AI and Ethics (2025) 5:4377–4387
https://doi.org/10.1007/s43681-025-00697-6

ORIGINAL RESEARCH

Beyond arbitrage: machine learning for financing impactful economic activity

Dalila Benachenhou[1] · Masud Cader[2] · Dario Mazzilli[3] · Andrea Tacchella[3]

Received: 9 January 2024 / Accepted: 27 February 2025 / Published online: 12 May 2025

Abstract
Institutional investors are looking for opportunities that will have a positive social and environmental impact as well as positive returns. Compared to the largely speculative trading in equity markets and financial markets, this is a significant change. Investing with impact requires a long-term perspective. Impact investments are becoming increasingly popular as governments expect institutional investors to contribute to the achievement of UN sustainable development goals by 2030. Several characteristics are included in these goals, including environmental sustainability, social inclusion, integration, competitiveness, and resilience. In this paper, new AI methods are presented that utilize network theory, complex fitness dynamics of networks, and machine learning techniques for sourcing investments more effectively and forecasting their likely impact more accurately. The paper discusses ethical considerations and safeguards that should be observed when deploying artificial intelligence for impact investment.

Keywords Economic complexity · Machine learning · Network validation · Trustful AI

1 Introduction

Low-interest rates of the past decade and the US fixed-income bull market since 1980 gave rise to speculative investment. This was a regime in which market patterns (of prices and volumes) could be data mined for profits irrespective of the actual productive performance of underlying companies. Concurrently, there was a realization that the world must aspire to allocate capital towards achieving social and environmental impact by reaching a set of consensus global goals (the UN Sustainable Development Goals or SDGs) and that achieving these goals required increased impactful investment by the private sector. AI, particularly ML and GenAI, can evolve from solely and simply capturing temporary market arbitrage opportunities to assisting in longer term investing which delivers impact and progress toward achieving the SDGs - yielding virtuous growth, lower inequality, and fewer negative impacts on the planet. Impact investment creates new ethical considerations around intentionality, distribution of welfare across generations and beneficiaries. Add to that the ethical complications of modeling, measuring, and expressing investment impact using ML and GenAI co-pilots, and we have a recipe for impact washing. This paper introduces the concept of a socioeconomic activity network (FermiKNet) leveraging the ubiquity of granular data capture (e.g., global trade networks, technology dependencies, skills, mobility, and others), discusses how such a knowledge network gives rise to the characterization of impact from competition, competitiveness, future growth, and production capabilities through a suite of fitness indicators (i) measuring importance in the knowledge network and providing typical measures of growth (e.g., GDP) with a capability dimension, (ii) which are rigorously developed with network null models, and (iii) deployed reliably – creating trust in collaborative use with human experts, dealing with signal recovery from noisy data, dimensionality reduction, as well as drift in the underlying data, bias, and MLOps governance. Our results show that AI models and network representations can be ethically developed and deployed, and add important private sector texture to impact investment methods.

✉ Andrea Tacchella
andreatacchella@gmail.com

[1] Melange Research, Rockville, MD, USA

[2] International Finance Corporation, Washington, USA

[3] Centro Ricerche Enrico Fermi, Rome, Italy

2 Machine learning and AI in finance

AI has ignited significant interest in algorithmic trading, asset and derivative pricing, automation, financial modeling, fraud detection, credit and insurance underwriting, prevention of financial risk, sentiment analysis, and trade settlement [5]. AI methods in finance have increasingly narrowed to the application of neural networks and deep learning methods in three main classes – supervised, unsupervised, and unsupervised with reinforcement learning [6]. Genetic algorithms, fuzzy logic, hybrid systems, and other branches of AI applied to finance have become increasingly niche [12].

Econometric models (typically multivariate linear regression) are used due to their simplicity and low signal-to-noise ratio [14]. Yet, economic and financial data has evolved beyond infrequent macro and low-dimensional market information (prices, volumes, market trades) to include alternative data of unstructured textual information [4, 8], voice recordings, news articles, social media posts, and satellite images [25].

ML models are used to decipher complex patterns in a high-dimensional data setting and can be applied to deeper exploration, prediction, causal inferencing, and visualization than traditional econometric models. ML can be used to overcome the limits of traditional econometric models in dealing with transformational nonlinearities and structural change [24], detecting outliers, extracting features, and performing classification and regression of complex data [16] while limiting the dangers of p-value hacking [13] by using randomization inference methods [2].

ML models have become increasingly sophisticated in recent years, and they are capable of predicting and anticipating the interconnected nature of economic systems and their multimodal characteristics. Despite their usefulness for econometric estimation, production functions that combine economic growth, output, capital, labor, and innovation/total factor productivity are aggregated (e.g., total labor hours which is the sum of hours worked by all workers) and do not offer sufficient granularity (e.g., the hours worked in activities related to making a particular output – e.g., a ball bearing) for expressing interactions that can be translated into profitable investment strategies which contribute to the Sustainable Development Goals. Knowledge networks or graphs are data constructs that people can understand well because they can depict complex interactions and relationships easily through visualization and provide structure as inputs to machine learning algorithms, therefore reducing the computational effort required to estimate these relationships when using conventional machine learning algorithms [15]. In this paper, the Fermi Knowledge Network serves as the basis for linking technological innovation with output in the form of trade [18]. There have been several Knowledge Network projects developed to address the Sustainable Development Goals, including SustainGraph [10] and SustainBench [26]. These projects provide valuable information regarding potential linkages between Sustainable Development Goals and progress tracking towards 2030 targets, but not about the disaggregated production function innate in FermiKNet.

3 The SDGs are changing the way investors allocate capital

The sustainable development goals or SDGs are seventeen goals signed by all countries of the United Nations in 2015 [22]. The agreement set a deadline of 2030 and invited the fulfillment of people's needs (social dimension), within the planet's physical limits (environmental dimension), while recognizing that commerce and jobs produced by the private sector and enabled by states are vital to sustainably achieving the economic dimension.

This perspective required a new form of AI and Machine Learning (ML) on a socioeconomic activity network that captures the dynamics of output growth, labor, and innovation and the complex and potentially nonlinear interactions across social, environmental, and economic dimensions or systems. This view of the economy realizes systems within systems, requires new ways of assessing competition and cooperation dynamics with network structures, and demands new measures that capture important dynamics enabling and predicting growth and opportunity pathways and impacts in the network at various geographic scales. AI could be used to develop joint actions aimed at preserving the environment, including climate action, life below water and life on land. AI advances will support understanding climate change and modeling its possible impacts, and low-carbon energy systems. AI can help improve ecosystems' health, such as by preventing and significantly reducing marine pollution, and by identifying desertification trends over large areas [23]. FermiKNet is a modest first step in this direction.

Finance is the allocation of capital to projects. Projects typically require ring-fencing and estimating future cashflows based on the successful production of targeted outputs. Anything outside of the project context is considered an externality and excluded from the project's valuation and pricing. Corporations are collections of projects that can be implemented with positive returns because of their capabilities and ability to integrate into and leverage the broader socioeconomic systems in which they operate [7].

The SDGs suggest that investors will no longer be able to price a project excluding impacts on third parties in the

Table 1 Key fitness framework indicators

Economic fitness framework indicator	Description & application
Product Progression Probability	Represents the feasibility (as a probability) of cultivating competitiveness in a specific product or sector over a five-year horizon, given existing productive capabilities
Complexity of a Product or Industry	Characterizes the level and nature of productive capabilities required to competitively export a product
Complexity Gain	A probabilistic expression of the expectation that success in one sector will unlock opportunities to sidegrade/upgrade to related products/sectors/industries in the future – that is, to enable the catalytic creation of new markets
Economic Fitness Indicator	Importance of a geographic location (typically country) based on diversity of economic activity and complexity of outputs

environmental or social dimensions – previously considered externalities. Investors will need to consider the relationships or theories of change and the positive and negative impacts across economic, social, and environmental dimensions. This change requires an unprecedented perspective afforded by big data infrastructure, the internet, and digitalization. While data is more ubiquitous, it requires organization. Yet, simply summing number of jobs created, profitability, inequality, gender and other aggregated impact indicators will not provide a wholly accurate perspective of project bankability with potential for impact across different system levels and dimensions as required by SDGs.

4 Fitness: a novel investment framework

The fitness suite of indicators (Table 1) provides forward-looking views on growth, production opportunities, innovation, and inequality at a granular level covering 191 countries with frequent updates over thousands of goods and services (collectively products), technologies, skills, research and education. The fitness indicators are estimated using techniques from dynamic and complex systems and machine learning on a knowledge network representing the main factors in an economy (Fig. 1) [18].

Knowledge networks provide a construct for storing socioeconomic data and their interactions in a granular and dynamic manner. Instead of an aggregation of price multiplied by volume and then averaged, it is now possible to track thousands of products, their precursors and correlates in terms of detailed job skills, technologies, and research areas among others. The web of associations across products, the activities which go into making products competitively, and the spillover markets they might create catalytically provides the basis for estimating conditional probabilities through time. These probabilities identify

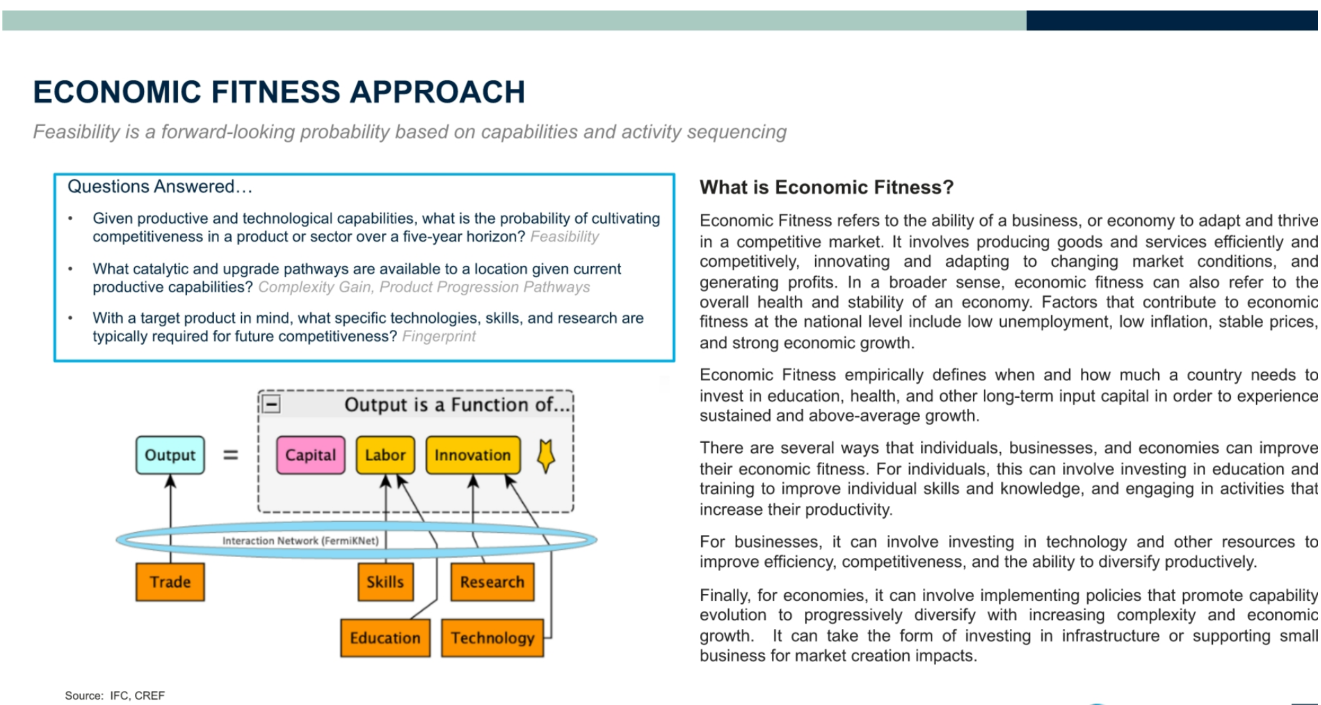

Fig. 1 Overview of fitness framework in economic analysis. Situates the interaction network of trade, skills, education, research, and technology relative to key factors of the typical economic production function. Fitness can be used at different scales from blocks of countries (trade blocks such as the EU), down to cities if the data is appropriately geocoded. Other information, such as corporate spending activity, provides fitness characterization of company-level management capabilities

opportunities with sustainable growth potential for both government policymakers and company strategists.

For example, feasibility in the fitness realm is the probability of producing a target product competitively in five years given the other products made and the technologies mastered in a location. While conventional economic aggregation serves a purpose, it reduces specificity, which is important to individual companies as they engage in potentially profitable projects using their own operational and management capabilities embedded within broader location capabilities. Since firm management requires granular information, and typical economic aggregation does not provide adequate detail, this typically requires separate strategy and microeconomic modeling, which is often reductionist (e.g., taking the project out of its environment – creating externalities - and potentially only capturing impacts through qualitative adjustments to discount rates used in the project's valuation).

5 From "how much you make?" to "what can you make?"

The Economic Fitness and Complexity (EFC) framework provides a way to depart from the traditional approach in the evaluation of economic performance. At its core, EFC shifts the focus from purely monetary considerations ("how much you make?") to a systemic perspective where economic value is seen in the diversity of capabilities and the ability to combine them to increase the diversity and complexity of outputs further ("what can you make?").

Examining economic data in terms of diversity of outputs, rather than only volumes, opens new perspectives and reveals patterns that are typical of ecological networks. Being diversified relates to stronger resilience, the ability to innovate faster, lower inequality, better carbon footprints and stronger economic growth, as measured by standard indicators.

Economic Fitness is a systemic indicator that quantifies the diversity and combinatorial effectiveness of the capabilities of an economic system. It can be computed based on international trade data, patenting activities, research publications and other granular outputs of the location. It is composed of two interconnected parts, the Fitness of economic actors, and Complexity of economic activities: the Fitness is a measure of the diversification of the actor's outputs, weighted by their Complexity; the Complexity of outputs is a nonlinear function of the Fitness of actors that can produce that output, with the one with the lowest Fitness providing an upper bound to the Complexity. In technical terms Fitness and Complexity are a non-linear centrality measure for bipartite networks, i.e., their values depend only on the topology of the actors-activites connections in a global or local economic system. For a more technical definition, see Appendix (Fig. 2).

Fitness is an extremely effective indicator that captures the ability of the system to produce outputs reliably and sustainably. When combined with intensive monetary indicators (e.g., purchasing parity-adjusted GDP-per capita) over time, it generates a low-dimensional description of such systems that captures the main features of medium-term development (5–10 years in the future). Dealing with

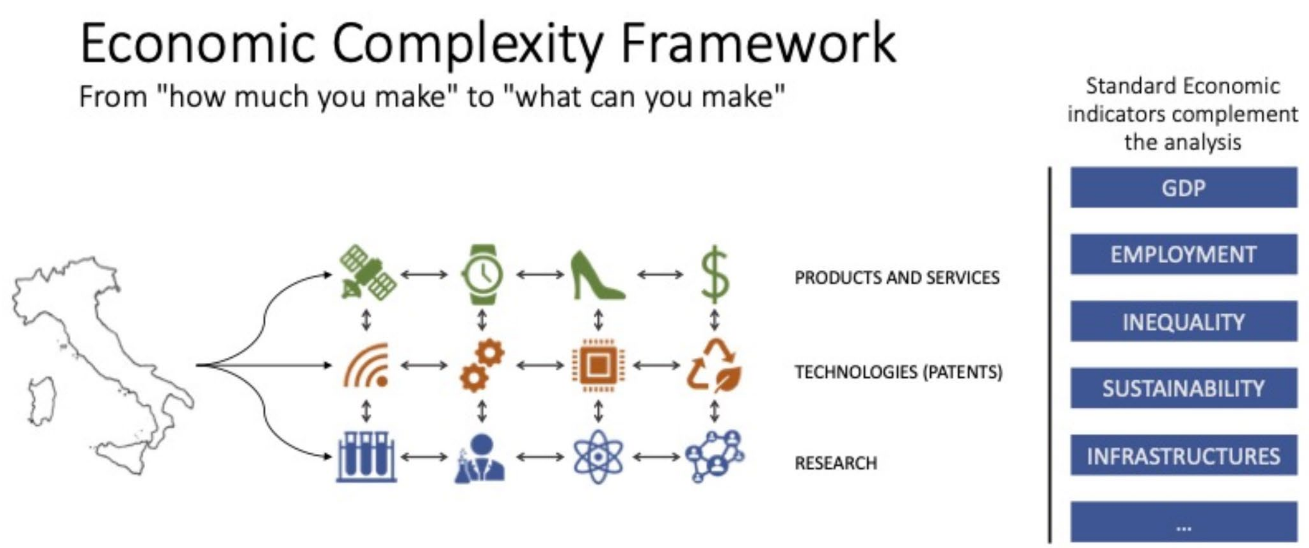

Fig. 2 Schematic of several FermiKNet layers (products, technologies, research) and interactions amongst them as key inputs to computing fitness indicators. Fitness and standard economic (or aggregated indicators) evolve in ways that explain and anticipate an actor's (e.g., location or firm in Italy) role in attaining stronger resilience, innovating faster, achieving lower inequality, improving carbon footprints, and contributing to stronger economic growth

low-dimensional systems allows the authors to develop simple, explainable ML approaches to forecast their growth trajectories. These systems are superior to high-dimensionality scorecards like the WEF competitiveness index, which now consists of 103 indicators (down from 114) [17].

These systems are about 25% more accurate [20] than standard high-dimensional methodologies, for example, the IMF WEO forecasts [11]. The authors' method, denoted by Velocity-SPS, consists of weighting the observed evolution of past trajectories in the GDP-Fitness space (analogues) by their distance from the current state of the country one wants to predict. The average evolution of the analogues weighted by their distance is the predicted evolution of the country. The method introduced in Tacchella [20] produces forecasting errors that are not only lower than those of the IMF growth forecasts but also uncorrelated with them. This means that the predictions from IMF and Velocity-SPS can be combined to produce an improved forecast with lower error than either original forecasts (See Table 2).

A country's capabilities development can be observed in the evolution of its outputs. The sequence of how outputs develop competitively is discoverable as a fine-grained set of pathways through which economic actors can increase the diversity and sophistication of their outputs. Outputs that require a similar set of capabilities tend to co-locate across geographies, and by leveraging data on such co-location patterns it is possible, through ML approaches, to infer a map among outputs. Economic actors are observed to move predictably in this space, diffusing towards related activities much more often, and much more successfully, than they nonlinearly progress towards far-away novelties.

The Product Progression approach uses this notion to assess the probability of actors successfully or competitively entering new activities (i.e., feasibility) and can be used as a quantitative tool to drive economic development [19]. Through the Product Progression approach, it is possible to assess the likelihood of a development project to succeed and most importantly its expected catalytic impact, in terms of new avenues for development that would be opened if success is achieved. For example, the competitive provision of hotel accommodation services opens the potential for developing management and legal skills appropriate for engaging in media services (e.g., music concerts, talent representation, etc.). This probabilistic framework allows us to steer conscious investment decisions.

Table 2 The fitness indicator based on the trade activity network topology provides more accurate forecasts for 5 years in the future of country GDP forecasts (e.g., MAE of 1.65% better than IMF's 2.09% & RMSE of 2.23% better than IMF's 2.63%). Fitness is based on trade network structure only, while IMF WEO forecasts as based on a high-dimensionality indicator approach with qualitative or judgmental inputs. Combined, these approaches would provide even more accurate forecasts

Method (+5 year GDP forecast)	Mean absolute error (%CAGR) (%)	Root mean squared error (%CAGR) (%)
IMF WEO process estimate	2.09	2.63
Velocity-SPS estimate	1.65	2.23
Combined: velocity-SPS & IMF	1.51	2.02

"Product Progression: a Machine Learning Approach to Forecasting Industrial Upgrading" outlines the method for predicting the emergence of new products in a country's export basket [1]. The aim is to predict which products a country will export several years from now. The approach involves forecasting the Revealed Comparative Advantage (RCA) which was introduced by Balassa [3]. RCA can be understood as the ratio of market shares to indicate whether a country's particular product trade share is higher than its overall international trade share – it is highly competitive in global terms for a particular traded good or service.

Multiple models were explored, but XGBoost was ultimately chosen due to its higher accuracy. In the datasets, there are strong correlations, both over time and across geographic areas, which complicate the validation of these methods. Several models with a variety of performance metrics are validated extensively in Tacchella [21].

One important takeaway from this setting is that training needs to be done, counterintuitively, on subsets of data where past examples from the country that we aim to predict have been removed. This procedure, analogous to what one would do in a cross-validation setting when performing model selection, avoids that the model learns to identify the structure of countries rather than the technological relationships between products. If the model learns to identify countries, in fact, it will generally predict that a given country has a likelihood of being competitive in a product in the future that is the same as how often it has been seen to be competitive in that product in the past. This makes the model largely unable to forecast structural changes in countries economies. By removing the target country from the training set, once presented with that country's export basket as input, the model will only be able to generalize from the learned relationships between products and is seen to be significantly better at predicting new country- product links.

6 Is the predictive ML model on a network graph lucky? Rigorously developing ML models on graph networks

Null models and signal filtering for networks are crucial instruments for data-driven approaches. Socioeconomic systems imply a mixture of information and their dynamics that one must isolate to understand important features of the structure. While dimensionality reduction aims to simplify the model by selecting the principal features that allow for a clear description of the system, null models and signal filtering aim to remove the effects of other dimensions on the data.

Describing a system as a network assumes a structure relevant to its dynamics. Understanding the significance of the structure observed in the real data is of the utmost importance. This is done by formulating and using null models. The goal is to quantify how much the presence of a certain structure or property of the network could be just a non-significant coincidence from a random configuration of the many other effects that may shape your network.

In the trade layer of FermiKnet, we wish to understand the importance of nestedness (hierarchical structures and actor cooperation and competition) in our bipartite country-product network. The two main ingredients of a null model for networks are the randomization procedure and the structural properties that should be fixed in the process.

- The randomization procedure consists of constructing, via simulation or analytically, an ensemble of random networks which represent fairly the entire ensemble without introducing bias. Once the ensemble is constructed, we can quantify how many times the structure or property we want to analyze appears in our typical random network. In the country-product example, we want to compare the nestedness of our network with the nestedness of all the possible networks in our ensemble and quantify how rare such a structure is in the ensemble.
- The constraints to fix in our random ensemble must be carefully selected. To understand the significance of our signal, we must decide which properties of our network may generate the signal we observe by chance. In the trade layer, nestedness may be a consequence of the degree distribution or the strength of the node or other factors. To exclude this case, we must constrain all our random networks to have the same degree distribution or strengths. Clearly, this can be relaxed, and we may require that these constraints are satisfied 'on average' by our ensemble. In the language of physics, this means using a canonical description instead of a model with 'hard' constraints.

In our approach, we use maximum-entropy methods from statistical physics theory that ensure, in most cases, an analytic approach, unbiased ensemble and constraints fixed in different ways. We use null models in all our network studies. For example, to determine the relationship between technology, research and output production of countries we rigorously check for random effects before establishing the strength of connections between the different layers [9], see a scheme in Fig. 3.

7 Managing trustful AI deployments

The product progression probability is estimated with an XGBOOST machine learning approach using the trade layer of FermiKNet network as input (see Albora et al. [1] for a technical description of the estimation method). There are several built-in checks and balances in a network like FermiKnet. In the trade layer, for example, reconciling product trade flows between exporters and importers provides innate error checking on the fundamental product flows. Trade flows can change due to sanctions, unexpected events like

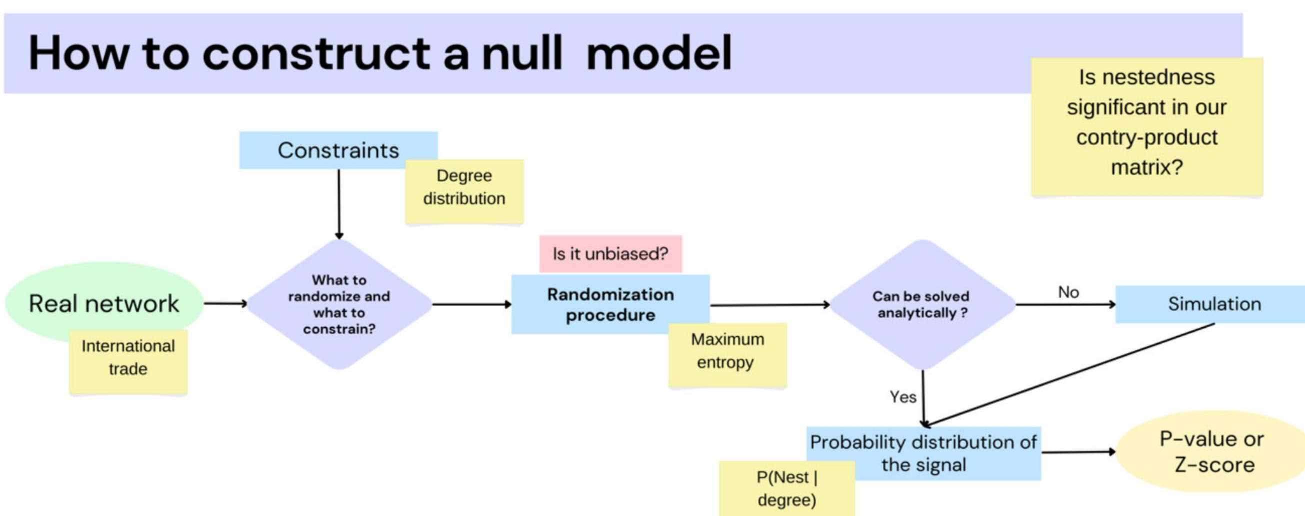

Fig. 3 Overview of null model construction methodology. This method assures that links in FermiKNet are not merely accidental or random, allows statistical estimation of link significance (p-values) in FermiKNet, and assures the reliability of product progression probabilities, thus providing granular yet robust estimates of opportunities and prospects for private sector actors

covid, etc. These changes can affect market share (revealed comparative advantage or RCA) and thus create the potential for drift in the underlying data. Overall, the network layers are robust to this type of "noise" by the nature of the network structure and constructed null models. The focus in this section is not the design process for the model, but the deployment considerations as this model is run continuously (in a typical Machine Learning Operations or MLOps environment).

If investors and businesses are to incorporate fitness and progression probabilities into their expansion strategies and board level decision-making, trust is required to ensure that the expected business outcomes are captured by the model design, implementation, and operation. Operationalizing the product progression forecasting model in a formal MLOps process increases trust because:

- It supports running anywhere via a containerized environment, which isolates and helps manage server differences, runtime library dependencies, and security.
- Based on an explicit policy, the deployment model can be run alongside several challenger models and switched in and out as needed based on formal processes and a priori-decided policy performance metrics.
- Trust is also expected during prediction. If the model is called to forecast, it must be available reliably for forecasting. The MLOps environment must also be sufficiently sophisticated to identify when the input data changes in unexpected ways (regime shift/non-stationarity caused by crises or radical policy changes, new expression of the input data – e.g., the World Economic Forum changing structure of the competitiveness index from 114 to 103 input data items in their scorecard implying that the index may not be strictly comparable through time and provides discontinuities if trend features are used in a forecasting model).
- Expect a trusted model to support its decision by providing an explanation—meaning which inputs (features) matter and the magnitude of these features' impact at the case level.

In short, trust means the model is consistent in providing a prediction, can explain its decision, and is secure. However, these are just the minimal outcomes needed to establish trust.

Automation in model design and MLOps can be seductive because of cost savings. It is critical to resist this impulse and ensure a qualified human is put in the loop appropriately, creating trust. Human in the loop means that someone who knows the real-world objective and understands the model's sensitivity to changes in the operating environment is specifically identified. The danger of not having a dedicated person is that there is no one to be accountable and no one to fix issues that will inevitably arise. While strong documentation and MLOps systems can mitigate some staff turnover risk, formal MLOps risk assessments should be implemented to underline the importance of the human role.

The trust process is enabled through the six-stage cycle in Fig. 4. The main objective of trust is that the end user feels confident that the data, model, and processes are credible because they are defensible, explainable, well-documented, and operated by knowledgeable professionals. This results in achieving key characteristics of trust: fairness, stability, robustness, computability guarantees, reliability, replicability, and interpretability.

First, we must have a real-world problem that can be solved using machine learning. This requires sophisticated staff who understand business (and international development) strategy and dynamics, and have a considerable understanding of machine learning architectures, risks, and trade-offs. Traditionally this has been satisfied by requirements engineering intermediaries working with business, modelers, and information technology staff. However, this approach regularly leads to implementations that suffer "just missing" the real-world purpose. Contemporary requirements intermediaries are primarily unaware of ML architectures that capture dynamics in knowledge graphs and how to deploy them for decision-making rather than typical reporting.

Second, the data source must be fit for purpose and reliable and have a set of processes to clean, transform, and situate it in the network. For example, the UN-COMTRADE database is a source for the product trade layer of FermiKnet and is updated eight times per day. While the data are updated at regular intervals, there is also one annual reconciled "fixing." Network data integrations and machine learning models are built for both continuous updates and the annual fixing using different data cleaning rules and error handling. Different trust criteria are valid for the annual fixing data. Decisions must be made on how trust is defined for continuously updated data which suffers from country lags due to different customs processes and inconsistent quality due to data management and local reporting procedures.

Third, the candidate ML model must be designed with a deep understanding of the objective. For example, a progression probability is intended to help identify expansion feasibility and could be implemented through a random forest or a gradient boost or XGBoost method. A random forest is a type of resampling learner (bootstrap + random features), while XGBoost is similar but sequential and thus better suited to high-dimensional, large and imbalanced datasets due to innate feature selection mechanism and ability to predict accurately even when features are interdependent.

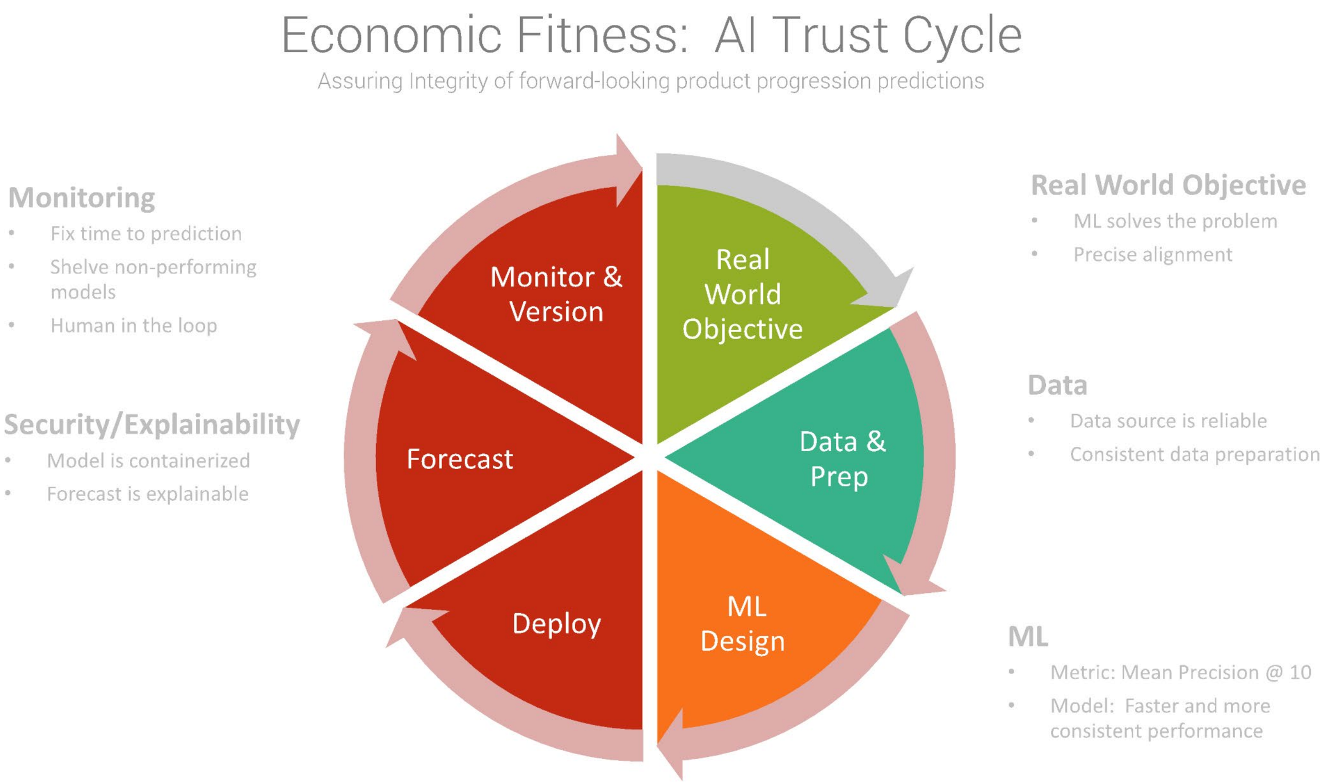

Fig. 4 Product progression and other economic fitness estimates are deployed via a six-step trust cycle. The first three steps (Real world objective through ML Design) are considered the realm of design. The real-world objective must be defined sufficiently to capture appropriate data, identify valuable features to the prediction task, and parameterize ML models. Once tested and approved as part of the model design, the MLOps phases (red-shaded segments) become critical for trust

Fourth, the candidate model must be deployed in a stable and secure environment – a MLOps environment (e.g., BreakLine, Iguazio and Mlflow, among others) which aids in achieving reliability and replicability across compute resources.

Fifth, we expect the model to explain its predictions in terms of accuracy and precision levels and in the evolution of country capabilities.

The sixth step in the cycle requires proactive monitoring and versioning. However, due to the complexity of the data we implemented a yearly continuous integration and deployment (CI/CD) process for the model, as shown in the figure above. The process involves ingesting the most up-to-date UN-COMTRADE data, cleaning and adding it to the previously cleaned data, retraining a new ensemble model with the updated data, and deploying it. At this stage, both the new and old models perform a forecast, and all models are saved (a step called versioning). We also monitor the old models and compare their forecasts to the new model's forecasts. This process is repeated for each new year of data and aids in achieving replicability.

By comparing the predictions of newer models to those of previous models, we can gain insight into the divergence in prediction and indirectly understand how networks may be changing over time. This approach can help us comprehend the complexity of networks and improve our forecasting approach. Therefore, for us, MLOps is about ensuring smooth production and gaining a deeper understanding of network change and prediction. Since we are forecasting for a 5-year period, we can only fully assess the model's performance once that time has elapsed, although a partial, rolling evaluation is possible even before. This step is critical in establishing trust and demonstrating the relevance of our work. For instance, once the 2023 UN-COMTRADE data is available, we can use it to confirm the performance of our model by comparing it to the mean Precision @10 of our forecast using data up to 2018 UN-COMTRADE. By monitoring the model's performance, we can also detect when it starts to drift and decide when to shelve it. Additionally, this approach may help us indirectly understand data drift and potentially lead to new discoveries.

7.1 Overview of ethical considerations

Several levels of ethical considerations were considered and implemented in the use of FermiKnet and fitness indicators for impact investing. These include (i) the formulation and expression of integration and competitiveness impact claims including both multigenerational and beneficiary perspectives, (ii) building, training and evaluating FermiKNet and fitness-based predictors, and (iii) deployment, monitoring and governance. Table 2 summarizes the ethical considerations for all three phases of the workflow.

Workflow stage	Ethical purpose Will AI actions be a net good to society?	Fairness Is AI discriminating on sensitive features (e.g., affecting inclusion)?	Transparency & Interpretability Is there sufficient disclosure to stakeholders in the models?	Accountability Are standards of governance applied based on risk?
Real World Objective (impactful investment)	Impacts are expected to increase affordability and access	Defined in theories of change	Fitness indicators are well defined, academically published, and code is available for inspection in public repositories	Impact investment ratings affected by fitness indicators are governed by written guidelines
Data & Prep	Identifiable personal data is not used	Network links are based on actual activity among entities. Prep does not exclude any entity from sources.	Preprocessing pipelines are well defined and documented and available to stakeholders upon request; Imputed data (smoothed) data is identified and its impact assessed with formal statistics	Data is governed by the sources (e.g., UN COMTRADE)
Network Estimation	A network is constructed to preserve diversity	Network structure is tested using random network null models which preserve the hierarchal and competitive structure of the empirical network	Network models and link structures are tested with relevant p-values computed	Statistical tests are defined and performed on the network. Results are made available for compliance purposes.
Predictive model development (build, train, evaluate)	ML models are designed for use as predictors/ causal effects	No protected features are present in the network (e.g., country income category is not used)	Results of performance, stability, and sensitivity tests are documented (e.g., via cross-validation)	Model development is governed through several artefacts (model documentation, model assumptions, alternativesconsidered, variable selection, performance)
Deployment	Model predictions are not directly traded. They are used to set impact expectations	The system is monitored for drift	Changes to the outputs (e.g., recalibration) are well documented and changes between version are identified and potentially re-rated (e.g., population stability reporting)	Version changes or substantial changes to the ways in which model outputs are used require changes to methodology and operational guidelines

Workflow stage | Ethical purpose Will AI actions be a net good to society? | Fairness Is AI discriminating on sensitive features (e.g., affecting inclusion)? | Transparency & Interpretability Is there sufficient disclosure to stakeholders in the models? | Accountability Are standards of governance applied based on risk?

8 Conclusion

Machine Learning and AI techniques are no longer only applied to financial time-series prediction of multivariate data to profit from momentary arbitrage opportunities. Increasingly AI is needed to learn transformative dynamics about global networks of economic activity to identify opportunities for expansion projects amenable to private sector actors which deliver on the 5Ps (People, Prosperity, Planet, Peace, and Partnership), and especially are likely to contribute positive impact towards reaching the SDGs. This paper has introduced one example of product progression probability estimation on a knowledge network of traded production, constituting one layer of FermiKNet (traded products, skills, technologies, research, and other aspects of economic growth). It highlights how risk is quantified (through network null model generation), and models are managed through MLOps and human-in-the-loop processes to provide trustworthy estimators for use in board rooms and corporate decision-making.

Appendix: Computation of fitness and complexity metrics

In this section we define mathematically the Fitness (F) and Complexity (Q) that have been introduced and explained in this paper. These definitions where first introduced in Tacchella [21] Let q_{cp} be the total export in US$ of country c in product p. We can define the Revealed Comparative Advantage (RCA) following [3] as $RCA_{cp} = \frac{q_{cp}}{\sum_{c}^{C} q_{cp}} \frac{\sum_{p}^{P} q_{cp}}{\sum_{c,p}^{C,P} q_{cp}}$ (where P and C are the set of products and countries respectively). In the literature a value of RCA greater than 1 is assumed to be the signal of a country having competitive capabilities with respect to other countries that enable the export of a given product. Therefore we define the elements of the matrix M_{cp} to be 1 whenever $RCA_{cp} \geq 1$ and 0 otherwise. The computation of the Fitness and Complexity measures is performed iteratively, as a function of the M_{cp} matrix. The equations are defined as

$$F_c^{(n)} = \sum_{p}^{P} M_{cp} Q_p^{(n-1)}$$

$$Q_p^{(n)} = \frac{1}{\sum_{\{c}^{C} \frac{1}{F_c^{(n-1)}}}$$

At each step the values are normalized such that the average over the countries of the values of F and the average over the product of the values of Q is always equal to 1. This coupled map always converges exponentially to a unique fixed point, regardless of the starting values for $n=0$, that can be chosen arbitrarily.

The meaning of F is therefore that of a “complexity-weighted diversification”, I.e. the more a country is diversified (sum over the rows of the M matrix) and the more complex are the products in which it is diversified, the higher the Fitness. The meaning of Q, I.e. the complexity of products, is slightly less obvious: the nonlinearity imposes a lower bound on the complexity, that is defined as the Fitness of the less fit country that is able to have RCA>1 in that product.

Acknowledgements D.M. and A.T. acknowledge the financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 104 published on 2.2.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union – NextGenerationEU – Project Title “WECARE – WEaving Complexity And the gReen Economy” – CUP 20223W2JKJ by the Italian Ministry of Ministry of University and Research (MUR). D.M. and A.T. acknowledge the financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 1409 published on 14.9.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union – Next Generation EU – Project Title “Triple T – Tackling a just Twin Transition: a complexity approach to the geography of capabilities, labour markets and inequalities” – CUP F53D23010800001 - Grant Assignment Decree No. 1378 adopted on 01.09.2023 by the Italian Ministry of Ministry of University and Research (MUR).

Declarations

Conflict of interest The authors declare no conflict of interest.

Disclaimer The views and opinions expressed in this work are those of the authors and do not necessarily reflect the views of the International Finance Corporation, the World Bank Group, its Board of Directors, or the governments they represent. The International Finance Corporation, and the World Bank do not guarantee the accuracy of data included in this work. The boundaries, colors, denominations, and other information shown on any map in this work do not imply any judgement on the part of the International Finance Corporation or the World Bank concerning the legal status of any territory or the endorsement or acceptance of such boundaries. The authors are responsible for any errors or omissions.

References

1. Albora, G., Pietronero, L., Tacchella, A., Zaccaria, A.: Product Progression: A Machine Learning Approach to Forecasting Industrial Upgrading (2021). arXiv:2105.15018v1

2. Athey, S., Imbens, G.W.: The state of applied econometrics: causality and policy evaluation. J. Econ. Perspect. **31**(2), 3–32 (2017)
3. Balassa, B., Noland, M.: ``Revealed'' comparative advantage in Japan and the United States. J. Int. Econ. Integr. **4**, 8–22 (1989)
4. Das, S.R.: Text and context: language analytics in finance. Facul. Book Gallery **54**, 66 (2014)
5. Deboeck, G.J.: Trading on the Edge: Neural, Genetic, and Fuzzy Systems for Chaotic Financial Markets. Wiley, New York (1994)
6. Dixon, M.F., Bilokon, P., Halperin, I.: Machine Learning in Finance: From Theory to Practice. Springer, Germany (2020)
7. Dosi, G., Mathew, N., Pugliese, E.: What a firm produces matters: processes of diversification, coherence and performances of Indian manufacturing firms. Res. Policy **51**, 104152 (2020). https://doi.org/10.1016/j.respol.2020.104152
8. Feldman, R., Sanger, J.: The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data. Cambridge University Press, Cambridge (2006)
9. Formichini, M., Cimini, G., Pugliese, E., Gabrielli, A.: Influence of technological innovations on industrial production: a motif analysis on the multilayer network. Entropy **21**, 126 (2019)
10. Fotopoulou, E., Mandilara, I., Zafeiropoulos, A., Laspidou, C., Adamos, G., Koundouri, P., Papavassiliou, S.: SustainGraph: a knowledge graph for tracking the progress and the interlinking among the sustainable development goals' targets. Front. Environ. Sci. **10**, 66 (2022)
11. Genberg, H., Martinez, A., Salemi, M.: The IMF/WEO Forecast Process. IEO Background Paper No. BP/14/03 (Washington: Independent Evaluation Office of the IMF) (2014)
12. Goodell, J.W., Kumar, S., Lim, W.M., Pattnaik, D.: Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. J. Behav. Exp. Finance **32**, 66 (2021)
13. Head, M.L., Holman, L., Lanfear, R., Kahn, A.T., Jennions, M.D.: The extent and consequences of P-hacking in science. PLoS Biol. **13**(3), e1002106 (2015)
14. Mullainathan, S., Spiess, J.: Machine learning: an applied econometric approach. J. Econ. Perspect. **31**(2), 87–106 (2017)
15. Nickel, M., Murphy, K., Tresp, V., Gabrilovic, E.: A review of relational machine learning for knowledge graphs. Proc. IEEE **104**(1), 11–33 (2016)
16. OECD: Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers (2021)
17. Olczyk, M., Kuc-Czarnecka, M., Saltelli, A.: Changes in the Global Competitiveness Index 4.0 Methodology: the improved approach of competitiveness benchmarking. J. Compet. **14**(1), 118–135 (2022). https://doi.org/10.7441/joc.2022.01.07
18. Pugliese, E., Cimini, G., Patelli, A., Zaccaria, A., Pietronero, L., Gabrielli, A.: Unfolding the innovation system for the development of countries: coevolution of science, technology and production. Sci. Rep. **9**(1), 16440 (2019)
19. Roster, K., Harrington, L., Cader, M.: Country case studies in economic fitness: Mexico and Brazil. Entropy **20**(10), 753 (2018)
20. Tacchella, T., Andrea, D.M., Pietronero, L.: A dynamical systems approach to gross domestic product forecasting. Nat. Phys. **14**(8), 861–865 (2018)
21. Tacchella, A, et al.: Relatedness in the era of machine learning. arXiv preprint arXiv:2103.06017 (2021).
22. United Nations: Transforming our World: The 2030 Agenda for Sustainable Development. United Nations, New York (2015)
23. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S.D., Tegmark, M., Fuso, N.F.: The role of artificial intelligence in achieving the sustainable development goals. Nat. Commun. **11**(1), 233 (2020)
24. Woloszko, N.: Adaptive Trees, a Novel Method for Economic Forecasting [Conference Presentation]. OECD New Approaches to Economic Challenges, Paris (2019)
25. Yeh, C., Perez, A., Driscoll, A., et al.: Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. Nat. Commun. **11**, 2583 (2020)
26. Yeh, C., Meng, C., Wang, S., Driscoll, A., Rozi, E., Liu, P., et al.: Sustainbench: Benchmarks for monitoring the sustainable development goals with machine learning (2021). arXiv preprint arXiv:2111.04724