate the mPrecision@10. We consider the 10 highest scores for each row, i.e. for each firm, and compute the fraction of true positives. Then we average over the firms. Since most of the firms do not show at least 10 realized activations, the global number is far from 1. We have computed the mPrecision also restricting ourselves only to the firms with 10 or more realized activations, finding similar qualitative results.

References

Giambattista Albora, Luciano Pietronero, Andrea Tacchella, and Andrea Zaccaria. Product progression: a machine learning approach to forecasting industrial upgrading. *arXiv preprint arXiv:2105.15018*, 2021.

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

- 1
2
3
4
5 Bela Balassa. Trade liberalisation and “revealed” comparative advantage 1. *The manchester school*,
6 33(2):99–123, 1965.
- 7 Charles H Berry. Corporate growth and diversification. *The Journal of Law and Economics*, 14(2):
8 371–383, 1971a.
- 9
10 Charles H. Berry. Corporate growth and diversification. *The Journal of Law & Economics*, 14(2):
11 371–383, 1971b. ISSN 00222186, 15375285. URL <http://www.jstor.org/stable/724951>.
- 12
13 Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- 14
15 Stefano Breschi, Francesco Lissoni, and Franco Malerba. Knowledge-relatedness in firm technological
16 diversification. *Research policy*, 32(1):69–87, 2003.
- 17
18 Charles D Brummitt, Andrés Gómez-Liévano, Ricardo Hausmann, and Matthew H Bonds. Machine-
19 learned patterns suggest that diversification drives economic development. *Journal of the Royal*
20 *Society Interface*, 17(162):20190283, 2020.
- 21
22 Tullio Buccellato. The competences of firms are the backbone of economic complexity. *Available at*
23 *SSRN 2827468*, 2016.
- 24
25 Giulio Cimini, Alessandro Carra, Luca Didomenicantonio, and Andrea Zaccaria. Meta-validation of
26 bipartite network projections. *Communications Physics*, 5(1):1–12, 2022.
- 27
28 Caspar J Fall, Atilla Töröcsvári, Karim Benzineb, and Gabor Karetka. Automated categorization in
29 the international patent classification. In *Acm Sigir Forum*, number 1, pages 10–25. ACM New
30 York, NY, USA, 2003.
- 31
32 Michael Gort. *Diversification and Integration in American Industry*. National Bureau of Economic
33 Research, Inc, 1962a. URL <http://EconPapers.repec.org/RePEc:nbr:nberbk:gort62-1>.
- 34
35 Michael Gort. *Diversification and integration in American industry: a study by the National Bureau*
36 *of Economic Research*. Princeton University Press, 1962b.
- 37
38 Bronwyn H Hall, Adam B Jaffe, and Manuel Trajtenberg. The nber patent citation data file: Lessons,
39 insights and methodological tools. Technical report, National Bureau of Economic Research, 2001.
- 40
41 César A Hidalgo, Bailey Klinger, A-L Barabási, and Ricardo Hausmann. The product space condi-
42 tions the development of nations. *Science*, 317(5837):482–487, 2007.
- 43
44 César A Hidalgo, Pierre-Alexandre Balland, Ron Boschma, Mercedes Delgado, Maryann Feldman,
45 Koen Frenken, Edward Glaeser, Canfei He, Dieter F Kogler, Andrea Morrison, et al. The principle
46 of relatedness. In *International Conference on Complex Systems*, pages 451–457. Springer, 2018.
- 47
48 Mohammad Hossin and Md Nasir Sulaiman. A review on evaluation metrics for data classification
49 evaluations. *International journal of data mining & knowledge management process*, 5(2):1, 2015.
- 50
51 Adam B Jaffe, Manuel Trajtenberg, and Michael S Fogarty. Knowledge spillovers and patent cita-
52 tions: Evidence from a survey of inventors. *American Economic Review*, 90(2):215–218, 2000.
- 53
54 Si Joo and Yeonbae Kim. Measuring relatedness between technological fields. *Scientometrics*, 83
55 (2):435–454, 2010.
- 56
57 Stuart A Kauffman. *Investigations: The nature of autonomous agents and the worlds they mutually*
58 *create*. Santa Fe Institute, 1996.
- 59
60

STRACCAMORE, PIETRONERO AND ZACCARIA

- 1
2
3
4
5 Seung Hwan Kim, Bogang Jun, and Jeong-Dong Lee. Technological relatedness: How do firms
6 diversify their technology? ., 2021.
- 7
8 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*
9 *arXiv:1312.6114*, 2013.
- 10
11 Sotiris B Kotsiantis, I Zaharakis, P Pintelas, et al. Supervised machine learning: A review of
12 classification techniques. *Emerging artificial intelligence applications in computer engineering*,
13 160(1):3–24, 2007.
- 14
15 Bart Leten, Rene Belderbos, and Bart Van Looy. Technological diversification, coherence, and
16 performance of firms. *Journal of Product Innovation Management*, 24(6):567–579, 2007.
- 17
18 Alessia Lo Turco and Daniela Maggioni. On firms’ product space evolution: the role of firm and
19 local product relatedness. *Journal of Economic Geography*, 16(5):975–1006, 2016.
- 20
21 Vittorio Loreto, Vito DP Servedio, Steven H Strogatz, and Francesca Tria. Dynamics on expanding
22 spaces: modeling the emergence of novelties. In *Creativity and universality in language*, pages
23 59–83. Springer, 2016.
- 24
25 Catalina Martínez. Patent families: When do different definitions really matter? *Scientometrics*,
26 86(1):39–63, 2011.
- 27
28 Lionel Nesta and Pier-Paolo Saviotti. Firm knowledge and market value in biotechnology. *Industrial*
29 *and Corporate Change*, 15(4):625–652, 2006.
- 30
31 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Pretten-
32 hofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot,
33 and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning*
34 *Research*, 12:2825–2830, 2011.
- 35
36 Edith Penrose. *The theory of the growth of the firm*. John Wiley, New York, 1959.
- 37
38 Edith T Penrose. The growth of the firm—a case study: the Hercules Powder Company. *Business*
39 *History Review*, 34(01):1–23, 1960.
- 40
41 OECD. Publishing, Organisation for Economic Co-operation, and Development Staff. *OECD sci-*
42 *ence, technology and industry scoreboard 2001: Towards a knowledge-based economy*. Organisation
43 for Economic Co-operation and Development, 2001.
- 44
45 Emanuele Pugliese, Lorenzo Napolitano, Andrea Zaccaria, and Luciano Pietronero. Coherent diver-
46 sification in corporate technological portfolios. *arXiv preprint arXiv:1707.02188*, 2017.
- 47
48 Emanuele Pugliese, Giulio Cimini, Aurelio Patelli, Andrea Zaccaria, Luciano Pietronero, and Andrea
49 Gabrielli. Unfolding the innovation system for the development of countries: coevolution of
50 science, technology and production. *Scientific reports*, 9(1):1–12, 2019a.
- 51
52 Emanuele Pugliese, Lorenzo Napolitano, Andrea Zaccaria, and Luciano Pietronero. Coherent diver-
53 sification in corporate technological portfolios—supplementary information. *PloS one*, 2019b.
- 54
55 Pouya Rahmati, Ali R Tafti, J Christopher Westland, and César Hidalgo. When all products are
56 digital: Complexity and intangible value in the ecosystem of digitizing firms. *Forthcoming, MIS*
57 *Quarterly*, 2020.

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

- 1
2
3
4
5 Samuel Pinto Ribeiro, Stefano Menghinello, and Koen De Backer. The oecd orbis database: Re-
6 sponding to the need for firm-level micro-data in the oecd. *OECD Statistics Working Papers*,
7 2010.
- 8 David L Rigby. Technological relatedness and knowledge space: entry and exit of us cities from
9 patent classes. *Regional Studies*, 49(11):1922–1937, 2015.
- 10 Richard P Rumelt. Strategy, structure, and economic performance. ., 1974.
- 11 Richard P Rumelt. Diversification strategy and profitability. *Strategic management journal*, 3(4):
12 359–369, 1982.
- 13 Robert W Rycroft and Don E Kash. *The complexity challenge: Technological innovation for the*
14 *21st century*. Burns & Oates, 1999.
- 15 Angelica Sbardella, Emanuele Pugliese, Andrea Zaccaria, and Pasquale Scaramozzino. The role of
16 complex analysis in modelling economic growth. *Entropy*, 20(11):883, 2018.
- 17 Brent Smith and Greg Linden. Two decades of recommender systems at amazon. com. *Ieee internet*
18 *computing*, 21(3):12–18, 2017.
- 19 Deborah Strumsky, Jose Lobo, and Sander Van der Leeuw. Measuring the relative importance
20 of reusing, recombining and creating technologies in the process of invention. Technical report,
21 Working Paper, 2011.
- 22 Deborah Strumsky, José Lobo, and Sander Van der Leeuw. Using patent technology codes to study
23 technological change. *Economics of Innovation and New technology*, 21(3):267–286, 2012.
- 24 Andrea Tacchella, Matthieu Cristelli, Guido Caldarelli, Andrea Gabrielli, and Luciano Pietronero.
25 A new metrics for countries' fitness and products' complexity. *Scientific reports*, 2:723, 2012.
- 26 Andrea Tacchella, Dario Mazzilli, and Luciano Pietronero. A dynamical systems approach to gross
27 domestic product forecasting. *Nature Physics*, 14(8):861–865, 2018.
- 28 Andrea Tacchella, Andrea Zaccaria, Marco Miccheli, and Luciano Pietronero. Relatedness in the
29 era of machine learning. *arXiv preprint arXiv:2103.06017*, 2021.
- 30 David J Teece, Richard Rumelt, Giovanni Dosi, and Sidney Winter. Understanding corporate
31 coherence: Theory and evidence. *Journal of economic behavior & organization*, 23(1):1–30, 1994.
- 32 Francesca Tria, Vittorio Loreto, Vito Domenico Pietro Servedio, and Steven H Strogatz. The dy-
33 namics of correlated novelties. *Scientific reports*, 4(1):1–8, 2014.
- 34 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine*
35 *learning research*, 9(11), 2008.
- 36 Bowen Yan and Jianxi Luo. Filtering patent maps for visualization of diversification paths of
37 inventors and organizations. *Journal of the Association for Information Science and Technology*,
38 68(6):1551–1563, 2017.
- 39 Hyejin Youn, Deborah Strumsky, Luis MA Bettencourt, and José Lobo. Invention as a combinatorial
40 process: evidence from us patents. *Journal of the Royal Society interface*, 12(106):20150272, 2015.
- 41 Andrea Zaccaria, Matthieu Cristelli, Andrea Tacchella, and Luciano Pietronero. How the taxonomy
42 of products drives the economic development of countries. *PloS one*, 9(12):e113770, 2014.
- 43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

STRACCAMORE, PIETRONERO AND ZACCARIA

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgements

We thank Lorenzo Napolitano for providing the patent data used for this work and Andrea Tacchella for insightful discussions. All the authors acknowledge the CREF project "Complessità in Economia".

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

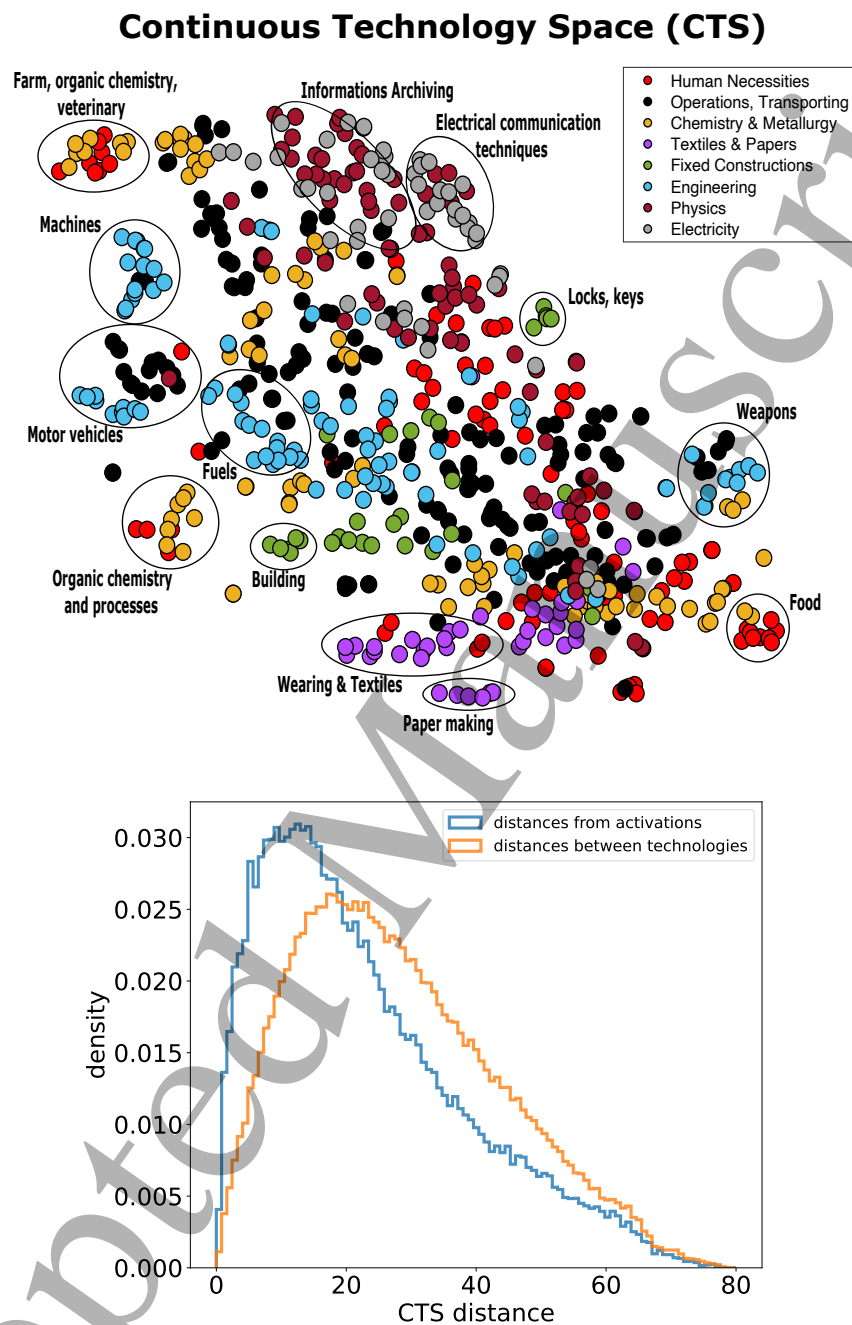


Figure 4: Above: representation of the Continuous Technology Space (CTS). Each point is a technology code and each colour is associated to the respective macrocategory, i.e. the first of the 4 digits. Clusters related to macro categories are evident. Below: comparison between two density distributions of distances. The first (cyan line) represents the distances between activation codes patented by a company from the closest code the company used; the second (orange line) is the density distribution of distances between points in the CTS.

STRACCAMORE, PIETRONERO AND ZACCARIA

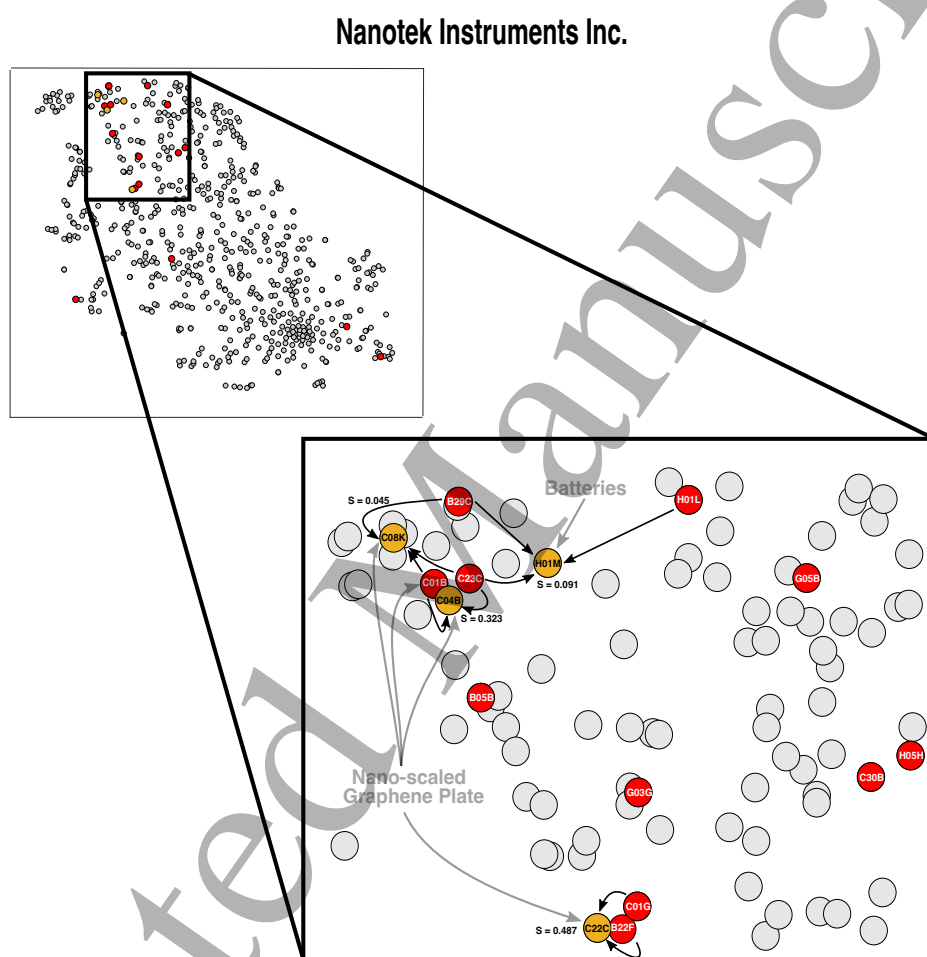


Figure 5: The exploration of the Continuous Technological Space by an American nanotechnology company. Starting from the red sectors, in which Nanotek patented in the past, this company moved nearby, patenting in the gold sectors in the next years.

WHICH WILL BE YOUR FIRM'S NEXT TECHNOLOGY?

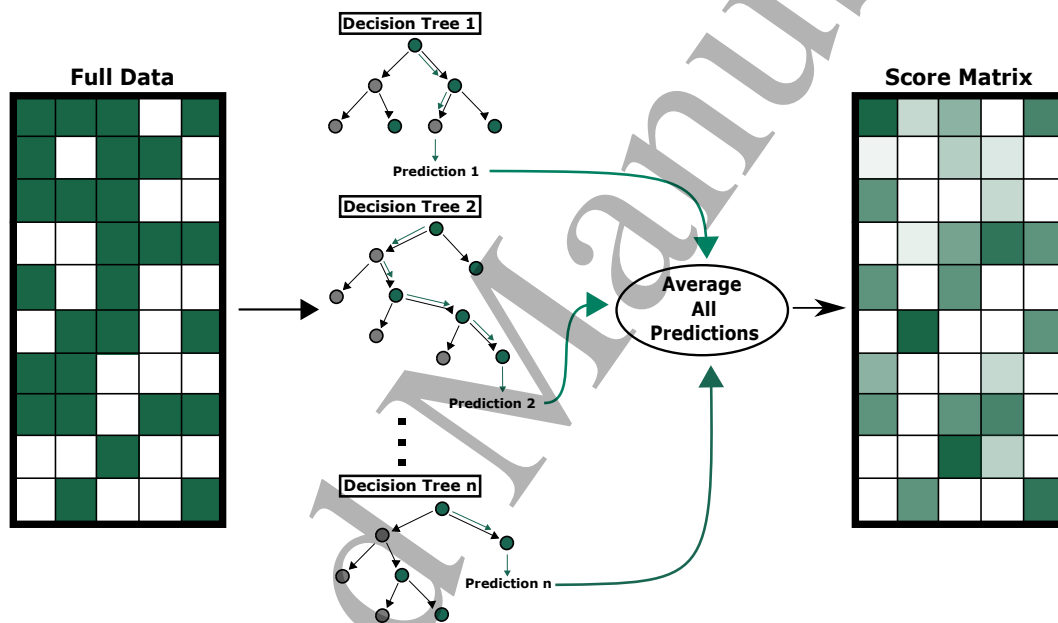


Figure 6: We show how the score matrix S^{2011} is obtained by the Random Forest by combining the predictions of different decision trees.