Irving Fisher Committee on
Central Bank Statistics

BIS

International Conference on "Statistics for Sustainable Finance", co-organised with the Banque de France and the Deutsche Bundesbank

14-15 September 2021, Paris, France, hybrid format

Accelerated Data Science, AI and GeoAI for Sustainable Finance in Central Banking and Supervision[1]

Jochen Papenbrock, NVIDIA GmbH, Germany; John Ashley, NVIDIA Corp., USA;
and Peter Schwendner, Zurich University of Applied Sciences

[1] This presentation was prepared for the conference. The views expressed are those of the authors and do not necessarily reflect the views of the BIS, the IFC or the central banks and other institutions represented at the event.

Accelerated Data Science, AI and GeoAI for Sustainable Finance in Central Banking and Supervision

Jochen Papenbrock, John Ashley and Peter Schwendner[1]

Abstract

Sustainable finance – the umbrella term for investing based on Environmental, Social, and Governance (ESG) considerations – requires the analysis and integration of massive amounts of complex data – including but not limited to geospatial terrain and elevation data, climate data, company news, market data, regulatory filings, supply chains & logistics. Traditional analytic techniques will not suffice. Here, we lay out the argument for computing platforms and strategies derived from scientific and industrial high-performance computing (HPC) and the emerging field of Artificial Intelligence (AI). We highlight some of the critical considerations in building such systems to support the data analytics and, equally importantly, the controls required to build supervisory systems for sustainable finance. Lessons learned from leading HPC, and AI systems point to GPU accelerated compute as being a key feature of these systems, along with data-centric system design. We also highlight the role of explainable and trustworthy AI as catalysts for accelerated computing and their challenges. Finally, we propose a concrete use case: a cross-regional, shared system that integrates geospatial data with economic and market data to illustrate the benefits of accelerated computing in building these systems.

Keywords: spatial finance, environmental exposures, sustainable finance, central banks, data processing, explainable AI, RAPIDS, ESG

JEL classification: Q58; Q57; Q56; E58; G28; C31; C81

[1] Respectively, Head of Financial Technology – EMEA, NVIDIA GmbH, Germany (jpapenbrock@nvidia.com); General Manager Financial Services and Technology, NVIDIA Corp., USA (jashley@nvidia.com); Head of Institute of Wealth and Asset Management, Zurich University of Applied Sciences (scwp@zhaw.ch).

Contents

Introduction and Overview

Sustainability and Artificial Intelligence (AI) are two of the megatrends that will shape the development of the financial services and the supervisory sector for the near future.

Sustainable finance considers environmental, social, and corporate governance factors when making investment and underwriting decisions. These factors are, in many cases, both qualitative and quantitative; combining this information into a decision-making framework is complex, and best practices are still evolving. While some of the data needed for this is available via standard, machine-friendly sources, a considerable proportion is uncurated and unstructured – in text, images, audio, or video. Artificial Intelligence is the subject of much debate in both definition and execution. Here, we use it to mean systems that learn how to do tasks by repeated training with relevant data, especially in cases where that data may be unstructured (text, speech, images, or video). It is compelling in uncovering and acting on complex patterns in and among datasets.

Collecting sustainable finance data and closing the data gaps will be crucial in implementing data-driven, evidence-based policymaking in areas like micro-prudential supervision, financial stability, macroeconomic analysis, and risk and reserve management. The amount, complexity and innovative character of the required data will pose enormous challenges as the data characteristics need to be aligned with the policy requirements.

Assessing the appropriate level of protection against sustainability risks is a challenge for central banks. Climate risk is the area that gained the most attention during the last 20 years. Public scrutiny is substantial: a legal opinion commissioned by Greenpeace Germany (Verheyen 2021) identifies a liability of ECB and Bundesbank to use monetary policy for sustainability objectives. Dafermos et al. (2021) suggest three policy scenarios to address the climate footprint of bonds in the ECB collateral portfolio.

Despite the massive impacts, the marginal contributions to climate risks are intrinsically difficult to measure accurately, notably due to the radical uncertainty that characterizes climate risks in the form of tipping points, non-linearities and regime shifts. Climate-related financial risks may directly impact central banks' counterparties and financial assets used in monetary policy operations and collateral management.

Beyond climate risks, the commonly accepted sustainable development goals (SDGs) consider environmental risks relating to water quality, deforestation, land use and biodiversity loss. In contrast to climate risks, the impacts of ecological footprints of companies happen in a local context and can have far-reaching consequences along with the supply and value chain. Upcoming regulations like extensions to the Taxonomy for Sustainable Finance will aim for the more forward-looking "net zero loss" instead of the current "do not harm" objective.

Existing sustainability ratings from different providers diverge significantly for the same investable asset (Berg et al., 2019). This implied risk for greenwashing. In position papers and public consultations, IOSCO (2021), FCA (2021) and BaFin (2020) already addressed this risk. A naive solution would be to select just one specific data source for each of the environmental (E), social (S) and governance (G) exposures across investable assets and rely entirely on this assessment. Pedersen et al. (2020) build an ESG-efficient frontier and argue which balances investors with different preferences between Sharpe ratios, and ESG scores will take. They also find a positive risk premium for a governance factor, no significant risk premium for a carbon-based environmental factor and a negative risk premium for a social factor. However, this single selection of ESG data would introduce substantial model risk, which might be inappropriate for more significant exposures, especially for public institutions with a complex structure of stakeholders.

An example in the literature for different interpretations even of the same ESG ratings is (Giese et al. 2021). They find MSCI ESG data to enable investment outperformance, especially for governance ratings and social and environmental ratings. However, (Bruno et al. 2021) find no alpha in the same data after adjusting for equity factor exposures: they attribute 75% of the outperformance to the standard "quality" equity factor that reflects simple financial statements information. They also point out a substantial tech sector bias in ESG ratings and suggest that ESG performance's main drivers are strong tech sector performance and increasing investor attention. These conflicting findings prompted a letter to the FT (Edmans, 6.5.2021) calling for more scientific rigour in ESG evaluation studies.

Therefore, portfolio managers should consider the full set of information to find the optimal portfolio that incorporates all external environmental, social and governance exposures and costs after evaluating diverse and conflicting data sources. From a theoretical viewpoint, the optimal portfolio position would solve a Bayes-optimal classification problem across an ensemble of different models and objectives. To implement such a solution, model ensemble techniques like stacking (Spilak and Härdle 2020) are promising for incorporating information from various sources.

Obviously, for these purposes, standard numerical, tabular and/or time-series data sets need to be combined with alternative data sets, mainly derived by remote sensing (Damm et al., 2020, Santos and Decker, 2020 and Schneider et al., 2017) with a mix of textual and tabular economic and financial data, along with network-based (e.g., transactions, supply chain and geospatial data) information. Patent data analysis can be used to detect investment opportunities in specific sustainable technologies (Guderian et al. 2021).

Even more challenging will be the merging of sustainable data sets with the existing ones used in central bank use cases like business cycle analysis, risk indicators and financial risk assessments, the behaviour of market participants, monitoring financial transactions and capital flows, surveillance exercises, market abuse, flagging misconduct and fraud, assessing system-wide risks. Modern AI tools and techniques can combine these datasets in a technical sense; exactly which datasets to connect and how to interpret the results are open questions.

Sorensen et al. (2021) point out the advantages of a quantitative treatment for ESG investing. They recognize the ESG rating divergence documented by Berg et al. (2019) as an opportunity for active investing. Furthermore, they argue for a quant-based expert system with four advantages relative to a discretionary investment process: better integration with MPT, better integration of big data using natural language processing (NLP), better systemizing into domain knowledge and better replication with digital estimation. Chen et al. (2021) warns of a bias in ESG ratings: companies that publish more ESG data tend to receive better ratings, so investors need to analyze the details of the rating process before including the resulting ratings in an automated investment process. Sokolov et al. (2021) propose an approach that combines modern machine learning techniques in NLP with portfolio optimization to incorporate views of companies' ESG performance. Martellini and Vallée (2021) discuss measuring and managing ESG risks in sovereign bond portfolios and implications for sovereign debt investing.

In implementing policies, monitoring tools, and stress scenarios in individual financial institutions and supervisors, one of the main challenges of this topic area will become acute: The availability and analysis of reliable data.

The collection and evaluation of this data and the establishment of deep insights and links to future policies can only succeed with massive use of new and already established technologies. This is where the two megatrends, Sustainable Finance and AI, that will shape the financial sector in the coming years overlap.

AI computing platforms will play a crucial role here as they are the key technology enabler to implement sustainable finance using AI, big data processing and effective data filtering and data science visualization. Such platforms can enable automation to process the required sustainability data, make insights, and model explanations available to all relevant parties.

Providers of such platforms and the technology companies building solutions on top of them will be the future trusted advisors and key technology providers on the journey to more data and AI-driven central banking, policymaking, and supervision.

These computing platforms need to be designed holistically, providing a scalable path to support growth in data storage, movement, and compute demand while meeting necessary security requirements., New sustainability data always need to be added, usually under high time pressure, requiring new processes and new tools to be integrated into the environment. Otherwise, there is a significant risk of an emerging patchwork of individual solutions. A fragmented IT landscape would not scale sufficiently and therefore would not meet the organization's long-term goals.

Establishing such computing platforms must start from a well-designed base, but evolve in parallel, going hand in hand with the data collection and curation process. Several factors drive this need for a co-evolutionary approach:

1. Data, IT systems, operations and policies must be fully aligned and integrated. As all eventual data sources and data quality and curation processes cannot be known on day one, the architecture must be designed with a degree of flexibility and an expectation that it will evolve over time. Ideally, this evolutionary process will be agile and scalable. There is an adaptive, auditable process to evolving data quantity and quality in which modern AI supports and automates the model building and updating steps, and additional data improves other models. A key issue is to ensure that nowcasts and predictions based on big data are accurate, auditable and interpretable, helping to identify specific explanatory causes or factors. These audibility and interpretability requirements are challenging to retrofit into models and architectures, and again, will need to evolve with the system. The potential to evolve is crucial for the communication and building up trust with the external stakeholders.
2. Visualization of data, models, and tools to create narratives around the models and their results will be increasingly required in direct proportion to the amount of data, variety of datasets, and complexity of the models being deployed. Large scale visualization, clustering and network analysis can support the explorative, visual inspection of substantial amounts of complex data, helping human understanding of complex systems and results.
3. Compute-intensive AI/ML techniques can help improve the data quality, detect outliers and bias in existing datasets. It can also help fill gaps, identify fewer valuable data with a low signal-to-noise ratio or point to data that should potentially receive more attention. An emerging and critical data science field is using AI to create "realistic" synthetic data. This has numerous motivations, like enriching existing datasets, filling gaps, creating adverse and stress data sets, and amplifying models. It can also sometimes be used to enhance explainability. Synthetic data can be calibrated with real data sets, and there are parameters to generate additional data that lies outside of the original distributions but still exhibits similar statistical properties and stylized facts as the original data. Agent-based simulation, using procedural and AI agents, can also be used where underlying assumptions and properties of the data generating processes suggest such an approach.
4. High-Performance Computing (HPC) incorporates general system design principles and platforms allowing organizations to run large-scale mathematical models and simulations. These simulations can include environmental, economic, and financial models to get more insights and understand the relationships. HPC also helps to solve complex multi-objective optimization problems to find the optimal policies under constraints.
5. The size of data and number of sources will increase, and so will the number, complexity, and frequency of updates of models. Storage, networking, compute, and operations will need to evolve with this growth.

There are human and environmental factors involved as well. Data scientists and technical personnel who can develop and manage these systems are expensive and in short supply, so it is also essential to build systems and processes that are efficient with their time. Large computing systems can consume significant amounts of energy, so any system should be as energy efficient as possible.

A modern AI computing platform needs to address the factors identified above. The recent draft of the AI Act of the European Commission characterizes computing platform by the following points:

- Integrated hardware and software stack
- Build and manage robust models at scale in an efficient way
- Include the human user in the process: simplified, transparent, explainable
- Enable strategies for AI model maintenance, risk management and incident management.

Fortunately, modern HPC and AI computing systems address both the European Commission's points and the factors we have identified above. An elementary building block in AI and high-performance computing platforms (HPC) are GPUs – "Graphics processing units". According to the IFC-BIS publication "Computing platforms for big data analytics and artificial intelligence" (see Bruno et al. 2020) ", Central banks' experience shows that HPC platforms are primarily developed to ensure that computing resources are used in the most efficient way, so that analytical processes can be completed as rapidly as possible. [...] A processor core (or "core") is a single processing unit. Today's computers – or CPUs (central processing units) – have multiple processing units, with each of these cores able to focus on a different task. Depending on the analytical or statistical problem at hand, clusters of GPUs (graphics processing units, which have a highly parallel structure and were initially designed for efficient image processing) might also be embedded in computers, for instance, to support mass calculations." Today the superb computing power of GPU clusters is widely used in research where most of the supercomputers are powered with GPUs. Industrial firms and other research organizations have long since adopted GPU computing to address high-performance computing requirements. (e.g. Cambridge 1).

The efficiency of such computing platforms is often measured in floating-point operations per unit of energy consumed. This measure is significant when supercomputing is used in support of Sustainable Finance. GPU-based super-computers lead the global top500 HPC list[2], and many regional GPU-powered leadership class supercomputers are considered 'greenest'[3].

The use case presented in this paper covers alternative, non-tabular data and advanced technologies that will be necessary to address Sustainable Finance and the more traditional statistics and data collection initiatives already underway or currently emerging. We cover approaches that are a mixture of Earth Science and especially geospatial/geostatistical data, geographic information systems (GIS), big spatial data, climate science, scalable high-performance computing, machine learning, artificial Intelligence, deep learning plus professional rendering and visualization, and massive simulation. These techniques are often referenced in the literature as AI-powered systems or platforms delivering on-demand, personalized and actionable Intelligence for enterprises, governments, regulators, and supervisors, using approaches like GeoAI, Climate Intelligence, Geo Knowledge Discovery, and Earth Science AI. Spatial science offers tools and technologies that enable us to understand, analyse, and visualize real-world phenomena according to their locations. Such systems process, analyse and synthesize petabytes of geo-related, sensor-based data to create a unified perspective on critical issues like climate risk at global and individual levels.

(Geo-)Spatial finance offers socio-economic and environmental insights that have the potential to enhance data transparency in the financial system. Sustainability related risks can be better managed when models include geospatial data. It can also assist in the analysis and management of other factors affecting risk and return in different asset classes or central bank collaterals and to support a transition to sustainable development and to monitor environmental and economic activity day and night. These technologies will generate insights and enable us to tackle a wide variety of global challenges in new ways and to forecast the impact of climate change and to respond

[2] https://top500.org/lists/green500/

[3] https://www.hpcwire.com/off-the-wire/meluxina-named-eus-greenest-supercomputer/

to societal challenges. In the financial domain this means that (spatially located) asset owners can verify their investments, asset managers can engage with their investment objects, corporates can verify internal data collection, compare performance with peers, and understand environmental risks and impacts within their supply chains. Regulators will be able to assess environmental and social systemic risks more accurately within the financial system and policymakers will be able to better track progress.

The selection and orchestration of the underlying IT platforms, processing units, and distributed, parallelized computing and algorithms will be crucial to process this kind and quantity of data and to support the policy and decision-making process. Virtualization will allow for faster computing to be accessible from any device anywhere and for remote collaboration. Advanced visualization and rendering allow us to create 3D/4D experiences in support of these efforts. Today it is already possible to create physics-aware, true-to-reality, physics-based simulations with digital landscapes and earth twins, realizing the full technical potential to address spatial, sustainable finance and full technical support for policymaking. Our paper is dedicated to outline and and describe these data and technologies which are already available today and how Sustainable Finance will benefit from them.

The structure of the paper is as follows:

- We cover “Tools, Platforms and Communities to access and develop Sustainability Data” and "Computing Platforms for SupTech, AI Explainability and Compliance".
- Then we present a specific use case for computing platforms in sustainable finance: "Empowering Regions regarding their Environmental Exposures"
- We also attach an Appendix: GPU-Accelerated Data Science and AI Computing Platforms at Central Banks.

Tools, Platforms and Communities to access and develop Environmental Sustainability Data

From the perspective of asset owners, responsibility for getting access to the necessary data for sustainable investing is clearly with their asset managers. Most ESG data providers primarily see asset managers, banks and real-economy corporates as their clients. Beyond the large commercial data providers, specialized technology companies, NGOs and the scientific community also offer platforms and tools. Innovative governments explicitly support fintech companies working on ESG data solutions and connect them with scientific institutions as they recognize the opportunities not only for the environment and for better portfolios, but also specifically for commercial innovation. Examples for state-funded initiatives are the Green Fintech Network[4] assisted by the Swiss State Secretariat for International Finance (SIF) and the Networking Event Series – Sustainable Finance Technology[5] organized by Zurich University of Applied Sciences (ZHAW) and Swiss Sustainable Finance (SSF) and supported by Innosuisse. The Business@Biodiversity[6] platform organized by the European Union is an example of a state-funded community to scale data access and awareness for corporate users. ENCORE[7] is developed by UNEP-WCMC in partnership with the Natural Capital Finance Alliance and funded by the Swiss government and the MAVA foundation. It shows how business sectors depend on ecosystem services and natural capital assets. Currently, it is expanded to also include the impacts on businesses on biodiversity and enables an impact assessment of corporate and financial investments. WWF Sight[8] is a global platform tool including spatial datasets, satellite data, information about species and industrial assets. Restor[9] is a new community-based data platform set up in 2021 specifically for ecosystem restoration projects. It is based on research from the Crowther Lab of ETH. IBAT[10] (Integrated Biodiversity Assessment Tool) is developed by a consortium of BirdLife, Conservation International, IUCN and UN-WCMC. It offers both reports and GIS data downloads based on a database of protected areas and protected and endangered species. IBAT includes the Species Threat Abatement and Restoration (STAR) metric. STAR documents the contribution of specific conservation and restoration actions in specific places by businesses, governments, civil society, and other actors towards global goals for halting extinctions. The geoFootprint[11] initiative launched by Quantis in collaboration with 25 other partner institutions is focused on the life-cycle environmental impact of agricultural crop production. The independent Responsible Mining Foundation (RMF) is focused on ESG impacts of the extractive industry, but announced it will close at the end of 2022[12]. Satelligence is focused on agricultural and industrial impacts on deforestation as this is measurable by remote sensing in high quality. CDC Biodiversité together with Carbon 4 Finance form a French joint venture to offer biodiversity

---

[4] https://www.sif.admin.ch/sif/en/home/dokumentation/fokus/green-fintech-action-plan.html

[5] https://www.zhaw.ch/en/sml/institute-zentren/iwa/veranstaltungen/translate-to-english-networking-event-series-sustainable-finance-technology/

[6] https://ec.europa.eu/environment/biodiversity/business/workstreams/methods/index_en.htm

[7] https://encore.naturalcapital.finance

[8] https://wwf-sight.org/

[9] https://restor.eco

[10] https://www.ibat-alliance.org

[11] https://quantis-intl.com/strategy/collaborative-initiatives/geofootprint

[12] https://www.responsibleminingfoundation.org/2022-status-announcement/

exposure assessments for financial sector clients. EconSight[13] analyses technology innovations using ML applied to patent data and therefore contributes to detecting interesting investment targets also in the sustainable technology space. RepRisk and Truevalue labs are established sophisticated ESG data platforms running large-scale news-based ML analytics.

Establishing forward thinking and efficient market structures on the implementation of sustainable financing strategies is a goal of the 'Green and Sustainable Finance Cluster Germany"[14] which is also part of the FC4S Network[15] where green finance is a central theme of the agenda, including the question of "how could financial centres contribute to the delivery of the Sustainable Development Goals and the Paris Agreement?". The Impact Festival[16] is a platform for sustainable technologies and innovations.
The Financial Big Data Cluster (FBDC) is the use case in the financial domain of European GAIA-X initiative. The FBDC represents the central use case for the financial sector. Closely linked to the establishment of the cluster is the research and development project "safeFBDC." During the safeFBDC project, data sets will be used to create five applications that will be researched, developed and prototypically validated using artificial intelligence and machine learning. One of the applications is on Sustainable Finance. The aim is to improve the currently often insufficient availability and quality of ESG data and to establish a local ESG data platform. The central approach is to test and further develop innovative AI and ML methods both to close identified ESG data gaps and to develop new methods for ESG data generation. Financial market participants, especially banks, should thus be able to better integrate sustainability risks into their operational risk modeling and relevant decision-making processes[17]. A related project is the Financial AI Cluster (FAIC[18]) addressing one of the central topics in the Finance & Insurance Data Space of GAIA-X, namely the implementation of trustworthy, explainable AI, AI Assurance and AI Audit based on computing platforms.
The Spatial Finance Initiative[19] has been established by the Alan Turing Institute, Satellite Applications Catapult, and the Oxford Sustainable Finance Programme to bring together research capabilities in space, data science and financial services and make them greater than the sum of their parts. 'Spatial finance' is the integration of geospatial data and analysis into financial theory and practice[20]. Earth observation and remote sensing combined with machine learning have the potential to transform the availability of information in our financial system. It will allow financial markets to better measure and manage climate-related risks, as well as a vast range of other factors that affect risk and return in different asset classes.
The BRIDGE Discovery project 'Spatial Sustainable Finance'[21] run by Zurich University of Applied Sciences and University of Zurich aims to set a global rating standard that enables financial institutions and companies to reduce their water and biodiversity footprint.

[13] https://www.econsight.ch

[14] https://gsfc-germany.com/en/

[15] https://www.fc4s.org/about/

[16] https://impact-festival.earth/

[17] https://www.fs-unep-centre.org/project/financial-big-data-cluster-fbdc-use-case-sustainable-finance-susfi/

[18] https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Blog/2020-12-16-the-innovation-potential-of-gaia-x-technologies-for-the-trusted-use-of-artificial-intelligence-in-the-finance-sector.html

[19] https://www.cgfi.ac.uk/spatial-finance-initiative/

[20] https://www.refinitiv.com/perspectives/ai-digitalization/using-spatial-finance-for-sustainable-development/

[21] https://www.spatial-sustainable-finance.ch

The Future of Sustainable Data Alliance[22] is looking to answer the question "What data do investors and governments need to deploy capital sustainably and in line with the requirements of regulators, citizens and the market now and in the future?" GeoWorks[23] is Southeast Asia's first geospatial industry centre.

We also see the first combined platforms for banks and insurers that are both dedicated to socially responsible investment (SRI) and compliant with the principles of transparent artificial intelligence as set out in the proposed European regulation[24], respecting the European principles of autonomy, interpretability, explicability, transparency, responsibility and robustness.. The platform combines algorithmic know-how with expertise in data management and decarbonization. Based on artificial intelligence and deep learning, the solution leverages financial and extra-financial data, including accurate and standardized ESG data. The attention paid to data integrity, combined with advanced management tools, allows the design, large-scale deployment and management of transparent machine learning models. Additionally, the models can be manually modified to take into account company policies and regulatory obligations (e.g. the European regulation on green finance SFDR).

There are other platforms that gives users access to real-time geo Copernicus satellite data and enables them to combine it with their own data and tools, to build new innovative products and services that integrate accurate and real-time information from satellites[25].

Monitoring risks in real-time is a classical task for NLP to support human analysts with. It is important to constantly monitor news and trends that could impact different markets around the globe. A data gathering and processing solution is needed that reduces the workload while keeping the quality of the information a risk analyst has to monitor on a daily basis. News and articles are automatically summarized and analyzed to be able to monitor hundreds of entities and real-time[26].

There are also real-Time ESG insights from vast amounts of structured and unstructured data with a no-code data science platform[27], using techniques like adaptive Natural Language Processing (NLP)[28]. In this way, native language extraction in dozens of languages can real-time coverage of public and private companies can be achieved. NLP can be used to tracking of climate impact and controversies, to anticpate risk and generate ESG scores, to detect ethical and social impacts, and to increase ESG transparency.

NLP platforms convert unstructured information into structured data and NLG (Natural Language Generation) turns this data into stories[29] and narratives. The following NLP-related technqiues can be used for ESG-related analyses: knowlegde graph, market sentiment, named entity recognition, summarization, question answering, topic modelling/clustering, patent and enterprise search, knowledge graphs, synonym search, etc.

Other platforms are dedicated to build Climate Intelligence[30]. Climate Intelligence is mass intelligence for managing climate decisions on valuable assets. providing insight into how risks such as flooding, droughts, and extreme temperatures will impact these assets.

---

[22] https://futureofsustainabledata.com/

[23] https://www.geoworks.sg/

[24] https://atos.net/en/2021/press-release_2021_04_21/atos-and-dreamquark-advance-responsible-finance-with-transparent-artificial-intelligence

[25] https://atos.net/en/2018/press-release_2018_06_21/atos-announces-satellite-data-platform-mundi-now-live

[26] https://tryolabs.com/work/allianz-GI/

[27] https://accern.com/esg-investing

[28] https://accern.com/adaptive-nlp

[29] https://www2.deloitte.com/us/en/insights/industry/financial-services/natural-language-processing-investment-management.html

30 https://www.linkedin.com/company/cervest/

Other platforms are capable of linking the sustainability factors to business performance. They combine AI and ESG big data to assess the performance and sustainability of corporations and to evaluate the alpha impact of sustainability, using self-learning quantitative models[31].
ESG data science leverages the combination of AI and machine learning with human intelligence to systematically analyze public information and identify material ESG risks[32].

[31] https://www.arabesque.com/ai/

[32] https://www.reprisk.com/

Computing Platforms for SupTech, AI Explainability and Compliance

The mandate of central banks and financial supervisors has become challenging. Pandemics, economic instability, cyber risks, sanctions, climate change, and sustainable investing are just a few examples of complex, emerging risks. Also, the digital, technological disruption of the financial ecosystem, as well as the speed of change and the growing complexity of the industry and its unbundling, increase the likelihood of supervisory blind spots.

For this reason, regulators will need to significantly adapt their operating model over the next decade to include technology and collaboration with (fin)tech firms. Many regulators are beginning to move in this direction with a variety of initiatives. Big Data Analytics, ML and AI will play a critical role, for example, in improving the quality and timeliness of risk identification and monitoring. Central banks already have access to vast amounts of valuable data, drawn from traditional, structured, and unstructured sources; streaming, complex, multi-layered, and alternative data to provide a holistic picture. Using advanced data collection and analytics techniques, certain areas of supervision will be able to use real-time monitoring of emerging risks and generate much earlier warning signals. Regulators will be able to nearcast the developments and provide a holistic picture of issues but also will be able to drill down into more granular micro developments and trends.[33]

While the authors (nor the firms with which they are affiliated) do not endorse all these points, they do offer valuable insights into areas of supervisory concern and identify valuable topics for discussion. According to a speech (https://www.bis.org/review/r180307d.htm) on AI and Banking by Prof. Joachim Wuermeling, Member of the Executive Board Deutsche Bundesbank, held at the 2nd Annual FinTech Conference in Brussels, there are six 'warnings' for supervisors sone of which we would like to highlight in this context:

- Do not miss the opportunities of Artificial Intelligence in finance ... – Human shortcomings in dealing with finance can be mitigated. As behavioral finance has taught us, biases, inertia, and ignorance lead to the malfunctioning of markets. AI allows humans to reach out beyond their intellectual limits or simply avoid mistakes.
- Central banks should embrace Artificial Intelligence – Central banks have access to vast amounts of valuable data stemming from market operations, supervision, payments, and statistics. They are well-positioned to tap the benefits of AI so they can enhance their ability to fulfil their mandate for price stability and the stability of the financial system.

The Bank of International Settlements (BIS) has published a survey report on the application of big data by central banks (see Serena et al. (2021)) with the following main conclusions:

Central banks have a comprehensive view of big data, they are increasingly using it, the range of big data sources exploited is diverse, and is effectively used to support central bank policies. Most central banks also report using big data for micro-level supervision

[33] For further reading and on the future of central banking see the following resources:

- https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2019/feb/Oliver-Wyman_Supervising-Tomorrow.pdf
- https://www.oliverwyman.com/our-expertise/insights/2020/apr/risk-journal-vol-9/redefining-business-models/future-proofing-financial-supervision.html
- https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2020/apr/future-proofing-financial-supervision.pdf
- https://www.forbesmiddleeast.com/leadership/opinion/how-big-data-and-ai-are-hailing-a-new-future-for-finance

and regulation. The survey also underscored the need for adequate IT infrastructure & human capital.

Supervisory Technology and Sustainable Finance Technology will be key topics in the digitalization of financial supervision (see Zeranski and Sancak (2020)).

A key issue in financial big data will be interpretable modelling. This is because to make evidence-based policy decisions central banks need to identify specific explanatory causes or factors which they can take action to influence. Furthermore, transparency regarding the information produced by big data analysis is essential to ensuring that its quality can be checked and that public decisions can be made on a sound, clearly communicated basis. Lastly, there are important legal constraints that reduce central banks' leeway when using private and confidential data; interpretable modelling helps address all these issues.

The focus of ML/AI models in central banking and supervision cannot be just predictive accuracy. Models must be trustworthy, interpretable, explainable, interactive, fair, robust, accountable, and secure. Proper risk management, data/AI governance, and compliance must be in place. For an in-depth description see Deutsche Bundesbank (2020).

A bank, commercial bank or supervisor using AI in production needs to overcome the explainability gap to produce transparent, appropriate governance, risk management, and controls over AI. The publication "Financial Risk Management and Explainable, Trustworthy, Responsible AI" (Fritz-Morgenthal et al. 2021) discusses further details.

More specific use cases can be found, moving the discussion from the realm of the general to the specific. One related use case can be found in Bussmann et al. (2020). The use case had been selected as the best AI case in the EU Horizon2020 project FIN-TECH (www.fintech-ho2020.eu) by the European financial services community including the European supervisors. Other related Explainable AI (XAI) use cases can be found in Jaeger et al. (2021) and Papenbrock et al. (2021). The developed approaches can help to implement Explainable AI using techniques like SHAPley values (a local, and global variable importance method with mathematical footings in co-operative game theory) even for large and complex models. For classical datasets, these methods can already substantially improve the transparency of portfolio allocation processes. They also enable the visualization of the variables and their influences of the entire data set in a single analysis. Clustering and network analysis of the variables and their influences are often used to find overall model structure and connections. Real-time monitoring of model drift in continuous learning machines is applied. Simulations and perturbations to test the robustness of the model can be run at large scale. Iterative and evolutionary approaches are now able to create and evaluate millions of models, allowing supervisors to select those that best balance prudential goals. It will also be necessary to meet the upcoming requirements from the European AI Act, especially the technical and auditing requirements for High-Risk AI[34]:

- Creating and maintaining a risk management system for the entire lifecycle of the system.
- Testing the system to identify risks and determine appropriate mitigation measures, and to validate that the system runs consistently for the intended purpose, with tests made against prior metrics and validated against probabilistic thresholds.
- Establishing appropriate data governance controls, including the requirement that all training, validation, and testing datasets be as complete, error-free, and representative as possible.
- Detailed technical documentation, including around system architecture, algorithmic design, and model specifications.
- Automatic logging of events while the system is running, with the recording conforming to recognized standards.
- Designed with sufficient transparency to allow users to interpret the system's output.

[34] See https://datainnovation.org/2021/05/the-artificial-intelligence-act-a-quick-explainer/

- Designed to maintain human oversight at all times and prevent or minimize risks to health and safety or fundamental rights, including an override or off-switch capability.

In summary, meeting the overall combination of the supervisory, legal, diverse stakeholder, and technical requirements will drive model development, deployment, monitoring and retirement process that features enhanced auditability, transparency, and explainability. The ever-increasing data volume, velocity, and variety – across structured and unstructured sources – combined with the rapid pace of AI development will drive an overall system architecture that is scalable, flexible, and secure. To be efficient with both people's time and energy, the system must strongly adopt lessons learned in the leading HPC and AI supercomputers of today, leveraging GPU accelerated compute and networking that can accelerate workloads across the diverse, end-to-end data science use cases of today and tomorrow.[35]

[35] For additional information, see “Computing Platforms for Big Data Analytics and Artificial Intelligence” (Bruno et al. (2020)), which highlights the experiences of central banks with respect to HPC platforms.

Use Case: Empowering Regions regarding their Environmental Exposures

We suggest an open-source-based accelerated platform to empower regions to join forces to tackle similar environmental challenges, to quantify and track changes in their environmental exposures in time, and to identify groups of regions with similar ecological exposures for policy and risk assessment purposes. Any common features and insights will also support more consistent and reliable supervisory analysis and reporting against sustainability goals.

Introduction

Legislation is provided on a national or supranational level, but the actual impact of environmental exposures happens on a regional level, as does the operational implementation of environmental protection. Regions with similar environmental challenges are interested in collaborating to exchange experiences, improve their negotiation leverage, and support each other in joint projects or ad-hoc emergencies like floods or wildfires. Improved transparency of regions' environmental exposures can also help decrease insurance and funding cost for public or private purposes.

Our suggested platform allows stakeholders to quickly build and deploy accelerated applications for large data sets of environmental exposures such as climate, deforestation, land use and population. It supports financial institutions, corporates, regulators, and supervisors (e.g., central banks) to visualize, assess and monitor the current and evolving physical risks at the level of regions and groups of regions with similar environmental exposures. This will allow regions, firms, and regulators to evaluate environmental risks more quickly, accurately, and more easily connect these risks to other macro and financial risks.

The tool's first objective is to quantify environmental exposures of regions using multiple layers of raster data from remote sensing with varying spatial resolutions as inputs. We apply unsupervised and semi-supervised learning to cluster regions based on the quantified environmental features within their borders. The identified region clusters can be used for coordinated policy development and implementation as well as for relative risk assessments even though the regions within the clusters may be geographically separated. Our proposed process allows connecting global raster data from diverse sources with macroeconomic and financial data to include environmental considerations into country risk analysis and policy considerations.

The tool's second objective is visualization. Large data from multiple sources can be interactively analysed on an ad-hoc basis without the need to pre-aggregate the population or use data subsets. The cluster analysis results from the machine learning steps are also displayed in the interactive dashboards and the macro and financial data per cluster.

The required High-Performance Computing (HPC) is done with GPUs: Graphics Processing Units. As we have mentioned previously, this allows us to take advantage of a variety of state-of-the-art tools, deliver faster time to insight, manage large and diverse datasets – all in an energy efficient fashion.

Data

Geospatial environmental data layers derived from satellite measurements have a raster format and need to be connected and enriched with relevant country and sector data to allow decision-makers from the public and private sectors analysis. The raster data is collected and displayed at open access sites like https://globalforestwatch.org or https://earthdata.nasa.gov/ but not yet further connected or processed. The open-access data contains a variety of data: land cover (tree cover, primary forests, tree cover height, forest landscape integrity indices, tree plantations and mangrove forests) land use (concessions for logging, mining, oil palm plantations and mills, oil, gas, and wood fiber; major dams; indigenous and community lands, population density and resource rights) climate change measurements (carbon flux, carbon density, potential carbon gains). Restor[36] is a data platform specifically for ecosystem restoration projects.

The data are of varying global coverage, spatial resolution, spatial granularity, and time resolution across datasets and sometimes even within one dataset. Therefore, evaluation across data layers and across time is not a trivial task.

Methodology

To connect the environmental raster data with region-specific data, we apply an unsupervised learning approach:

1. For each environmental data layer, and for each snapshot in time: Project raster image data to a political map.
2. Compute the mean raster values within the boundary of each region. This leads to regional features as their mean raster value, i.e., to a table (rows: regions, columns: mean raster values of each feature layer).
3. Cluster regions by the similarity of their feature vectors.
4. Characterize the regions by their mean feature vectors across all regions of this cluster.
5. Enrich cluster features with regional macro data and financial market data within the clusters.
6. Determine changes in cluster composition and cluster characteristics in time to allow for a cross-sectional comparison of regions and groups of regions.

Implementation and Technology Platform

Licensed under Apache 2.0, the RAPIDS suite of open-source software libraries and APIs gives the ability to execute end-to-end data science and analytics pipelines entirely on GPUs, including data engineering and preparation, ML/AI, graph analytics, cross-filtering, and visualization. RAPIDS exposes GPU parallelism and high-bandwidth memory speed through user-friendly Python interfaces. See the appendix for further information.

Cuxfilter is a real-time cross-filtering library on GPUs. This empowers data scientists to extract valuable insights by working independently across the whole stack — from raw data to user interface — and quickly deliver interactive dashboards without the need for large developer and IT teams. Computed live using GPUs, there is no need to pre-aggregate the population or use data subsets.[37]

---

[36] https://restor.eco

[37] Related visualization examples from developer blogs:

- https://developer.nvidia.com/blog/interactively-visualizing-a-drivetime-radius-from-any-point-in-the-us/
- https://developer.nvidia.com/blog/an-interactive-2010-census-plotly-dash-visualization-accelerated-by-rapids/
- https://medium.com/rapids-ai/plotly-census-viz-dashboard-powered-by-rapids-1503b3506652
- Related videos on the visualization technologies proposed: https://www.youtube.com/watch?v=3tSLwK8p210 (covid related dashboard w/ plotly

Appendix: GPU-Accelerated Data Science and AI Computing Platforms at Central Banks

Data science workflows have traditionally been slow and cumbersome, relying on CPUs to load, filter, and manipulate data and train and deploy models. GPUs reduce infrastructure costs and provide superior performance for end-to-end data science workflows using RAPIDS open-source software libraries. GPU-accelerated data science and AI workloads is available regardless of the location where GPUs are deployed, whether in the laptop, the workstation, in the data center, at the edge or in the public cloud.

In the following we will briefly discuss some use cases based on large financial data that benefit from several aspects of a GPU accelerated computing platform. The reference architecture is the following: enterprise-class GPUs using the Ampere Architecture (A100), featuring CUDA and Tensor cores, ultra-high memory bandwidth, NVLink for inter-GPU communications and the resource sharing and security isolation feature called MIG. Details are provided in Choquette et al. (2021). It is impossible to analyse the given amount of data in reasonable timeframes with traditional non-accelerated systems based on conventional CPUs.

- Using large-scale accounting data for financial statement audits: monitoring trustworthiness of financial statements and detecting potential misstatements by applying neural networks to learn representations of accounting data that constitute a representative audit sample and using these systems to detect potential anomalies.

- Deep learning (DL) for anomaly detection methods is used to identify opportunities and risks across many industries. Accurate methods are required to produce actionable predictions that don't simply add to an already noisy data environment. These techniques are especially applicable in fighting financial crimes like fraud and AML.

- Accelerating deep learning product recommendations with the transformer-based model BERT: a deep learning-based product recommendation system that provides a personalized ranked list of products to their sellers that help them to manage, run, and grow their business in the most efficient way possible.

- Neural networks for exotic options and risk: complex derivative modelling can be significantly accelerated using GPUs and neural nets. The unavailability of analytical solutions, the higher dimensionality for complex interest-rate and foreign-exchange products (volatility surfaces, multiple curves), and the requirements to compute sensitivities for hedging pose unique challenges. The oracles of traditional modelling can do complex computations, but they are expensive and slow, and are traditionally done on computer grids overnight. A modern GPU accelerated computing platform can obtain accurate valuations in well under a second. Buehler et al. (2019) developed deep hedging approaches to replace sensitivities-based hedging with reinforcement learning.

- The application of generative adversarial networks (GANs) for simulating market data for direct predictive analytics or for training other machine learning models. Generative adversarial networks are one tool for developing synthetic financial datasets that can be used as training data across many classes of machine learning models. GANs can complement and augment the value of Monte Carlo simulations and replicate regime-specific conditions to better prepare models for more robust predictive analytics. Using a generator and discriminator, with care, the model will transition the simulated data so that it converges to an empirical distribution in a Nash or Quasi-Nash process, preserving more of the real-world temporal characteristics consistent with the targeted market regime.

For DL workloads, GPU-based computing platforms have set records for the MLPerf[38] benchmark (an industry-standard set of benchmarks across a variety of AI modelling tasks), handily surpassing all other commercially available systems (see Mattson et al. (2020)). Comparable results are documented with respect to the STAC-A2™ Benchmark suite[39] which is the industry standard for testing technology stacks used for compute-intensive analytic workloads involved in pricing and risk management.

According to Li et al. (2020), gaining performance from multi-GPU scaling is not trivial, mainly because traditionally, inter-GPU communication shares the same bus interconnect as CPU-GPU communication, such as PCIe. This situation changed with the introduction of NVIDIA's DGX server line due to the introduction of dedicated GPU-oriented interconnects such as NVLink, and NVSwitch. The NCCL library allows multiple applications to efficiently optimize their communications within and across systems at the software layer. NCCL automatically understands the underlying hardware network topology to enable orchestrated mapping of functions and updates; it handles the issue of synchronization, overlapping and deadlock; and allows applications to prefer different performance metrics (e.g., latency-oriented for small transfers but bandwidth-oriented for large transfers).

RAPIDS

The RAPIDS software suite of open source[40] accelerated Python libraries give the ability to execute end-to-end data science and analytics pipelines entirely on GPUs. Integration into existing workflows normally requires only a few lines of code, because its API was deliberately designed to be consistent with existing data science utilities (e.g., Pandas DataFrame, SciKit Learn). RAPIDS is incubated by NVIDIA based on extensive hardware and data science experience both internal to NVIDIA and from customers and other contributors. RAPIDS wraps NVIDIA CUDA primitives for low-level compute optimization and exposes GPU parallelism and high-bandwidth memory speed through user-friendly Python interfaces. RAPIDS focuses on common data preparation tasks for analytics and data science. This includes a familiar data framedataframeintegrates with a variety of machine learning algorithms for end-to-end pipeline accelerations without paying typical serialization costs. RAPIDS also includes support for multi-node, multi-GPU deployments, enabling vastly accelerated processing and training on much larger dataset sizes. RAPIDS projects include cuDF, a pandas-like dataframe manipulation library; cuML, a collection of machine learning libraries that will provide GPU versions of algorithms available in scikit-learn; and cuGraph, a NetworkX-like accelerated graph analytics library. RAPIDS also provides tight integration with the key deep learning frameworks. This means data processed by RAPIDS can be seamlessly pushed to deep learning frameworks that accept an array interface or work with DLPack, such as Chainer, MXNet, and PyTorch.

RAPIDS democratizes the power of GPU accelerated data science for everyone with observed accelerations from CPU to GPU that can range from a factor of 10x to 1000x in some cases.

According to Entschev, Kirkham, and Ronaghi (2020), the Oak Ridge Leadership Computing Facility (OLCF) provides a variety of tools to perform data wrangling and data analysis tasks, CPU-only Python based tools often lack scalability, or the ability to fully exploit the computational capability of OLCF's Summit supercomputer. OLCF started a detailed performance evaluation on NVIDIA RAPIDS and Dask and to exploit how it helps to distribute data analytics workloads from personal computers to heterogeneous supercomputing systems. The performance evaluation of NVIDIA RAPIDS also included a subset of RAPIDS libraries, i.e., cuDF, cuML, and cuGraph, and Chainer's CuPy, and their multi-GPU variants when available.

[38] https://mlcommons.org/en/

[39] https://stacresearch.com/

[40] Apache 2.0 License

According to Liu, Tunguz, and Titericz (2020), RAPIDS has been proven to be a very efficient for black-box optimization algorithms. According to the paper the authors implemented a fast multi-GPU accelerated exhaustive search on the DGX-1 to find the best ensemble of optimization algorithms. The ensemble algorithm has been generalized to multiple optimizers and the proposed framework scales with multiple GPUs. They evaluated 15 optimizers by training 2.7 million models and running 541,440 optimizations. On a DGX-1, the search time is reduced from more than 10 days on two 20-core CPUs to less than 24 hours on 8-GPUs. Multi-GPU-optimized framework to accelerate a brute force search for the optimal ensemble of black-box optimization algorithms by running many experiments in parallel.

According to Yang, Buluc, and Owens (2020), there is a mismatch between the high-level languages that users and graph algorithm designers would prefer to program in (e.g., Python) and programming languages for parallel hardware (e.g., C++, CUDA, OpenMP, or MPI). To address this mismatch, many initiatives, including NVIDIA's RAPIDS effort, have been launched in order to provide an open-source Python-based ecosystem for data science and graphs on GPUs.

Ocsa (2019) present BlazingSQL, a SQL engine build on RAPIDS open-source software, which allows accelerated query of enterprise data lakes with full interoperability with the RAPIDs stack.

According to Richardson et al. (2020), the lines between data science (DS), machine learning (ML), deep learning (DL), and data mining continue to be blurred and removed. This ushers in vast amounts of new capabilities, but it also brings increased complexity and a vast increase in the number of tools/techniques that must be learned. It is common for DL engineers to use one set of tools for data extraction/cleaning and then pivot to another library for training their models. After training and inference, it is common to then move data yet again to another set of tools for post-processing. The RAPIDS suite of open-source libraries not only provide a method to execute and accelerate these tasks using GPUs with familiar APIs, but it also provides interoperability with the broader open-source community and DL tools while removing unnecessary serializations that slow down workflows. GPUs provide massive parallelization that DL has leveraged for some time, and RAPIDS provides the missing pieces that extend this computing power to more traditional yet important DS and ML tasks (e.g., ETL, modelling). Complete pipelines can be built that encompass everything, including ETL, feature engineering, ML/DL modelling, inference, and visualization, all while removing typical serialization costs and affording seamless interoperability between libraries. All experiments using RAPIDS can effortlessly be scheduled, logged and reviewed using existing public cloud options.

RAPIDS SPARK 3.0

Given the parallel nature of many data processing tasks, it's only natural that the massively parallel architecture of a GPU should be able to parallelize and accelerate Apache Spark data processing queries, in the same way that a GPU accelerates deep learning (DL) in artificial intelligence (AI). NVIDIA has worked with the Apache Spark community to implement GPU acceleration through the release of Spark 3.0 and the open-source RAPIDS Accelerator for Spark[41]. The RAPIDS Accelerator for Apache Spark uses GPUs to:

- Accelerate end-to-end data preparation and model training on the same Spark cluster.
- Accelerate Spark SQL and DataFrame operations without requiring any code changes.
- Accelerate data transfer performance across nodes (Spark shuffles).

As ML and DL are increasingly applied to larger datasets, Spark has become a commonly used vehicle for the data pre-processing and feature engineering needed to prepare raw input data for the learning phase. Because Spark 2.x has no knowledge about GPUs, data scientists and engineers perform the ETL on CPUs, then send the data over to GPUs for model training.

The Apache Spark community has been focused on bringing both phases of this end-to-end pipeline together, so that data scientists can work with a single Spark cluster and avoid the performance penalty of moving data between Spark based systems for data preparation and PyTorch or TensorFlow based systems for Deep Learning. Apache Spark 3.0 represents a key milestone, as Spark can now schedule GPU-accelerated ML and DL applications on Spark clusters with GPUs, removing bottlenecks, increasing performance, and simplifying clusters. In Apache Spark 3.0, you can now have a single pipeline, from data ingest to data preparation to model training on a GPU powered cluster.

NGC

The NVIDIA GPU Cloud (NGC)[42] is a container repository with GPU optimized container images for popular Deep Learning and machine learning frameworks and comes bundled with the CUDA and DL libraries. With the introduction with NVIDIA NGC, deployment and development was dramatically simplified as described in Radhakrishnan, Varma, and Kurkure (2019). The traditional non-virtualized approach typically is to access GPU in bare-metal or native environments. This entails installing the DL software stack (DL framework, libraries, CUDA libraries) along with the GPU drivers on the bare metal system. A containerized DL solution bundles the DL framework, referenced run-time libraries and CUDA libraries in a tested, supportable container for ease of deployment and portability.

---

[41] https://nvidia.github.io/spark-rapids/

[42] https://www.nvidia.com/en-us/gpu-cloud/

Acknowledgements

This work partially relies on the support from the European Union's Horizon 2020 research and innovation program "FIN-TECH: A Financial supervision and Technology compliance training programme" under the grant agreement No 825215 (Topic: ICT-35-2018, Type of action: CSA).

Thanks to Thomas Schoenemeyer, Wolfram Peters, and Allen Enemark for the discussions.

References

Aguerzame, Abdallah, Benoit Pelletier, and Francois Waeselynck. 2019. “GPU Acceleration of Pyspark Using Rapids Ai.” INSTICC; SciTePress. https://doi.org/10.5220/0008191404370442.

BaFin (2020) Guidance Notice on Dealing with Sustainability Risks. https://www.bafin.de/SharedDocs/Downloads/EN/Merkblatt/dl_mb_Nachhaltigkeits risiken_en.html

Berg, Florian and Kölbel, Julian and Rigobon, Roberto, Aggregate Confusion: The Divergence of ESG Ratings (May 17, 2020). Available at SSRN: https://ssrn.com/abstract=3438533 or http://dx.doi.org/10.2139/ssrn.3438533

Bok, Brandyn, Daniele Caratelli, Domenico Giannone, Argia M Sbordone, and Andrea Tambalotti. 2018. “Macroeconomic Nowcasting and Forecasting with Big Data.” Annual Review of Economics 10: 615–43.

Bruno, Giovanni, Mikheil Esakia, and Felix Goltz. 2021. "Honey, I Shrunk the ESG Alpha": Risk-Adjusting ESG Portfolio Returns — April 2021. Scientific Beta.

Bruno, Giuseppe, Hiren Jani, Rafael Schmidt, and Bruno Tissot. 2020. "Computing platforms for big data analytics and artificial intelligence." IFC Reports 11. Bank for International Settlements. https://ideas.repec.org/p/bis/bisifr/11.html.

Buehler, Hans, Lukas Gonon, Josef Teichmann & Ben Wood (2019) Deep hedging, Quantitative Finance, 19:8, 1271-1291, DOI: 10.1080/14697688.2019.1571683

Bussmann, Niklas, Paolo Giudici, Dimitri Marinelli, and Jochen Papenbrock. 2020. "Explainable Machine Learning in Credit Risk Management." Computational Economics, September. https://doi.org/10.1007/s10614-020-10042-0.

Campiglio, Emanuele, Yannis Dafermos, Pierre Monnin, Josh Ryan-Collins, Guido Schotten, and Misa Tanaka. 2018. "Climate Change Challenges for Central Banks and Financial Regulators." Nature Climate Change 8 (6): 462–68.

Chakraborty, Chiranjit, and Andreas Joseph. 2017. "Machine learning at central banks." Bank of England working papers 674. Bank of England. https://ideas.repec.org/p/boe/boeewp/0674.html.

Chen, Mike, von Behren, Robert and George Mussali (2021). The Unreasonable Attractiveness of More ESG Data. The Journal of Portfolio Management Novel Risks 2021, jpm.2021.1.281 DOI: https://doi.org/10.3905/jpm.2021.1.281

Choquette, Jack, Wishwesh Gandhi, Olivier Giroux, Nick Stam, and Ronny Krashinsky. 2021. "NVIDIA A100 Tensor Core Gpu: Performance and Innovation." IEEE Micro.

Committee, Irving Fisher. 2017. Big Data. Bank for International Settlements. https://EconPapers.repec.org/RePEc:bis:bisifb:44.

———. 2019. The Use of Big Data Analytics and Artificial Intelligence in Central Banking. Bank for International Settlements. https://EconPapers.repec.org/RePEc:bis:bisifb:50.

Corey J. Nolet, Edward Raff, Victor Lafargue, and Joshua Patterson. 2020. "Bringing Umap Closer to the Speed of Light with Gpu Acceleration." http://arxiv.org/abs/2008.00325.

Covert, Ian, Scott Lundberg, and Su-In Lee. 2020. "Explaining by Removing: A Unified Framework for Model Explanation." http://arxiv.org/abs/2011.14878.

Dafermos, Yannis, Daniela Gabor, Maria Nikolaidi, Adam Pawloff and Frank van Lerven. 2021. Greening the Eurosystem collateral framework. New Economics Foundation.

Damm, A., Paul-Limoges, E., Kückenbrink, D., Bachofen, C., & Morsdorf, F. (2020). Remote sensing of forest gas exchange: Considerations derived from a tomographic perspective. Global Change Biology, 26(4), 2717–2727.

Deutsche Bundesbank (2020). Policy Discussion Paper The Use of Artificial Intelligence and Machine Learning in the Financial Sector. Retrieved from https://www.bundesbank.de/resource/blob/598256/d7d26167bceb18ee7c0c296902e42162/mL/2020-11-policy-dp-aiml-data.pdf

Doerr, Sebastian, Leonardo Gambacorta, and Jose Maria Serena Garralda. 2021. "Big data and machine learning in central banking." BIS Working Papers 930. Bank for International Settlements. https://ideas.repec.org/p/bis/biswps/930.html.

Edmans, Alex. "Letter: ESG studies must show scientific rigour to avoid feeding biases". Financial Times, 11.5.2021.

Entschev, Peter, John Kirkham, and Zahra Ronaghi. 2020. "Performance Evaluation of Python Based Data Analytics Frameworks in Summit: Early Experiences." In Driving Scientific and Engineering Discoveries Through the Convergence of Hpc, Big Data and Ai: 17th Smoky Mountains Computational Sciences and Engineering Conference, Smc 2020, Oak Ridge, Tn, Usa, August 26-28, 2020, Revised Selected Papers, 1315:366. Springer Nature.

FCA (2021). Guiding principles on design, delivery and disclosure of ESG and sustainable investment funds. https://www.fca.org.uk/news/news-stories/guiding-principles-on-design-delivery-disclosure-esg-sustainable-investment-funds

Flood, Mark, H. V. Jagadish, and L. Raschid. 2016. "Big Data Challenges and Opportunities in Financial Stability Monitoring." Financial Stability Review, no. 20: 129–42. https://EconPapers.repec.org/RePEc:bfr:fisrev:2016:20:14.

Fritz-Morgenthal, Sebastian and Hein, Bernhard and Papenbrock, Jochen, Financial Risk Management and Explainable Trustworthy Responsible AI (June 25, 2021). Available at http://dx.doi.org/10.2139/ssrn.3873768

Giese, Guido, Zoltán Nagy and Linda-Eling Lee. 2021. Deconstructing ESG Ratings Performance: Risk and Return for E, S, and G by Time Horizon, Sector, and Weighting. The Journal of Portfolio Management February 2021, jpm.2020.1.198; DOI: https://doi.org/10.3905/jpm.2020.1.198

Guderian, Carsten .C., Posth, Jan-Alexander., Grob, Linus (2021): "Utilizing Patent Analytics for Investment Decisions in COVID-19 Technologies: A Multi-Case Study", Working Paper.

Hammer, Cornelia, Diane Kostroch, and Gabriel Quiros. 2017. "Big Data : Potential, Challenges and Statistical Implications." USA: INTERNATIONAL MONETARY FUND. https://www.elibrary.imf.org/view/IMF006/24476-9781484310908/24476-9781484310908/24476-9781484310908.xml.

Han, Chuan-Hsiang. 2020. "GPU Acceleration for Computational Finance." In Handbook of Financial Econometrics, Mathematics, Statistics, and Machine Learning, 1519–32. WORLD SCIENTIFIC. https://doi.org/10.1142/9789811202391_0039.

IOSCO (2021). Recommendations on Sustainability-Related Practices, Policies, Procedures and Disclosure in Asset Management. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD679.pdf

Jaeger, Markus, Stephan Krügel, Dimitri Marinelli, Jochen Papenbrock, and Peter Schwendner. 2021. "Interpretable Machine Learning for Diversified Portfolio Construction." The Journal of Financial Data Science 3(3). https://doi.org/10.3905/jfds.2021.1.066

Khraisha, Tamer. 2020. "A Holistic Approach to Financial Data Science: Data, Technology, and Analytics." The Journal of Financial Data Science 2 (2): 64–84. https://doi.org/10.3905/jfds.2020.1.031.

Krishna, Suresh, and Ravi Krishna. 2020. "Accelerating Recommender Systems via Hardware "Scale-in"." http://arxiv.org/abs/2009.05230.

Krukovets, Dmytro, and others. 2020. "Data Science Opportunities at Central Banks: Overview." Visnyk of the National Bank of Ukraine 249: 13–24.

Li, A., S. L. Song, J. Chen, J. Li, X. Liu, N. R. Tallent, and K. J. Barker. 2020. "Evaluating Modern Gpu Interconnect: PCIe, Nvlink, Nv-Sli, Nvswitch and Gpudirect." IEEE

Transactions on Parallel and Distributed Systems 31 (1): 94–110. https://doi.org/10.1109/TPDS.2019.2928289.

Liu, Jiwei, Bojan Tunguz, and Gilberto Titericz. 2020. "GPU Accelerated Exhaustive Search for Optimal Ensemble of Black-Box Optimization Algorithms." http://arxiv.org/abs/2012.04201.

Martellini, Lionel, Vallée, Lou-Salomé. 2021. "Measuring and Managing ESG Risks in Sovereign Bond Portfolios and Implications for Sovereign Debt Investing". The Journal of Portfolio Management Aug 2021, jpm.2021.1.290; DOI: 10.3905/jpm.2021.1.290

Mattson, Peter, Christine Cheng, Cody Coleman, Greg Diamos, Paulius Micikevicius, David Patterson, Hanlin Tang, et al. 2020. "MLPerf Training Benchmark." http://arxiv.org/abs/1910.01500.

Mitchell, Rory, Eibe Frank, and Geoffrey Holmes. 2020. "GPUTreeShap: Fast Parallel Tree Interpretability." http://arxiv.org/abs/2010.13972.

Ocsa, Alexander. 2019. "SQL for GPU Data Frames in RAPIDS Accelerating end-to-end data science workflows using GPUs." LatinX in AI Research at ICML 2019. https://hal.archives-ouvertes.fr/hal-02264776.

Papenbrock, Jochen, Peter Schwendner, Markus Jaeger, and Stephan Krügel. 2021. "Matrix Evolutions: Synthetic Correlations and Explainable Machine Learning for Constructing Robust Investment Portfolios." The Journal of Financial Data Science, March, jfds.2021.1.056. https://doi.org/10.3905/jfds.2021.1.056.

Pedersen, Lasse Heje, Shaun Fitzgibbons, Lukasz Pomorski, Responsible investing: The ESG-efficient frontier, Journal of Financial Economics, 2020, ISSN 0304-405X, https://doi.org/10.1016/j.jfineco.2020.11.001.

Radhakrishnan, Ramesh, Yogesh Varma, and Uday Kurkure. 2019. "Evaluating Gpu Performance for Deep Learning Workloads in Virtualized Environment." In 2019 International Conference on High Performance Computing & Simulation (Hpcs), 904–8. IEEE.

RAPIDS-Spark. 2021. "RAPIDS Accelerator for Apache Spark." GitHub Repository. https://github.com/NVIDIA/spark-rapids; GitHub.

Raschka, Sebastian, Joshua Patterson, and Corey Nolet. 2020. "Machine Learning in Python: Main Developments and Technology Trends in Data Science, Machine Learning, and Artificial Intelligence." Information 11 (4): 193. https://arxiv.org/abs/2002.04803.

Richardson, Bartley, Bradley Rees, Tom Drabas, Even Oldridge, David A Bader, and Rachel Allen. 2020. "Accelerating and Expanding End-to-End Data Science Workflows with Dl/Ml Interoperability Using Rapids." In Proceedings of the 26th Acm Sigkdd International Conference on Knowledge Discovery & Data Mining, 3503–4.

Santos, M. J., & Dekker, S. C. (2020). Locked-in and living delta pathways in the Anthropocene. Scientific Reports, 10(1), 19598. https://doi.org/10.1038/s41598-020-76304-x

Schneider, F. D., Morsdorf, F., Schmid, B., Petchey, O. L., Hueni, A., Schimel, D. S., & Schaepman, M. E. (2017). Mapping functional diversity from remotely sensed morphological and physiological forest traits. Nature Communications, 8(1). https://doi.org/10.1038/s41467-017-01530-3

Serena, Jose Maria, Bruno Tissot, Sebastian Doerr, and Leonardo Gambacorta. 2021. "Use of big data sources and applications at central banks." IFC Reports 13. Bank for International Settlements. https://ideas.repec.org/p/bis/bisifr/13.html.

Sokolov, Alik, Kyle Caverly, Jonathan Mostovoy, Talal Fahoum and Luis Seco (2021). Weak Supervision and Black-Litterman for Automated ESG Portfolio Construction. The Journal of Financial Data Science Summer 2021, 3 (3) 129-138; DOI: https://doi.org/10.3905/jfds.2021.1.070

Sorensen, Eric, Mike Chen and George Mussalli (2021).The Quantitative Approach for Sustainable Investing. The Journal of Portfolio Management August 2021, 47 (8) 38-49 DOI: https://doi.org/10.3905/jpm.2021.1.267

Spetko, Matej, Ondrej Vysocky, Branislav Jansik, and Lubomir Riha. 2021. "DGX-A100 Face to Face Dgx-2-Performance, Power and Thermal Behavior Evaluation." Energies 14 (2). https://doi.org/10.3390/en14020376.

Spilak, Bruno and Wolfgang Karl Härdle. 2020. Tail-risk protection: Machine Learning meets modern Econometrics. https://arxiv.org/abs/2010.03315

Team, RAPIDS Development. 2018. "RAPIDS: Collection of Libraries for End to End Gpu Data Science." https://rapids.ai.

Tissot, Bruno. 2018. "Big Data for Central Banks." In International Workshop on Big Data for Central Bank Policies–Bali, 23:25.

Tissot, Bruno, Timur Hulagu, Per Nymand-Andersen, and Laura Comino Suarez. 2015. "Central Banks' Use of and Interest in "Big Data"." IFC Reports. Bank for International Settlements. https://EconPapers.repec.org/RePEc:bis:bisifr:3.

Verheyen, Roda. 2021. Legal opinion: Legal options for implementing climate criteria in the monetary policy of the European Central Bank. https://www.greenpeace.org/eu-unit/issues/climate-energy/45725/legal-options-for-implementing-climate-criteria-in-the-monetary-policy-of-the-european-central-bank/

Wibisono, Okiriza, Hidayah Dhini Ari, Anggraini Widjanarti, Alvin Andhika Zulen, and B. Tissot. 2019. "The Use of Big Data Analytics and Artificial Intelligence in Central Banking - an Overview." IFC Bulletins Chapters 50.

Yang, Carl, Aydin Buluc, and John D. Owens. 2020. "GraphBLAST: A High-Performance Linear Algebra-Based Graph Framework on the Gpu." http://arxiv.org/abs/1908.01407.

Zammit-Lucia, Joe. "Letter: Investors should be wary of ESG virtue signalling". Financial Times, 6.5.2021.

Zeranski, Stefan, and Ibrahim Ethem Sancak. 2020. "Digitalisation of Financial Supervision with Supervisory Technology (Suptech)." J. Intl. Banking L. & Reg., no. 8. https://ssrn.com/abstract=3632053.

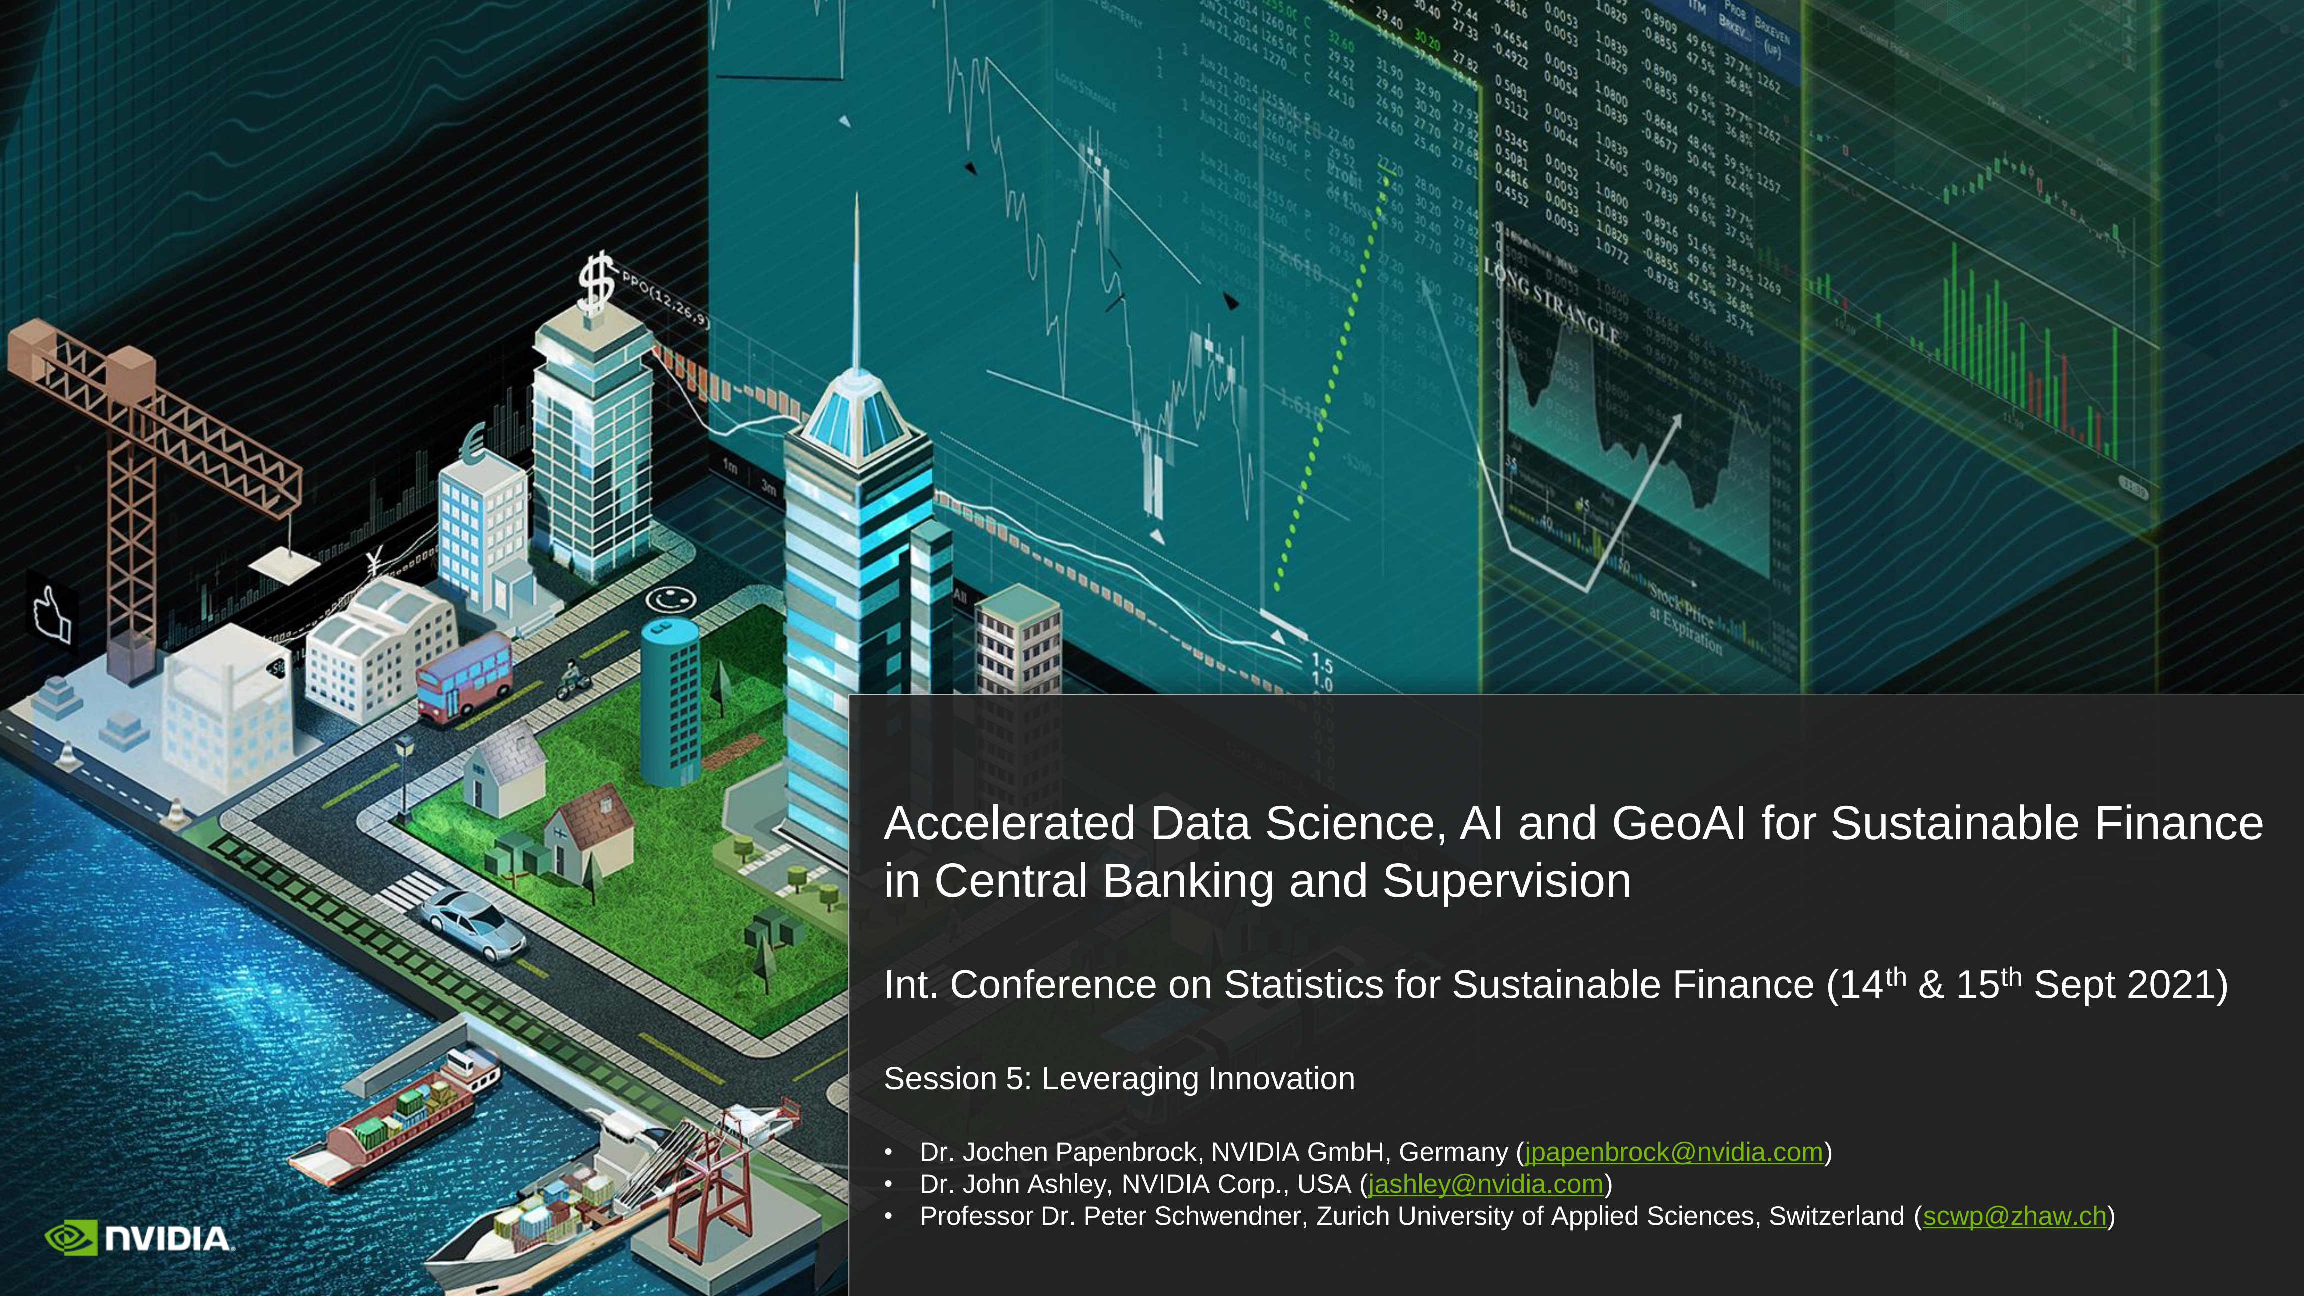

Accelerated Data Science, AI and GeoAI for Sustainable Finance in Central Banking and Supervision

Int. Conference on Statistics for Sustainable Finance (14th & 15th Sept 2021)

Session 5: Leveraging Innovation

- Dr. Jochen Papenbrock, NVIDIA GmbH, Germany (jpapenbrock@nvidia.com)
- Dr. John Ashley, NVIDIA Corp., USA (jashley@nvidia.com)
- Professor Dr. Peter Schwendner, Zurich University of Applied Sciences, Switzerland (scwp@zhaw.ch)

NVIDIA

LEVERAGING INNOVATION

Abstract

- Sustainability and Data Science / Artificial Intelligence (AI) are the two megatrends in FSI and supervision
- Use of massive amount of complex data
- Demand for computing platforms, HPC, and ‘AI’
- We see systems with data centric design, explainable/trustworthy AI models, GPU acceleration as key feature,
- Practical examples in ESG investing and data-driven policy design and investment/risk management

TRANSFORMED SUPERVISORY MODELS

Supervisory Technology (SupTech)

- synthetizes the vast quantities of structured and unstructured data
- improves actuality and timeliness (‘real-time’) of (emerging) risk identification, monitoring, early warning/intervention.
- supports both ‘full picture’ as near cast and granular, zooming level
- The same applies to Sustainable Finance Technology

NEED FOR COMPUTING PLATFORM

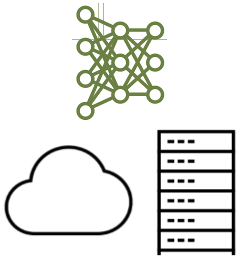

Establishment of computing platforms going hand in hand with the data collection and curation process. Several factors drive this need for a co-evolutionary approach:

1. Data, IT systems, operations and policies must be fully aligned and integrated: adaptive, auditable process to evolving data quantity and quality; validation, interpretability, narrative
2. filling data gaps and improve data quality (outliers)
3. HPC for simulating data and for complex optimization
4. large scale visualization, clustering and network analysis for data exploration
5. size of data and number of sources will increase, and so will the number, complexity, and frequency of updates of models

Compliance with the proposed AI Act for high-risk AI might also require a computing platform

RISE OF GPU COMPUTING

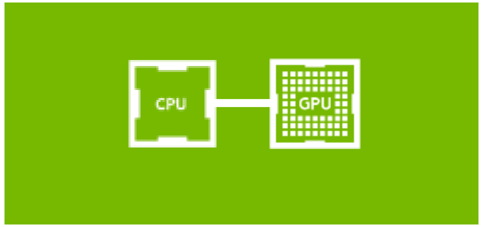

'Computing platforms for big data analytics and artificial intelligence.'
IFC Reports 11. Bank for International Settlements.
https://ideas.repec.org/p/bis/bisifr/11.html

"Central banks' experience shows that HPC platforms are primarily developed to ensure that computing resources are used in the most efficient way, so that analytical processes can be completed as rapidly as possible. […]

A processor core (or "core") is a single processing unit. Today's computers – or CPUs (central processing units) – have multiple processing units, with each of these cores able to focus on a different task. Depending on the analytical or statistical problem at hand, clusters of GPUs (graphics processing units, which have a highly parallel structure and were initially designed for efficient image processing) might also be embedded in computers, for instance, to support mass calculations"

HOW GPU ACCELERATION WORKS

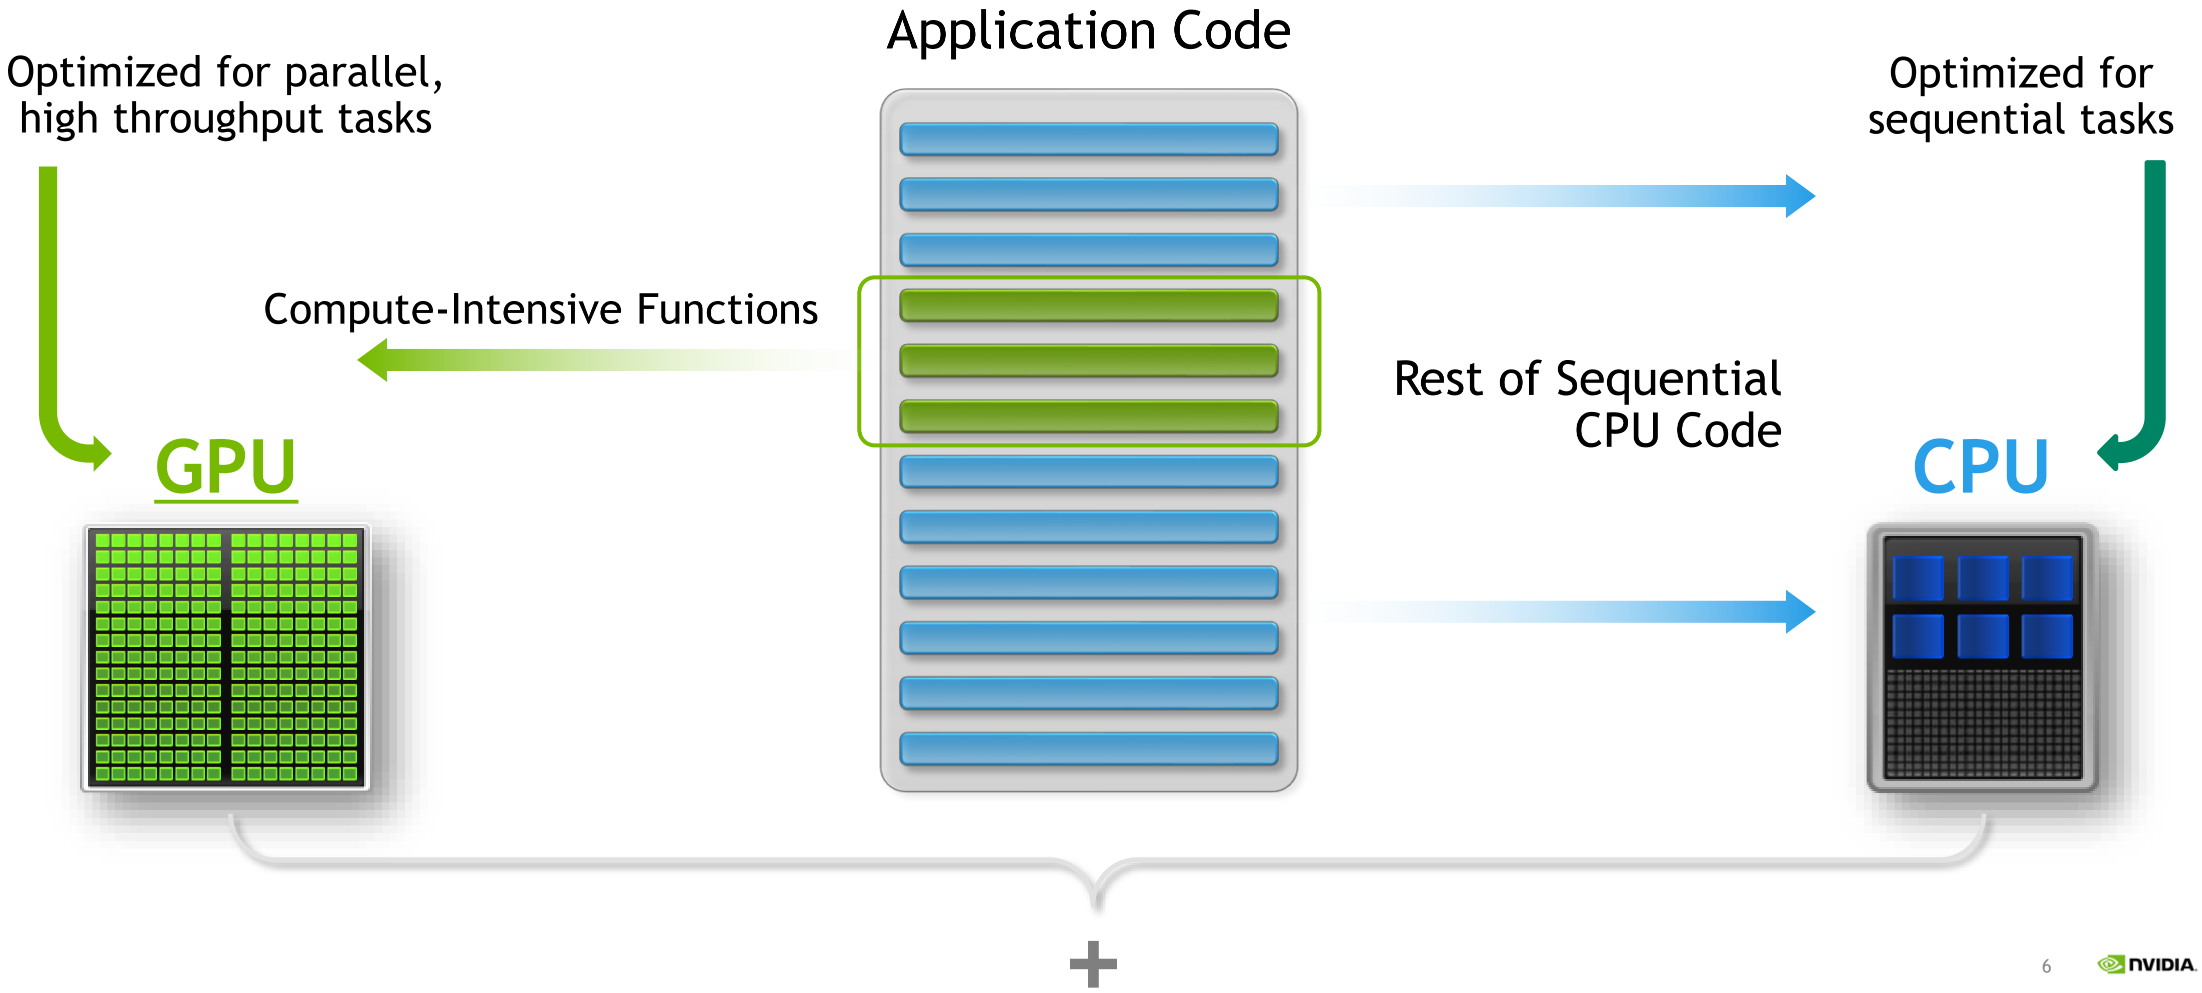

PYTHON TOOLS HAVE DEMOCRATIZED DATA SCIENCE

ACCESSIBLE, EASY TO USE TOOLS ABSTRACT COMPLEXITY

Python is the most-used language in Data Science today. Libraries like NumPy, Scikit-Learn, and Pandas have changed how we think about accessibility in Data Science and Machine Learning.

While great for experimentation, PyData tools lack the power necessary for enterprise-scale workloads. This leads to substantial refactoring to handle the size of modern problems, increasing cycle time, overhead, and time to insight.

These pain points are further compounded by computational bottlenecks of CPU-based processing.

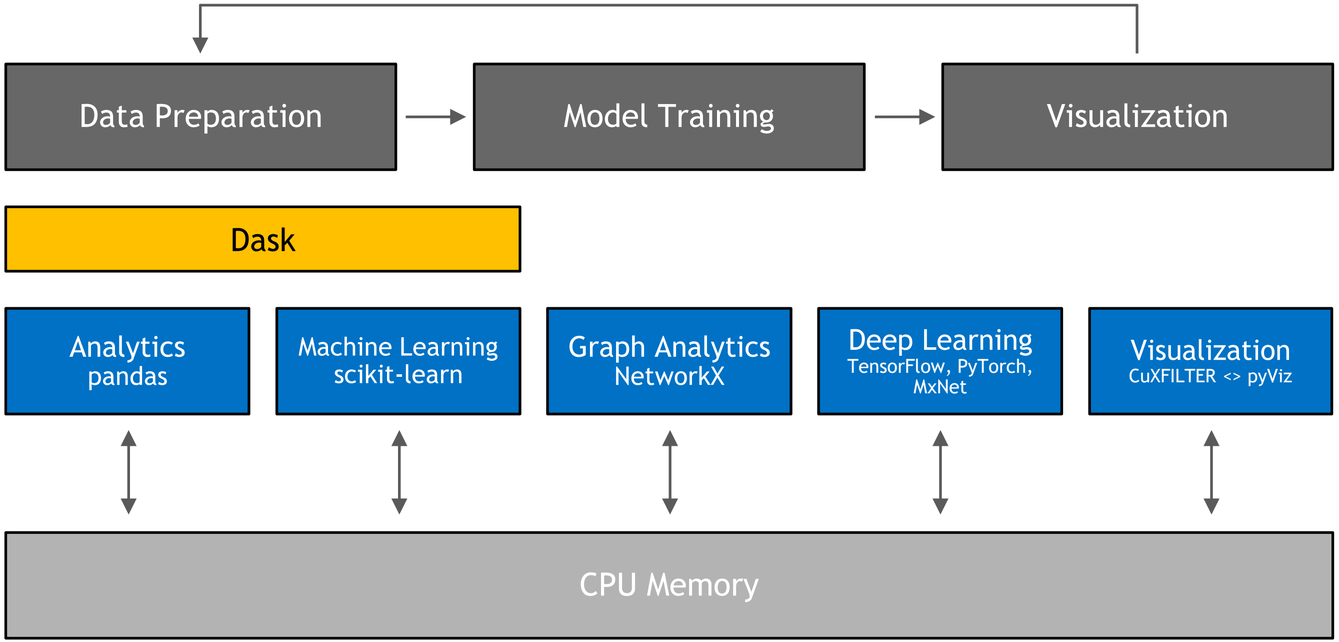

Code refactors and inter-team handoffs decrease data-driven ROI

RAPIDS ACCELERATES POPULAR DATA SCIENCE TOOLS

DELIVERING ENTERPRISE-GRADE DATA SCIENCE SOLUTIONS IN PURE PYTHON

The RAPIDS suite of open source software libraries gives you the freedom to execute end-to-end data science and analytics pipelines entirely on GPUs.

RAPIDS exposes GPU parallelism and high-bandwidth memory speed through user-friendly Python interfaces like PyData.

With Dask, RAPIDS can scale out to multi-node, multi-GPU cluster to power through big data processes.

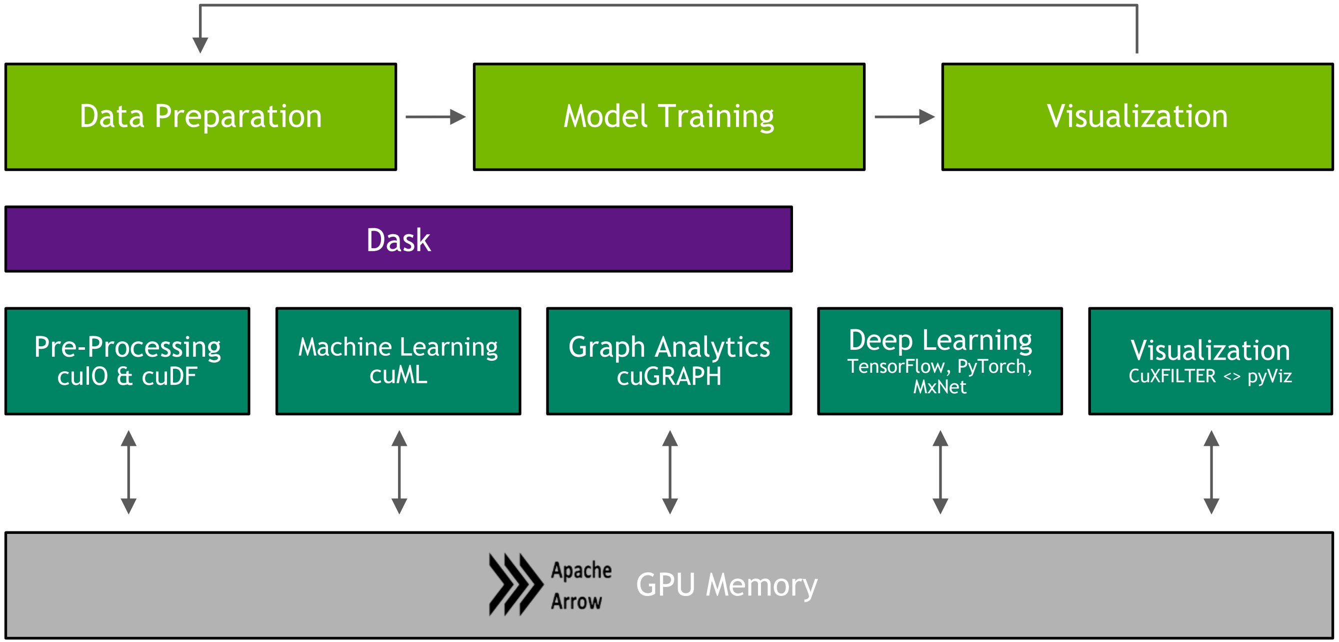

RAPIDS enables the PyData stack with the power of GPUs

XAI USE CASE ON ENTERPRISE PRODUCTION LEVEL

NVIDIA acceleration of XAI use case accelerated with RAPIDS

- Use case on credit portfolio risk with realistic data (11.2 million dataset from Fannie Mae)
- General speedups of up to 19x for SHAP values, and 340x for SHAP interaction values
- use case based on best AI use case in EU Horizon2020 project FIN-TECH *)

FINTECH RISK MANAGEMENT www.fintech-ho2020.eu

European Commission

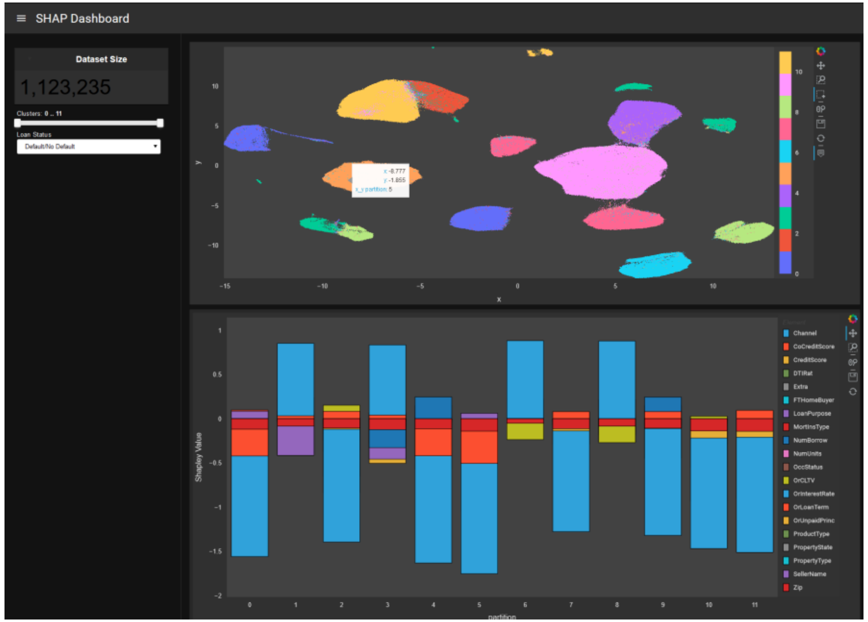

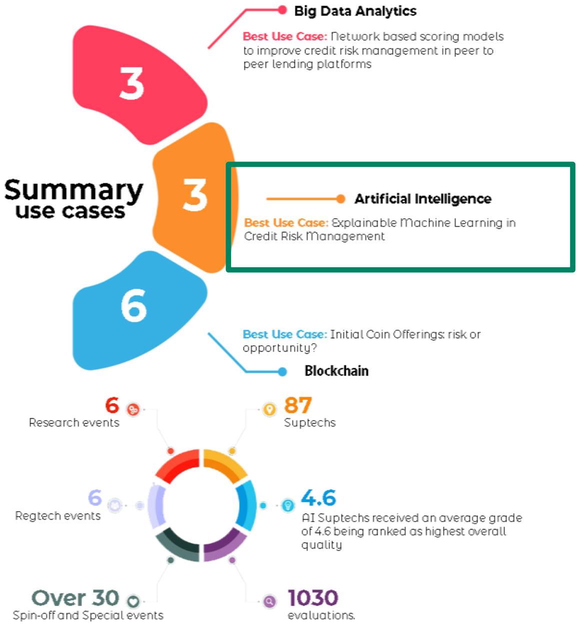

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 825215 (topic ICT-35-2018, Type of action: CSA)The content reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains.

NVIDIA

XAI FOR PORTFOLIO CONSTRUCTION

- Collaboration with Munich Re
- 3 Publications in top Journal of Financial Data Science
- Implementing a new workflow and combine XAI with synthetic market data generation to enhance the explainability
- Additional paper with cryptos is in the pipeline:
  “Can adaptive seriational risk parity tame crypto portfolios?”

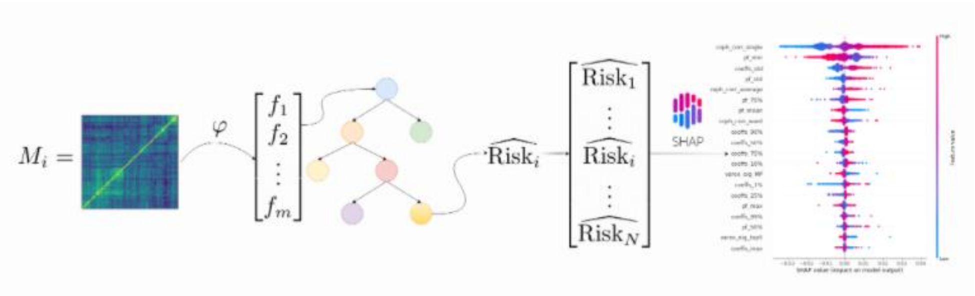

SUSTAINABLE FINANCE AND ESG INVESTING

Problem

- assessing the appropriate level of protection against sustainability risks is a challenge for central banks
- existing sustainability ratings from different providers diverge significantly for the same investable asset (Berg et al. 2019). This implies risk for greenwashing.
- also conflicting findings of ESG alpha (Bruno et al. 2021 vs Giese et al. 2021);

➡ calling for more scientific rigour in ESG evaluation studies (Edmans 2021)

Solution

- consider full set of data and information from multiple source
- process them with explainable AI/NLP models and aggregate them in a large-scale, transparent portfolio construction and optimization process
- This is where the two megatrends, Sustainable Finance and AI, that will shape the financial sector in the coming years overlap; computing platforms will play a crucial role

EMPOWERING REGIONS REGARDING THEIR ENVIRONMENTAL EXPOSURES (1/3)

Use Case

Purpose: empower regions to jointly tackle environmental challenges; identify regions with similar environmental exposures (climate, deforestation, land use, population, etc.)

- Why?
    - improve and benchmark ESG data
    - exchange experiences
    - improve their negotiation leverage
    - support each other in joint projects or ad-hoc emergencies like floods or wildfires
    - decrease insurance and funding cost for public or private purposes
    - monitor transition and physical risk

EMPOWERING REGIONS REGARDING THEIR ENVIRONMENTAL EXPOSURES (2/3)

Use Case

Data

- Geospatial environmental data on land cover, land use, climate change measurements, etc.
- data layers derived from satellite measurements have a raster format and need to be connected and enriched with relevant country and sector data (e.g. also with economic and financial data)
- potential sources: https://globalforestwatch.org, https://earthdata.nasa.gov/, https://www.restor.eco/, etc.
- the data are of varying global coverage, spatial resolution, spatial granularity and time resolution across datasets
- Non-trivial task

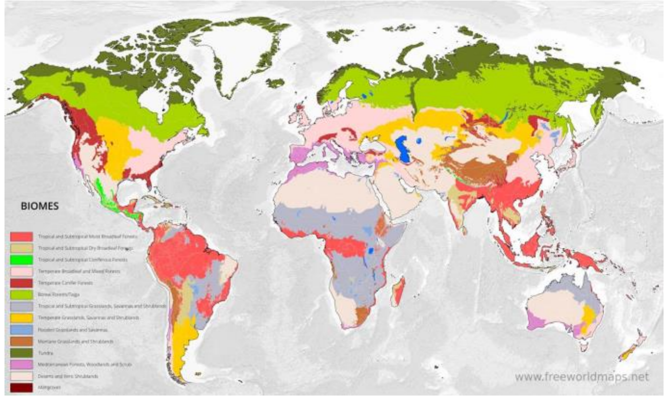

https://www.freeworldmaps.net/biomes/biomes-map.jpg

Method

- Similarity Clustering, Network Analysis, Visualization and real-time cross-filtering, typical RAPIDS use case

EMPOWERING REGIONS REGARDING THEIR ENVIRONMENTAL EXPOSURES (3/3)

Use Case

Potential implementation with **RAPIDS**

Data Preparation → Model Training → Visualization

Dask

Pre-Processing
cuIO & cuDF

Machine Learning
cuML

Graph Analytics
cuGRAPH

Deep Learning
TensorFlow, PyTorch, MxNet

Visualization
CuXFILTER <> pyViz

Apache Arrow GPU Memory

- cuxfilter
  github.com/rapidsai/cuxfilter
- Datashader
  datashader.org/
- Plotly Dash
  plotly.com/dash/
- Holoviews
  holoviews.org/
- Bokeh
  bokeh.org/

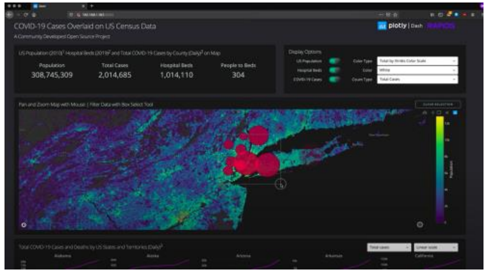

NVIDIA

OUR ENGAGEMENT IN XAI IN FINANCIAL SERVICES

Webinar series on XAI

frontiers in Artificial Intelligence | Artificial Intelligence in Finance

Editorial and publishing activity

Extension of a XAI use case of EU Horizon2020 project

Initiated a project on XAI

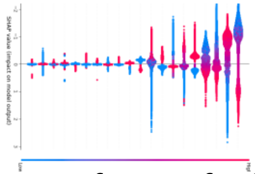

Accelerating open-source software for XAI