# European data spaces: towards sovereign and sectoral AI

## Conference report

Conference organized by Dauphine Governance and Regulation Chair

Paris Dauphine-PSL University, November 18th, 2025

DEBATE

Dauphine | PSL

CHAIRE GOUVERNANCE ET RÉGULATION

# Conference organized by Dauphine Governance and Regulation Chair

# November 18th, 2025

Synthesis n°105
Paris Dauphine-PSL University

# European data spaces : towards sovereign and sectoral AI

## Speakers

Jakob Rehof | Professor, TU Dortmund University; Director Lamarr Institute for Machine Learning and Artificial Intelligence; director Research Strategy, Fraunhofer Institute for Software and System Engineering

Emmanuel Bacry | Research Director, CEREMADE, CNRS Paris Dauphine-PSL; Chief Scientist, Health Data Hub

Tobias Guggenberger | TU Dortmund University; Fraunhofer Institute for Software and System Engineering

Hubert Tardieu | Member of the Board, GaiaX

Selma Souihel | Deputy Director of Inria's AI Program, INRIA

Anastasia Stasenko | Co-founder, Pleias

Laurent Lafaye | Co-CEO, Dawex

Claude Le Pape-Gardeux | Fellow Data Scientist, Schneider Electric

## Moderators

Eric Brousseau | Governance and Regulation Chair, Paris Dauphine-PSL University

Joëlle Toledano | Governance and Regulation Chair, Paris Dauphine-PSL University

# Summary

**Conference organised as part of Dauphine Digital Days 2025.**

**How can we build sovereign and sector-specific AI to strengthen innovation, autonomy and competitiveness in Europe?**

**European data spaces are paving the way for sovereign and sector-specific artificial intelligence, promoting innovation and competitiveness. This event brought together experts, decision-makers and industry players to explore the technical, economic and regulatory challenges involved in building these strategic infrastructures, which are essential for digital autonomy and data exploitation in Europe.**

# Introduction

## Eric Brousseau | Governance and Regulation Chair, Paris Dauphine-PSL University

The quality of AI engines is critically depending upon the quality of the data that are used to train them. While the US has invested heavily in computing capabilities and the development of very large language models (LLMs), Europe has instead focussed on the development of high-quality data spaces. These facilitate the training and fine-tuning of LLMs, and make it possible to develop vertical AI systems aimed to respond to specific needs. Even when there is a common political or industrial interest in data sharing, the concern for data security and the realities of industrial competition mean that it is not a straightforward task. As such, the governance of these data-sharing ecosystems is critical.

The first session will explore the interplay between data sharing and the development of AI systems, with particular attention to vertical AI, while the second session will focus on the development of sovereign AI systems.

# Session 1

# Industrial Data and Vertical AI

## Data sharing and Vertical AI : An Opportunity for European R&D

**Jakob Rehof | Professor, TU Dortmund University; Director Lamarr Institute for Machine Learning and Artificial Intelligence; director Research Strategy, Fraunhofer Institute for Software and System Engineering**

The relationship between data sharing and vertical AI provides a critical opportunity for European research and development. The Lamarr Institute, which is a joint venture between the universities of Dortmund and Bonn, is one of the new six research centres established in Germany since 2022. It has two Fraunhofer institutes: one focuses on research into logistics; the other focuses on machine learning and AI. The network brings together more than 40 principal investigators and professors and around 150 to 200 PhD students and postdocs. From the outset, the Lamarr Institute has sought to go beyond simple questions of data for machine learning by exploring the background knowledge that is specific to vertical domains, and the context of AI use, including resource use and sustainability.

The Institute's key research areas are horizontal, fairly generic fields that relate to the technical core of AI and machine learning and seek to bring relevant background knowledge into contact with the learning mechanism. Resource-aware machine learning aims to make machine learning more efficient. The recent DeepSeek development in China shows the potential to distil large LMMs into smaller packages that retain as much of the original functionality as possible. This could have a range of horizontal applications.

Trustworthiness is a horizontal topic that is relevant to almost any application of machine learning. Human-centred AI systems must be able to bridge the gap between human cognition and mechanical cognition. For example, they must be able to explain why a machine-learning model came to its conclusions or behaved in a certain way, or represent large data sets in ways that can be understood by humans.

### The role of vertical AI in science, engineering and industry

The Lamarr Institute also focuses on research topics with a vertical nature, including planning and logistics, physics, industry and production, life science and health, and natural language processing. Some examples will serve to illustrate the power of verticality in AI, especially in science, engineering and industry.

In physics, the IceCube project in the Antarctic used sensors to collect petabytes of data about neutrino radiation over a period of ten years. Physical theory was used to create a simulation model that was used to train a neural network to recognise the signatures of interactions that physical theory would predict to behave in certain ways under certain conditions. When the trained model was given access to the real data, it was able to identify the source of neutrino radiation in the Milky Way. This breakthrough was cited by the committee awarding the 2024 Nobel Prize for Physics as a prime example of what can be achieved by machine learning in physics.

The 2024 Nobel Prize in Chemistry was awarded to a team that had developed a machine-learning-based tool for analysing protein folding. The Lamarr Institute is using vertical AI to create a similar LLM-style model for particular protein structures. It could be used to study how such structures behave, and to discover new structures, for example, for drug discovery purposes. The third example concerns industrial production, specifically the ways in which production tools interact with physical materials. By placing sensors around the tool and the material and capturing large volumes of data about these interactions at a very fine-grained level, trained models can provide useful information about the conditions that could enhance or decrease the quality of the end product. Another machine-learning model could automate the process of optimising production processes.

## European data-sharing ecosystems and vertical AI

Machine learning techniques have been used for scientific and technical applications in production systems since before 2015, when the real revolution in this field began. For example, physicists at CERN used neural network technology and data analysis to deal with massive data sets produced by the Large Hadron Collider. Europe is knowledgeable about vertical, domain-specific applications of this type and has significant potential to tap into it through research and technical production.

Europe does not, however, have the resources required to compete directly with the US on extremely large, extremely expensive, horizontal, generic platforms. Instead, it has the potential to put mind over matter and mathematics over hardware in order to better marry insights and AI. While hardware investments will be required, value creation will come from human intelligence rather than pure hardware muscle. To achieve this and integrate vertical AI into industrial practices and organisations, data sharing – and the availability of ecosystems to enable data sharing – will be critical.

A number of computer science topics are important for vertical AI. Common domain ontologies are required to support the interpretation of data and ensure that all parties are able to understand and interpret shared information in the same ways. Shared data can be used to identify shared models for vertical purposes. For example, similarities between protein structures and linguistic structures should make it possible to fine-tune a general language recognition LLM to recognise protein structures. When a model is furnished with highly specific data, it can become highly specialised in a given field.

Retrieval augmented generation (RAG) processes aim to augment models by providing contextual information, such as specific documents, that pertains to a specific vertical application. To reduce the likelihood of hallucination and improve the focus of responses, a model designed to provide legal advice might have direct access to, and be directed to focus on, relevant legal documents. The organisation of these processes is a key topic within machine-learning studies.

Finally, agentic tools are required to wrap these machine-learning techniques in larger software systems that are accessible to domain experts. This might involve the creation of connections between specific workflow processes in a given domain and the integration of automated tools that enable people in that environment to access the resulting information. AI agents are simply software systems that take advantage of AI models to automate particular types of vertical process.

## Integrated data and AI value chain: leveraging high-quality vertical data

Data spaces like GAIA-X are important for facilitating trustworthy, secure and sovereign use of data. Governmental structures, policy structures and organisational structures must be designed to enable trust to develop between different parties, in order that they can confidently share data and use data provided by other parties. A pipeline approach could enable basic AI models to be fine-tuned for sector-specific applications and then, with further fine-tuning, become tailored for company-specific requirements. It is important, however, to consider the origins of the basic AI model. If the foundational model is a generic LLM created by one of the large American players, for example, there is a risk that it will not be suitable for use within a data ecosystem built on trust and data sharing. This is an interesting and complicated topic.

Domain-specific reasoning data exists within sector-specific AI models. With verticalization, it typically becomes necessary to augment these purely statistical models with some variety of logical AI. For example, a model working on physics requires access to foundational physical logic that establishes boundaries around the kinds of conclusions that can be drawn from certain data. Physical knowledge must be complemented by chains of reasoning, but statistical models tend to be less good at this type of logical reasoning processing than old-fashioned logical AI.

Again, there are many interesting questions in this field.

### Some research problems and opportunities

The concept of vertical AI in the context of data sharing raises a number of research problems and opportunities. Value creation can only arise when there is an interesting problem to solve and an interesting solution to offer within a productive context. As such, it is worth noting that machine-learning research has done little to exploit tabular data, such as information contained in databases. Most LLMs are trained on free-floating text found on the internet. Tabular data tends to be highly structured and the format and content of the table provides additional information and context. It appears to offer possibilities for machine learning that cannot be achieved with free-floating text, particularly in terms of the speed of the learning process and the quality of the results that can be obtained. Colleagues at the Institute are working on an interesting new technique, Tabular Prior-data Fitted Network (TabPFN), that has emerged in the last couple of years.

Federated learning research explores how systems can be developed to learn and create value from heterogeneous data sets without compromising the security and privacy of each set. For example, a vertical AI that learns from data from one hundred hospitals will be richer than a model that learns from one hospital's data. However, the data from each hospital must remain private to that hospital, and the model must not be able to reconstruct sensitive data from each of those sources. Research that can solve this challenge will be valuable for the construction of AI ecosystems. An AI ecosystem should go beyond data sharing and start to learn as an ecosystem.

There are significant issues around hardware in Europe. Although the EU system is enabling fairly significant investments in expensive, specialised hardware for machine learning, it falls a long way short of the astronomical investments being made in the US. Even if it puts mind over matter and mathematics over hardware, Europe will still need ramp up spending on hardware if it is to remain in the competition. No single entity in Europe has the economic clout to do this by itself; resources must be pooled. New giga-factories and AI factories across Europe will need to be able to work together. The idea of creating shared ecosystems will need to go beyond data and expand into hardware.

Finally, automated data interoperability is an interesting challenge for research and development. Effective data sharing requires shared methods of interpreting data that can be represented in myriad ways and come from myriad sources. It may be possible to use generative AI (genAI) to solve this problem automatically, for example by generating code that can transform between different representations of the same information, or map datasets from a vertical domain into a shared abstract information model. There is potential to use genAI models to solve the data interoperability problem and inside data-sharing software agents. A positive feedback loop of this type would solve a multitude of problems, particularly if the ecosystem was fed by a horizontal LLM foundational model based on trusted data and everything happened within a trusted ecosystem.

Vertical AI is an important project, particularly for Europe. The time for action is now.

## Data sharing and AI in healthcare

### Emmanuel Bacry | Research Director, CEREMADE, CNRS Paris Dauphine-PSL; Chief Scientist, Health Data Hub

Health is an excellent use case for AI. It is an area where AI can have many applications, potentially with significant benefits for us all, but also where problems with AI come into sharp focus.

### Significant promise

AI is already used for a number of purposes in healthcare. For example, while a senior radiologist always performs the second reading of a mammogram, the first reading is now often performed by AI rather than a junior radiologist. The AI is competent, but this change makes it difficult for junior radiologists to acquire useful experience. AI is also used in cancer prevention, by predicting which women are more likely to develop cancer in the next five years so they can be monitored more closely. There is huge potential for AI to deliver personalised and precision medicine, enabling drugs and pathways to be designed for individuals rather than populations or pathologies. Personalised cancer treatment pathways are already being delivered based on the genetic characteristics of a patient's tumour. Public health, pharmacovigilance and pharmaco-epidemiology are being improved with the development of huge databases, such as the French Carte Vitale database. This database provides a snapshot of the healthcare pathways of almost everyone in the country. It has already been used to identify drugs with hidden side effects and underpinned some of the world's largest Covid studies. When used with care, genAI also has significant potential, for example to support doctors with report writing and diagnostic brainstorming. There are major potential applications in drug discovery, drug design and clinical trials.

### Significant issues

Although AI offers many potential benefits, it is not without problems. The 'black box' issue is extremely significant for health applications. For example, several years ago, a deep-learning algorithm was developed that appeared to successfully detect pneumonia from x-ray. It emerged that it was actually focusing on the corners of the x-ray, where the name of the hospital and type of x-ray machine was coded, and predicting pneumonia based on the number and severity of cases that the hospital typically received. It is extremely important to know what the AI is doing and how it works.

The Carte Vitale database contains the data of around 68 million people in France. Large US insurance databases contain the data of around 20 million people, most of whom are relatively rich, young and healthy. All data has bias, but this level of significant bias could affect outcomes. We are developing an occidental AI, especially in health, because our databases are dominated by white, male Europeans.

There is a tendency to fail to consider the hidden aspects of AI. It has been shown that, from a chest x-ray, an AI can predict the ethnic origin and approximate wealth of a patient with a reasonable degree of accuracy. If you train AI to discover a pathology, there is a risk that bias will feed into the process.

The final major issue with AI in health relates to evaluation. When an AI is deployed, particularly in health, it is deployed as part of a chain that includes humans. Papers tend to evaluate the precision of the algorithm, but they fail to consider its performance within a chain that involves humans. An AI was deployed to alert for severe kidney failure in the intensive care units of six hospitals. A randomised trial assessed 3,000 people with the alert system and 3,000 without. In two of the hospitals, there were twice as many deaths when the AI alerts were in place. There was no clear reason for this, but it must be assumed that it was due to the human factor. Studies show that AI models can have very different levels of success depending on who is using them.

### Data sharing

Most of the AI revolution in health and other verticals is based on data. New methodologies are relevant, but data is central. Three issues are universal to almost all verticals. First, data fragmentation: knowing where data are and what they are is a huge challenge across industries and institutions, and at national and European level. Second, governance is a huge problem at a global level, particularly in health data. Data holders tend to feel that they own the data, they understand that it is valuable, and they are reluctant to share it. During Covid, this sense of ownership was extremely difficult to manage and limited data sharing between institutions, researchers and private companies. The more important the data, the less available it was. This is a major issue that needs to be addressed.

The third issue is interoperability. Some aspects of this problem cannot be solved automatically.

During Covid, there were attempts to monitor the dynamic of the pandemic in France by tracking data from the emergency services. This quickly hit a major stumbling block: all of the emergency services were coding Covid differently, some when they had done a test, others when they identified certain symptoms. This is not a challenge that can be resolved with an LLM. The same applies to rare diseases: if different people use different definitions, it is not possible to do anything with the data. The Health Data Hub is doing its best to address this and other challenges.

### The European Health Data Space

The European Health Data Space is a unique initiative in Europe that aims to solve the challenges of the data revolution. This ambitious new regulation was set up at roughly the same time as the AI Act and is now being rolled out. It has two aspects: primary use of data and secondary use for research. It states that all health data, whether it is produced by public or private institution, must be shared for public-interest research. This is very significant: data from clinical trials, genetic data, data from devices and so on must all be shared for public-interest research, which is a very wide concept that encompasses research by private companies.

The regulation will set prices for accessing data. A key principle, which is very sensitive, says that while providers can cover their costs, they are not allowed to make a profit from giving access to data. It also sets a maximum delay for providing access to data. Researchers that wish to use this space will be able to run queries on the meta data through a central portal in each member state. This has been designed to address the issue of data fragmentation. They will then be able to request data through a single form, which will apply to all of the databases across all of the member countries. Each body is required to respond to the request within four months. If the response is positive, they must provide access to the data within four months or pay a fine. The data space will be extremely difficult to deploy and the timelines are extremely tight. I am not confident that it will work, but it is a unique initiative that is attracting a lot of interest from across the world. I hope it will succeed.

## Data spaces: the competitive advantage of Europe for AI

### Hubert Tardieu | Member of the Board, GaiaX

Europe needs to prepare for vertical AI and ensure that at least some of its data spaces will be economically viable. Data is vital and these data spaces are central to Europe’s future competitive advantage. Nothing will happen without massive investment at national and European levels.

The European data union has a data-space reference architecture, support centres and a large international manufacturing community. Two standards are being prepared on trusted data transactions and data spaces. A major data space has been established in the automotive sector in Germany and two are being built for the aeronautics domain, where Europe is the world leader, and the nuclear domains, where France is in the top three countries globally. Establishing data spaces will help to maintain these positions. Data will be shared through clearing houses, with automated compliance.

### The challenges ahead

Four years after work on data spaces started, it is clear that many productive data ecosystems lack viable business models or suffer from slow adoption. It seems clear that a data space should not be subsidized unless there is clarity about potential use cases and about mechanisms and timescales to achieve economic viability. It makes no sense to establish data spaces that fail the moment subsidies are withdrawn. Together with Dauphine, Gaia-X has identified trajectories that are likely to result in viable ecosystems but further work is required to answer with clarity on all cases.

## Data spaces: the state of play

The phase of working through ecosystem formation, technology stack, feasibility and pilots, and initial investment is over. It is not feasible to continue to finance projects in Europe in this way without clarity about sustainability, business viability, economies of scale, and return on investment. The Gaia-X governmental advisory board met last week and agreed that data spaces cannot be subsidizedsed unless there is clarity about economic viability. This is a complete change of attitude and mentality.

## Why: the need for continuous support and funding

The Draghi Report highlighted ten domains – including automotive, aeronautics and nuclear – where vertical AI is necessary. To achieve this goal, Europe must identify solutions in these areas within the next three years and make ongoing support conditional upon the viability of the project.

Taxpayers' money must be invested responsibly. France and Germany are taking different approaches to this challenge. In France, any project that cannot demonstrate viability within a 36-month break-even period will be discarded. In Germany, the government's huge investment in Catena-X was intended, from the outset, to prioritize regulation use-cases, typically around carbon footprint or digital product passports. Aeronautics and nuclear tend to focus more on product than regulation.

Through our work with Dauphine, we have identified two markets: user communities that are willing to exchange data at a price that is reasonable for them; and data-sharing intermediaries that are able to facilitate data-sharing at an acceptable cost. Again, there is a difference of attitude between France and Germany. Germany has invested heavily in their intermediary and are now trying to shore up its economic viability by extending its reach from the automotive industry to the chemical industry. This is a very complex challenge.

## What: industrial AI components

The German government is keen to marry the three nodes of infrastructure, foundation model and data ecosystem through a very large project involving at least four European countries and the European Commission. Although Europe will remain behind the US on infrastructure, it has the potential to be good in foundation models, especially open source models. Europe must ensure that it has economically viable data spaces in each key vertical AI domain within the next three to four years. A huge amount of work remains to be done: positive intentions must now be translated into real projects. Data spaces are key to vertical AI ecosystems in terms of data sovereignty and data quality.

## How: from follower strategy to leading edge

Due to its competitive advantage in data sharing, Europe still has a slim opportunity to move from a follower strategy to a leading-edge position. It will require viable business models, regular health-checks, and the development and application of business model patterns for different types of ecosystem. While large companies with deep pockets, like EDF and Airbus, can launch data spaces and resolve any issues that arise, the data space for European agriculture collapsed in 2024 because its participants were primarily farmers rather than financially secure machinery manufacturers.

To scale adoption, there is a need to increase ecosystem participants' data readiness and establish cross-ecosystem governance to drive the reuse of commodities. Ultimately, vertical AI will need to develop ontologies and interoperability. As the economic viability of data spaces is now being prioritized, applications are now being developed by use case. After 36 months, there will be data definitions that are local to each use case. This is one of the main tasks that is being undertaken with Airbus. Data and AI economy is also extremely important.

### Recommended action for ecosystem stakeholders

Over the next six months, we must ensure that policymakers and governments in at least four major European countries, plus the Commission, agree on this approach. It is devastating to see the Commission issuing calls to tender that request a list of use cases as a deliverable after three years. This phase is finished: we know how AI works and we need real spaces that work. We must ensure that SMEs join these data spaces, as they generate most of the data but often lack technical maturity or have concerns about sharing data with large competitors. We believe that EDF's goal of having 20% of its SMEs connected to its new data space after three years is realistic.

The idea that it will be possible to deliver a robust ecosystem after three years is becoming widely accepted in Europe. We believe that it is the way forward.

## Why AI Needs Data Spaces – and Europe Needs Both!

**Tobias Guggenberger | TU Dortmund University; Fraunhofer Institute for Software and System Engineering**

### Europe's strategic context

Although Europe has a very strong industrial history and well-established industrial sectors, it also suffers from significant fragmentation: there are lots of SMEs and fewer large companies, especially in the digital economy. Compared to the US and China, Europe lacks innovations that are market-ready and have the potential to grow the digital economy. Data remains siloed and many organisations have low data readiness, fail to collect usable data or struggle to share data. Organisations also refuse to share data for a number of reasons, including intellectual property rights. The market value of open data – excluding private data – in Europe was estimated at €184bn in 2020. AI leadershipcorrelates less with intelligence than with data concentration. Data network effects mean that the more data is collected on a platform, the better the processes and the user experience. Thus, data concentration directly leads to a competitive advantage.

At present, the European response to these issues is to build sovereign data infrastructures composed of a landscape of data spaces. Many European initiatives are working to promote this and drive forward European digital sovereignty, or strategic autonomy, based on a shared vision of a federated, trusted infrastructure called the Single Market for Data. Foundational work on interoperability, governance, connectors and usage control is being developed in these data spaces, alongside the technical and organisational structures required to share data.

### The structure problem

AI, particularly vertical, sector-specific AI, has a structural problem in that it usually fails the tabular problem and struggles to make sense of industrial data that is organised in certain ways, such as timelines. Furthermore, model providers lack high-quality data for training on structural problems of this type. Legal barriers also limit the ability of organisations to introduce these data into training sets for LLMs. Due to the low availability of this type of training data, SMEs and public institutions struggle to train models that could support for business operations. These factors have created a structural issue for Europe, which is losing traction on a daily basis.

The use of sector-specific, industrial AI is blocked by three factors. First, unclear data rights and obligations that are often caused by poor data governance within companies. Second, a lack of mutual trust frameworks that clarify the actors and components within a given system and show who has access and how it is being used. Although organisations like Gaia-X are working to resolve this challenge, many high-value data in Europe remain completely unused. Third, there is a need for a sovereign, scalable and inclusive data foundation that can be used to enable AI across all sectors. This will require four layers: data ecosystems where data spaces and platforms compete; business applications that can be used inside and outside the data space; transactions between components; and an infrastructure layer for global data-sharing, which includes computing capacity and the infrastructure required to connect different entities into a single ecosystem. Work is underway to create a European ecosystem that will make this infrastructure possible.

## Potential use case and conclusion

Manufacturing is a heavily fragmented sector with many small companies that employ 10-20 workers. These companies do not collect data but do need to maintain machinery. Data sharing to support federated learning could be very useful in this environment, as it could enable companies with a single machine to improve performance by learning from multiple machine data points across Europe.

If we can make the data spaces use and include AI, and ensure that the data is available for AI, we can create a trusted environment for data sharing. This environment would not compromise sovereignty, but it would enable transactions within sectors and create a business layer that supports multi-actor AI value creation in order to increase the size of Europe’s economic cake.

We explain this to industries that are developing data spaces. Their first task should be to develop relevant use cases that add value, as this incentivises parties to contribute with data to those data spaces, and to foster adoption of data sharing.

# Session 1

## Discussion

### From the floor

What is the impact of the Digital Omnibus? How can we safeguard European data spaces against the influx of American and Chinese companies?

### Joëlle Toledano | Governance and Regulation Chair, Paris Dauphine-PSL University

This is a very technical question. Who is the orchestrator of the European Data Health Hub?

### Emmanuel Bacry | Research Director, CEREMADE, CNRS Paris Dauphine-PSL; Chief Scientist, Health Data Hub

This is a regulation that has been voted by the European Commission.

### Jakob Rehof | Professor, TU Dortmund University; Director Lamarr Institute for Machine Learning and Artificial Intelligence; director Research Strategy, Fraunhofer Institute for Software and System Engineering

This is a question about protecting data against American hyperscalers and other similar outsiders. I see this as one of the fundamental reasons for having data-space governance structures that allow us to establish rules that must be respected by anyone participating in these data-sharing ecosystems. These parties can come from any country and already include large companies like Microsoft and Huawei. Gaia-X brings together policy, rules and architecture, and a trust framework to establish a complex set of rules that people must respect if they want to operate in this space. As long as those rules are respected, it does not matter where a company is based or who participates. That is the basic logic of how to deal with this challenge within data ecosystems.

### Emmanuel Bacry | Research Director, CEREMADE, CNRS Paris Dauphine-PSL; Chief Scientist, Health Data Hub

The exact impact of the Omnibus is not yet clear. It could have a significant impact on health data, but that will depend on details in a law that has yet to be voted. Almost all personal health data are treated as pseudonymised data under GDPR, which means they are sensitive data that are subject to specific rules. This is not the case for health data in the US. This is a major issue for genAI, because if an LLM sees a medical report – even pseudonymised – the LLM itself must be treated as sensitive data that cannot be shared. The Omnibus says that pseudonymised data could be anonymised for some people, so the status of the data will depend on the user. We need to understand this is in more detail.

We all say that data is extremely important, which is true. But we must remember that the basic act of organising data spaces with high-quality, interoperable data could result in significant progress and have a major impact on public health. Introducing AI and new methodologies could offer additional gains, but we must not overlook what we could achieve simply by bringing together high-quality data.

### Eric Brousseau | Governance and Regulation Chair, Paris Dauphine-PSL University

We explain this to industries that are developing data spaces. Their first task should be to develop relevant use cases that add value, as this incentivises parties to contribute with data to those data spaces, and to foster adoption of data sharing.

### Hubert Tardieu | Member of the Board, GaiaX

I agree that we should not rush into expensive data centres until we understand the data that we are using. In the case of the data spaces that are being created for aeronautics and nuclear, the use of American cloud providers is another relatively important challenge. Airbus and EDF would like the data intermediary to organise a multi-cloud approach within the limits of the regulations.

In France, people speak about digital sovereignty but in Germany, they talk about data sovereignty. The distinction between these terms is important: digital sovereignty refers to the tools being used, while data sovereignty is about mastering the data that is being used. We must start thinking in terms of data sovereignty if we want to respond effectively to the challenges we face.

### Tobias Guggenberger | TU Dortmund University; Fraunhofer Institute for Software and System Engineering

We must increase data readiness to facilitate the sharing of useful, good quality data. Organisations can be reluctant to share poor-quality data. Use cases that are not for the public good should be funded by private money. We must not only depend solely on public money. The governance and management of these systems must take account of value creation and value distribution systems in order to ensure that the data economy is fair and attractive to private funding. Otherwise, people will not want to participate, and financial sustainability is not reachable.

# Session 2

# The challenges of sovereign AI

### Joëlle Toledano | Governance and Regulation Chair, Paris Dauphine-PSL University

This round table discussion will consider what sovereign AI is and the challenges it might pose for public policy, the development of various activities, and national and international businesses. We will explore different perspectives and the implications of various options.

## An ecosystem approach to sovereign AI

### Selma Souihel | Deputy Director of Inria's AI Program, INRIA

Since 2018, INRIA has been responsible for coordinating the research component of the National AI Strategy. INRIA is building sovereign AI through an ecosystem-design approach that integrates everything from skills and software companies to compute and evaluation. It will have applications from industry partnerships to public policy alliances. INRIA's coordination mandate was renewed in 2024 and extended to include higher education and research. This was a turning point for INRIA, which was traditionally a research institute but is now tasked with structuring a national project in light of AI's relevance to sovereignty and technological autonomy.

#### The challenge

INRIA's mandate is clear: to build a sovereign, competitive and innovation-driven AI ecosystem for France and Europe. This is not AI in isolation. The first task is to build a public-private ecosystem rooted in science, connected to industry and aligned with society. This reflects the philosophy of the program agency that INRIA has built in recent years: its architecture mirrors the reality of modern AI as a continuum that starts from AI, takes in cybersecurity, cloud data infrastructure and high-performance computing, and is closely tied to applications such as healthcare.

The agency's mission is to work alongside the state to develop a long-term vision and manage national programs involving academics, universities and the industrial sector. Its addresses digital sovereignty challenges to strengthen society, increase competitivity and transform industrial capacity.

#### Skills

Sovereign AI requires a systemic approach. Skills are a critical priority. France has a distributed academic backbone, strengthened by AI clusters, that is working to educate, attract and retain talent at scale as a core pillar of France 2030 and the national AI programme. Hubs, universities and national labs have a clear target to move from 10,000 trained students each year toward an annual capacity of 100,000 by 2030, while anchoring research, industry collaboration and international visibility. INRIA plays a role in the coordination of AI clusters.

Research structuring is a second key aspect of skills. INRIA co-leads the national priority research programme on AI alongside the CEA and CNRS. This multi-year programme targets scientific foundations through embedded AI and trustworthy AI, with the purpose of structuring scientific communities and funding frontier research.

Finally, skills are underpinned by organisational understanding. The LaborAI programme, developed in collaboration with the employment ministry, explores the impact of AI on work, skills and social dialogue and connects researchers with HR functions, unions and public administrations. To build sovereign AI, we must build sovereign skills and understand how AI is reshaping organisations.

### Public-private partnerships

There has been a critical shift in INRIA's approach to innovation and partnership building. It has been reoriented towards strategic partnerships with French actors, primarily strong technology actors. These bilateral collaborations aim to root AI in the fabric of French national industry and enable co-development, technology transfer, industrial adoption and real-world validation of new models.

INRIA has various accelerators and incubators in its ecosystem. The INRIA Startup Studio works to accelerate start-ups coming directly out of public research and build initiatives with strong IP strategies, long-term industrial technologies and positive national outcomes when acquired or scaled by French industrial leaders. It is also working to strengthen the deep tech pipeline. Sovereign AI needs sovereign start-ups as a strategic pillar, not as an afterthought.

### Software infrastructure

In terms of software infrastructure, the P16 project is the 16th measure of the national strategy for AI. This strategy aims to build digital commons for machine learning and AI, leveraging the success of previous projects such as scikit-learn which showed that high-quality, open-source software developed through public research can have a real impact. INRIA is now working in collaboration with Probabl, an INRIA spin-off that aims to experiment with a new governance and industrialisation model. This model involves the development of a set of services and products that accompany the use of open-source, open-access libraries that are applicable by the industry. Ecosystem thinking is central to these projects: libraries do not live in isolation and their success depends on integration with compute platforms, services and the developer communities.

### Compute resources

AI needs compute. It must be compute that we design, understand and control, not compute that is rented on the other side of the world. INRIA is meeting this need co-leading AI Factory France, a national hub that brings together a number of leading French universities, industry partners and public institutions. AI Factory France provides access to vital supercomputers, as well as training, technical support and pre-published AI models. Its ambition is to offer a sovereign continuum where compute, data and expertise serve researchers, start-ups, industry and government. In the future, there will also be a need for AI factories and specialised data centres optimised for training large models with energy-efficient architectures.

### Funding

The next pillar of the strategy is funding. Rather than attempting to build a sovereign AI agenda with a traditional funding model, we will need to implement mission-driven funding with strong, bold goals, fast, agile decision cycles, and the capacity to stop or to accelerate projects. This is the logic that underlines an ARPA-style project factory that INRIA is trying to implement to support R&D projects that have disruptive potential and sectoral impacts that could contribute to national and European sovereignty. Sovereign AI depends in part on the ability to taking risks by funding disruptive projects.

In terms of sectoral impact, INRIA aligns with public policy priorities for sectoral applications. If AI sovereignty is confined to a technological bubble, it will be meaningless. Instead, it must be constructed through real-world applications in sectors where national resilience is at stake. For example, in 2020, INRIA established a security defence mission to tackle critical challenges in cyber defence, cyber resilience, multi-modal data processing, trustworthy AI and multi-agent and robotic systems. It combines INRIA's scientific expertise with defence actors' operational requirements to address genuine operational needs. In health, where France is now recognised as a leader in federated learning, INRIA is working with the hospital network in Marseille to build federated infrastructures capable of powering clinical research, diagnostics and personalised medicine without exposing patient data.

### Regulation, trust and evaluation

At INRIA, we strongly believe that innovation and regulation are compatible. To implement regulation, we need innovation. There is a need for technical systems that are capable of evaluating AI in terms of robustness, risks, fairness and environmental impact. To meet this need, the National Institute for AI Evaluation and Security (INESIA) was launched in 2025. It aims to bring together expertise on system resilience and threat analysis with scientific know-how, AI research, mathematical modelling and regulatory and normative insight. INESIA is also working closely with the market surveillance authorities through a recent strategic partnership with the French Data Protection Authority.

### Collaboration

Sovereignty is not self-sufficiency or protectionism: alliances, collaboration, and the ability to choose and negotiate with partners are all critical. France and Germany have decided to work together to advance a joint vision for sovereign AI; INRIA is contributing to their strategic agenda on open, trustworthy European AI. Franco-German momentum is embodied by INRIA's collaboration with the Fraunhofer Institute, a pillar of applied research in Germany.

Investment in multi-lateral frameworks is also essential. France is working with Canada and Japan to build a global partnership on AI. INRIA coordinates the scientific aspects of this collaboration, which is developing into a new centre of expertise dedicated to global policy guidance, safety norms and evidence-based regulations. Europe must be able to offer a third path between the surveillance-driven AI model of China and the platform-dominant, attention-seeking AI model of the United States.

### Cultural issues

A sovereign AI cannot only be about experts. It also requires a digitally literate society, as citizens who do not understand AI or see it as a black box will feel excluded and drive a backlash in education, the workplace and democracy. Technologists have a responsibility to explain, to engage, and to make AI approachable for all, particularly given the rise in societal questions around the impact of AI on attention, learning and misinformation. We are only at the beginning of this process. There can be no sovereign AI without societal and social trust.

### Conclusion

France and Europe face a simple choice: use other nations' AI or shape the next technological wave. Our French policies, programs, partnerships and research start-ups are all aimed at achieving one goal: to make Europe a place where AI is invented, not imported.

## Tomorrow's Sovereign AI: Agentic SLMs, Synthetic Environments and Context Infrastructure

### Anastasia Stasenko | Co-founder, Pleias

It is a pleasure to be here and to share some pragmatic ideas of what sovereign AI could look like. Pleias is a very small start-up based in France, which was created 18 months ago by three co-founders who had developed use cases for the French government, particularly for the "Albert" AI which is already being used by almost 10,000 French public servants. We established Pleias because we saw a need for more transparent, more efficient models and because the devil is in the data.

### A data-efficient approach to training

The vast majority, almost the totality, of AI labs pre-train their genAI models on web archives.

Web archives are a source of massive amounts of data and have enabled these labs to prove that, with sufficient data and compute, it is possible to train models that can generalise. This is great, but it also poses problems in terms of legal compliance. It is not possible to filter copyrighted content at scale or ensure that models are not being trained on copyrighted data. Even a carefully constructed, accessible corpus will contain copyrighted material. This is an issue in Europe, where the concept of fair use does not apply. These issues around legal compliance and copyright were compounded by the terms of the AI Act and AI Code of Conduct.

At Pleias, we were also keen to hack the system a little and identify more data-efficient ways to train models which would perform well in enterprise settings. Companies do not speak like Reddit users and they do not need models that cover the full scope of human knowledge. Instead, they need models that can work with specific data, including legalese, and comprehend complex PDFs. This data is not readily available in web archives and most web crawlers do not crawl PDFs that are bigger than five megabytes. Some of the data in these PDFs will get into a model; most of it does not.

We explored the potential to deploy a data-centric, provenance-based approach. In other words, we choose our data sources and ensure that they correspond to certain criteria that we identified at the outset: all data should be copyright compliant under European law, of good quality, pertinent and diverse. By diverse, we mean that it should be multilingual and include both legal and cultural reasoning. We were able to do this by creating a common corpus which has been used to train more than seven sovereign LLMs. In this instance, 'sovereign' refers to projects developed by nation states. These countries include Switzerland, the Netherlands and Spain.

Our experience showed that it was possible to be very efficient with data and develop small models that would not necessarily be able to answer every possible question in a B2C environment, but could be a highly efficient solution in enterprise settings, particularly for specialised use cases.

### Sovereign AI and competitive AI

It is interesting to consider what sovereign AI could and should mean in the near future, specifically in Europe. When people talk about sovereign AI, they are usually thinking about analogues of the frontier models that might replicate the performance of a very big model like ChatGPT. It should be noted that ChatGPT still cannot do a lot of things, which may explain why studies show that enterprise adoption is still fairly low, particularly for agentic AI.

It is also worth asking what competitive AI could be. It can be tempting to think of competitive AI in terms of geopolitical competition around model parameters. We believe that competitive AI should empower enterprises in Europe to build competitive offerings, based on a competitive and nuanced business model. Frontier models, like ChatGPT, Claude or DeepSeek, would not be able to replicate the offering of this competitive AI native.

Data spaces and data play a huge role in this vision, but raw data is not sufficient. Reasoning models need to know and memorize data, but they also need to follow specific reasoning patterns that are introduced alongside the data. Health records, for example, include information that has been deduced and transformed but which does not necessarily appear in the final database. As well as knowing the data, models must be able to reason with this data as a starting point. We are currently trying to develop this kind of model, namely a very small model that is trained exclusively on data pertaining to the use case, with data volumes augmented using synthetic approaches.

### Synthetic data

To create a synthetic environment, we rebuild from scratch the environment in which our model, reasoner or agent must exist, operate or behave. This might involve rebuilding Salesforce or alegacy software environment or any product that constitutes a company's competitive advantage. The idea is to create a synthetic environment where data can be modelized without interference or monitoring by OpenAI or any other third party. Models follows different paths and take different decisions within this environment and a small model is trained using the resulting data.

We have used this approach, which is fun and technologically feasible, to train a 50-million parameter model and a 300-million parameter model, which is two times smaller than ChatGPT-2. These models were trained on the 50,000 articles that make up the cognitive core of Wikipedia: the articles that Wiki admins believe are the most important articles of knowledge for humanity. We augmented these data synthetically to develop 200 billion tokens, which is a relatively small number by modern standards. By training our models using these tokens, we have been able to reach standards on industry reasoning benchmarks for maths, like MMLU and GSM8K, which models trained on over 10 trillion tokens cannot achieve.

This exercise demonstrates that very small models can be competitive in specific use cases. We believe that we must consider whether, for the whole field of AI, sovereignty might be best attained through very specialised, decentralised small models that could drive value creation that benefits a wider audience than a couple of platforms or AI labs. We also believe that there is still huge potential to improve model training and that the future of agentic AI should be based on a large number of highly specialised small language models. For industries, big enterprises and vertical markets, this could generate a real competitive advantage and demonstrate that a sovereign AI is, indeed, better than a generic API from OpenAI.

## Beyond slogans: practical realities underpinning sovereign AI

### Laurent Lafaye | Co-CEO, Dawex

The recent reports by Enrico Letta and Mario Draghi are driving essential conversation in Europe at present. Both conclude that, if Europe wants to stay in the game, it must invest massively across the entire AI value chain, from network and cloud to data, compute and skills. But this grand macroeconomic ambition meets a very concrete question: Can we demonstrate, end-to-end, what our AI systems truly depend on in terms of data and software?

#### What is sovereign AI?

Sovereignty is everywhere today: we talk about energy sovereignty, industrial sovereignty, digital sovereignty and now AI sovereignty. Instead of treating it as a slogan, let's explore a practical definition. Sovereign AI means three things: Understanding and controlling to identify the origin of data, infrastructure, and software; the auditability and contestability to be able to explain and challenge AI actions, as required by the AI Act; and the selection of the dependency selection to retain the power to decide with whom to cooperate, on what data, and according to what rules and governance.

The Letta report highlights the need to invest in digital infrastructure, including data spaces. "Sovereign AI" will not emerge simply by funding more models or graphics processors. The key topic is traceability: it relies upon our ability to provide answers, at any time, to the three points outlined above.

#### Sovereignty through traceability

The idea of traceability has two aspects which are often discussed as separate concerns but which, for AI, must be considered together. The first is traceability of data. 'Data traceability' often makes people think about log, metadata, data lineage and tools. These are necessary but not sufficient for sovereignty, liability or compliance. From the point of view of a regulator or even a judge, the core questions are more legal and economic. With questions like, who provided the data, under which contract, which access policies, and for which purpose. How are these constraints propagated in the chain, when different datasets are combined, enriched, re-packaged and used to train or retrain models?

In other words, there is no credible data traceability without explicit contractualisation of access and use. In a data space or data ecosystem, one AI system will typically use internal data, data from other companies, data from public bodies or regulators, or specialized third party data. Each of these sources comes with a contract or policy, such as a sectoral limitation, geographical limitation, reuse prohibition, data protection rules or trade secret clauses. In practice, traceability will be weak if this information resides only in human-readable documents. Under these circumstances, it is not possible to say in real time whether a given model respects all of the conditions attached to each data exchange.

We are living through a critical transition moving from human-readable specifications to precise automation-ready, machine-readable specifications. That's where we at Dawex try to contribute. Concretely, this means that contracts, terms of use, internal policies and reference frameworks will need to be modelled and turned into machine-interpretable policies that data exchanges and data spaces can apply automatically.

Once every data transaction is linked to a machine-readable policy and every use of data via AI pipeline leaves a trace link to that policy, it will become possible to ascertain whether a given model was trained in full respect of the conditions attached to each data source and that no data was used outside its contractual perimeter. It will be possible to respond to situations like a data provider withdrawing consent or a change in the law. Sovereign AI in this strong sense requires the ability to tool the legal and technical traceability of data coming from many sources and multiple providers. This is consistent with the European idea of Common Data Spaces, i.e. an architecture where governance and usage rights are the core of the design, not afterthoughts.

Software traceability is the second layer in the software supply chain. A modern AI system runs on a cloud stack, a data stack and an AI stack. Each of these layers uses hundreds or thousands of open source components. If we speak honestly about sovereignty, we must also ask very basic questions here. It is important to know which open-source component the AI system release depends upon and how the components are developed and maintained. We need to know which communities, countries and firms have an interest in investing in these bricks and where our critical points of dependency lie in a given ecosystem. For AI, we also need a full software lineage for AI systems, just as we need a lineage for data, we need lineage for the code upon which our open-source communities, countries and AI systems depend. This is very important both for security and for industrial policy.

If Europe wants to follow Draghi's recommendation to close its competitiveness gap in digital technology and adopt Letta's ideas of freedom for research, innovation and data, then we must decide which open-source stack we want to invest in. If we combine these two dimensions, then trust in AI decisions depends on data traceability and software traceability. The key to both is the same: a transition from human-readable description to machine-readable specifications for both data contract and software descriptions so that traceability becomes automatable, auditable and enforceable.

### Measuring and steering dependencies

It is interesting to consider how these dependencies might be measured and steered. The question of AI traceability is close to a classic economic intuition: you can only control what you can measure. The new French Digital Resilience Index (DRI), launched by an alliance including RTE, Docaposte, Ascend partners and the think tank Digital New Deal, aims to measure the digital dependencies of organisations in terms of software, data, infrastructure, technology, core assets, internal skills, governance and resilience to shocks. It is designed as a reference index, with a specific committee; from 2026, it should be available as a common standard that European companies can use to map their dependencies and their resilience.

This development shows that we are entering a new phase. Rather than simply saying that we need more sovereignty, we are starting to measure and map dependencies precisely, and to use them as steering tools in boards and risk committees. I believe that we need similar, explicit metrics for data dependency in AI. For example, we should know what share of our critical data depends on non-European actors, and have substitution scenarios ready in case our access is cut or restricted. Metrics for software dependence might consider which non-substitutable components are in our AI pipelines, or assess our exposure to unilateral changes in license or governance. Metrics for institutional dependence could ask to which standard contractual framework and jurisdiction our data and AI governance schemes are effectively tied.

If we agree with Letta and Draghi that Europe must invest massively in AI, then part of this investment must go into foundational elements like data spaces and data exchange infrastructures, tools for data and software traceability, and metrics and indices for digital dependency and resilience. Data spaces are, from this perspective, observation points on data transactions, contracts, policies and levers for actions because they can implement machine-readable policies, link data exchanges to AI requirements and integrate resilience metrics in the infrastructure itself.

### Conclusion

I will conclude with three messages. First, sovereign AI is not only a technology issue, it's also an institutional architecture problem. It depends on how we define rights of access to data, how we write and apply contracts & usage policies, how we govern data spaces and manage software and open-source dependencies. Second, traceability is at the heart of AI sovereignty. Traceability of data transactions, based on contracts and policies, as well as traceability of the software supply chain, so that we really know which code, which communities and which countries our AI systems depend upon. Third, we need indicators and investments at the right level, this requires indicators of dependence and resilience, such as DRI, as well as the large-scale investments that Letta and Draghi recommend but directed to less visible layers like data governance, data space infrastructure, traceability and audit tooling.

As a technology leader, Dawex works on the infrastructure for data exchanges and data spaces. Academic research in economics, computer science and law is essential to formalize this notion of dependence and resilience, design incentive and governance models that work in practice, and steer public and private investment towards the right links in the AI value chain. If we manage to link this macro-economic vision (Letta, Draghi), diagnostic of resilience (DRI) and concrete infrastructure (data spaces, data and software traceability) then by the end of this decade, we may talk less about "sovereign AI" as a slogan and more as a conscious choice: building and operating AI in Europe, with evidence backed up.

## Data and Artificial Intelligence for Energy Management

### Claude Le Pape-Gardeux | Fellow Data Scientist, Schneider Electric

Sovereignty is usually thought of in geopolitical terms. Schneider Electric has activities throughout the world; it also has significant expertise in data and AI. Although there are some frontiers between countries in terms of our relative abilities and capacities to build our own AI systems, a cursory examination soon reveals that the frontiers are not only between countries, but also between us and our customers, us and our partners, and us and our suppliers.

Our energy management solution business aims to help our customers make the most of their energy. This might involve identifying situations where their energy use is not efficient and finding ways to eliminate waste. To do this, we use AI to compare past data with current and comparable data. We also work to optimise energy consumption. Energy is generally consumed for a good reason, such as keeping a building warm, manufacturing products in a factory or delivering services. Data about energy consumption can be commercially sensitive and our customers are sometimes reluctant to share them. A steel producer, for example, might be reluctant to share energy data as it immediately reveals a lot about their production methods and the state of their business.

The energy sector is dealing with very significant changes. In the past, electricity was produced by a handful of major utilities, which sent electricity into the transmission network and distribution network and onwards to the end user. Now, clean energy is produced by a multitude of small producers in a variety of locations who may choose to use, sell or store the energy that they generate. We are adapting our business to reflect this new reality. A key first step is accurate forecasting that enables us to balance production and consumption efficiently by predicting future energy requirements and the market price of electricity. Even the best forecasts will not be accurate all of the time, however, so we also need to optimise our behaviours and predictions as reality unfolds. AI, machine learning and reinforcement learning are extremely valuable tools in this process.

Data helps customers to make informed choices. For example, a customer wanting to install solar panels or batteries needs to understand the financial implications of their decision in terms of installation cost and potential return on investment. They also need to understand that a new solar installation or insulation could reduce their CO2 emissions, but that it will take time to offset the CO2 that was emitted during the construction and transportation of the panels or the insulating materials. Data allows us to perform financial, environmental and life-cycle evaluations of the different solutions we propose and support our customers in their decisions. We are also conscious that asking an AI system to run these calculations is not carbon-neutral; we are studying this impact.

We depend on our customers to provide us with a lot of data about their energy production, energy consumption, energy storage capacity and stored energy volumes. To help them to optimise their energy consumption and their carbon footprint, we need to understand how and why they use energy. If they are increasing energy consumption to increase production in a factory, it is important for them to understand whether the costs would outweigh the additional profits they earn. Businesses have to make commercial decisions and balance them against carbon emissions, water consumption and other environmental impacts of their activities. Data is also a source of private information. Even in a house, the fact that a boiler is running to make hot water for a shower immediately tells you something about the lives of the inhabitants. It is not surprising that people are not always happy to share their data.

For us, the problem of sovereignty relates to the fact that each actor owns their data, which limits our ability to use it to improve financial and environmental outcomes.

But the transition to a decarbonized energy supply need to increase data exchange and enhance networking within the increasing numbers of energy market players (Flexibility service providers, local energy community, aggregators, prosumers). And new emerging collaborative environment such as Data space is an effective way to foster and facilitate data sharing between the fragmented energy market and the diversity of DER, while ensuring data sovereignty to the owner of the data. In essence, data spaces can serve communities of organizations gathered around a common objective in a particular domain, such as energy flexibility services where several energy market players are requesting more flexibility from prosumer to better balance the overall electrical system and mitigate network congestions. Schneider Electric has already made the connectivity with these distributed energy resources for energy efficiency optimization and using same AI combined with data spaces environment then can now expose new flexibility services for energy market ecosystem as a sovereign AI.

Other energy use cases in the short term are also identified where data sovereignty is a challenge and where we need also to be prepared for regulations changes (DataAct, Digital Product Passport). For these regulatory use cases, Data space is also foreseen as a game changer to be compliant with regulation, avoiding complex and costly infrastructure and data monopolized by big tech players.

Data spaces appear as an option to let each data owner truly own their data and decide how they will be used and shared. We are interested in the potential of data spaces to provide secure, standardized and easy access to the data that we need from our partners in order to provide additional Trust and sovereignty functions that Europe and other countries need in order to foster collaboration between energy stakeholders to make the electrical network more flexible and sustainable. We are already engaging in data space initiatives such as Data4Inudstry-X project with Valeo, Dawex and others to demonstrate regulatory use cases related to manufacturing domain and we are active in the energy data-X project, performing some energy flexibility scenario with the four German TSOs on how data can be exchanged to improve balancing services using behind the meter flexibility.

Colleagues in Schneider Electric tell me that they used to think that the problem was data interoperability, but that increasingly they see that the challenge lies in creating a functioning data economy that supports both the business and our environmental performance.

# Session 2

## Discussion

### From the floor

Beyond basic metrics like accuracy and precision, recent studies show that generic AI models like ChatGPT and Claude often provide wrong answers when they are asked to give their sources. This is a typical example of hallucination, when the model provides the most probable answer rather than the most accurate answer. What metrics would you use to determine that an AI is competitive?

### Anastasia Stasenko | Co-founder, Pleias

The first aspect is scalability, specifically the amount of energy that is required to deploy the model. At present, we use very powerful models even when we are only asking very simple questions. We are not yet in a production mode that would demonstrate that this truly is the new industrial revolution. The scalability and efficiency of models and the AI systems based on them is a very important criteria.

The second aspect relates to metrics, which relate to the quality of the outputs. First, we must reconsider how we how we work with evaluation in the world of AI. These systems are not built using conventional software frameworks, but it is still possible to inject a degree of verifiability. They are trained on reasoning data and with reinforcement learning algorithms, which creates opportunities to achieve variability. Secondly, in terms of evaluation, we do not yet have good benchmarks for what it would mean for a model to be good in a particular industry, like telecoms. There are groups working on this. It is not a trivial question: for a model to be successfully deployed in the telecoms industry, it would need telecoms use cases that do not currently exist and training on industry-specific data formats, data ontology, requirements and so on. The same applies to other industries, like nuclear.

Industry benchmarks would make it possible to evaluate off-the-shelf models and create use-case-specific evaluations, albeit with the involvement of significant expertise. Companies like Scale AI can demand huge fees because they have the human expertise that AI labs need to evaluate models. Companies with advanced AI strategies, such as JPMorgan, adopt a ‘human in the loop’ approach but this is not sufficient. As well as having a human in the loop, it is necessary to have a scalable way of creating evaluation and metrics. Synthetic data has a role to play here.

What worked in previous machine-learning environments is no longer sufficient. For truly competitive AI at the industry or enterprise level, it is necessary to have standardised approaches that define what AI evaluation is, how it is created, and how it can be put into production.

### From the floor

With regard to nuclear data and data availability, what can we do to find data and develop synthetic data, and how can we go further with it?

### Anastasia Stasenko | Co-founder, Pleias

We currently envision having an initial ‘seed’ corpus, consisting of diverse original data that adequately represents the reality of the use case. In nuclear, for example, this would include publicly available standards and documentation about internal processes. The seed corpus can be relatively small; we will soon publish a technical report on a recent corpus based on 50,000 Wikipedia articles. We then create exercises by creating queries about the seed corpus and back-translating them to create even more exercises. Rather than generating the synthetic data from nothing or a single chunk, the aim is to create a lot of exercises and different examples of how this data could be used. For example, we simulate RAG, multi-churn, completely adversarial tasks, and so on. This makes it possible to preserve the diversity in the synthetic data so to avoid model collapse.

Two points must be highlighted. First, it is vital to be data aware and have a data strategy and data governance for your use cases. Second, it is important to have methodology for synthetic data generation that is scalable, applicable to different data types and trusted. You need to do all this and then you need to have some compute or, in the case of synthetic data, lots of compute.

### Laurent Lafaye | Co-CEO, Dawex

Dawex has explored whether, in the US, the data marketplace and data spaces could be the market. Most discussions about data are actually discussions about US data, because it is the source of most data-processing initiatives. Ten years ago, there were a lot of personal data brokers in the US, as well as data aggregators who handled financial or alternative data for the stock market and hedge funds. At this time, the trusted third party was big tech.

In Europe, in the absence of a big trusted third party, we are obliged to create data spaces or trusted third-party territories where we can exchange data in different verticals, private data from industries, and data from companies with trusted third-party status. This is why the European Commission pushed for the creation of the Data Governance Act creating the framework for Data Intermediary Service Providers which act as a trusted third party. It is possible that we will see similar data spaces emerge outside Europe in the coming years. Last year, China decided to launch 100 data-spaces initiatives. Europe is promoting the concept of data spaces in other parts of the world like Japan, South Korea, Australia and even the manufacturing sector in the US, because by creating technical, regulatory and economic frameworks, we are making it possible to build this kind of trusted data ecosystem.

### Eric Brousseau | Governance and Regulation Chair, Paris Dauphine-PSL University

Thank you to all of our speakers for sharing their experience and their expertise. We will organise other events on data sharing, AI and related fields. Thank you to our audience for their support.

Chaire Gouvernance et Régulation
Fondation Paris-Dauphine
Place du Maréchal de Lattre de Tassigny - 75016 Paris (France)
https://chairgovreg.fondation-dauphine.fr/