TechSonar

2022-2023 REPORT

EDPS

European Data Protection Supervisor

Title: TechSonar reports

Editors: Massimo Attoresi, Stefano Leucci.
Authors: Marco Anders, Christian Ivanov, Robert Riemann, Xabier Lareo, Stefano Leucci

Photo credits:
Pag. 7 and 8 - https://unsplash.com/photos/9vdDUUuLi5g
Pag. 7 and 11 - https://unsplash.com/photos/4lSz1Jv0Vkc
Pag. 7 and 13 - https://unsplash.com/photos/LqKhnDzSF-8
Pag. 7 and 15 - https://unsplash.com/photos/-tikpxRBcsA
Pag. 7 and 17 - https://unsplash.com/photos/ix9KDnqJNjA

PDF QT-03-22-131-EN-N ISBN 978-92-9242-710-8 doi:10.2804/06518

TechSonar, in an increased complex world

By Wojciech Wiewiórowski

It was at the end September 2021 when the EDPS first launched its first foresight-related project, TechSonar. At that moment, we already believed the world to be a highly-complex place. Yet from September 2021 onwards, we stood by and watched in disbelief the events that unfolded at the eastern border of the European Union, which threatened the very core of our European democracy and values. Today, one thing is certain: new challenges will not cease to emerge.

This led the EDPS to a major question: based on the events of the last two years, what lessons can we learn about technological developments within our community of data protection regulators?

In our day-to-day work, we are confronted with challenges that arise due to continuous innovation. In the growing complexity of our digital society, we need to consider an additional factor: **the speed of change**.

Someone hinted at this decades ago. In 1965, Gordon Moore hypothesised that the number of transistors in microprocessors would double every 12 months or so. Moore was wrong. The growth has often been faster and more unpredictable than the rate he anticipated. The world has increased in its complexity. From this we can learn a first lesson: we cannot predict the pace of technology evolution, but **we can prepare for a diverse set of plausible scenarios**.

If we consider this impressive evolution - from Fintech to the metaverse, from artificial intelligence to biometrics - we can see **how complex it has become to find effective ways to intervene in these processes and take timely actions**.

The more we advance in our work, the more we are convinced that we – as data protection authorities, but also as data controllers and processors – need new tools and skills. **We will not be able to carry out our mission to supervise and regulate the use of technologies effectively without being able to anticipate and guide their evolution**.

To do this, foresight methodologies are key. The more we interact with foresight experts, the more we see an urgent need to integrate this domain with data protection.

For this reason, over the past year and since **the launch of TechSonar**, we have continued to develop our anticipatory mindset.

We co-organised the **closing conference of the Panelfit project** on 31 March 2022, we joined a panel dedicated to anticipatory techniques at the **Computer Privacy and Data Protection conference** on 24 May 2022, and we discussed our preliminary achievements at the **30thEuropean Conference of Data Protection Authorities in Dubrovnik on 20 May 2022**. Foresight was also one of the main themes of the EDPS **conference on "The Future of Data Protection: Effective Enforcement in the Digital World"** that we organised on 16 & 17 June 2022, in Brussels and online.

Continuing on this path, today we are publishing the outcome of our second TechSonar edition, with an updated set of technologies that we consider to be of primary importance to increase the preparedness of stakeholders in the field of personal data protection.

This new release of TechSonar has been enriched with a proof-of-concept analysis tool, created together with the publicly accessible **Competence Center on Text Mining and Analysis of the European Commission's Joint Research Center**.

The tool supports the information collection process, analysing the most important academic papers, as well as patents and projects funded by the European Union that concern the technologies we selected.

TechSonar is just a first step towards a wider, forward-looking perspective on our future. We are convinced that an effective approach to data protection regulation needs to take into account anticipatory and proactive ways to tackle its supervisory and advisory tasks, and to support the value-creation process of privacy enhancing technologies. Let me reiterate what was said during the **Panelfit conference**: we need to start considering the use of anticipatory and foresight techniques as a "new normal" in our future data protection efforts.

The EDPS firmly believes that a multi-stakeholder conversation that anticipates risks and damages to our future digital world is one of the most effective way to enforce the fundamental rights to privacy and data protection.

Data Protection Technology Sonar: augmenting machine intelligence through human collaboration

By Massimo Attoresi and Stefano Leucci

2.1 One year of foresight at the EDPS (and beyond)

When we decided to kickstart the TechSonar project in 2020, we noticed that the interest in foresight practice was scarcely present in the data protection domain. However, after only one year from the launch of the first TechSonar iteration, we were confronted with a different setting. At the conference **"Effective enforcement in the digital world"**, held in June 2022 in Brussels, we realized that TechSonar, together with other pilot projects in the data **protection field**, contributed to put foresight in the spotlight, and caught the attention of the data protection community.

The reason behind this sudden interest can be explained by the difference between foresight and the practices that have been traditionally used in data protection until now. While the main approach in data protection has been to try to anticipate the impact of technology on the rights and freedoms of natural persons, and in particular their right to data protection, foresight provides a wider approach that does not focus on one single aspect – i.e. the impact on data protection rights – but ensures a more comprehensive analysis of possible future scenarios.

The forward-looking approach that the EDPS has decided to adopt is in accordance with Regulation (EU) 1725/2018. In line with it, the EDPS should "monitor relevant developments, insofar as they have an impact on the protection of personal data, in particular the development of information and communication technologies". That is why foresight for the EDPS is not simply an operational approach. Rather, it is a strategic direction and, as such, it has been included as the first pillar of the **EDPS Strategy 2020-2024**.

As a supervisory authority, independence is a key aspect of every activity we carry out. This applies also to our exploration and execution of foresight. We are convinced that our forward-looking approach will help us to promote a more through and fruitful engagement within the data protection community of practitioners and experts while remaining true to our requirement of independence and transparency. We have already received proof of this from the growing number of requests to be kept up-to-date on our foresight work by many organisations and entities that routinely work with data protection.

This improved collaboration also manifested inside our organisation. Driven by the need to develop the necessary monitoring and assessment skills that are necessary to the performance of TechSonar, the officers involved in the project begun developing specific lines of expertise that could benefit the EDPS in a wide range of activities. In fact, it is becoming part of a wider effort to intervene in various fields beyond the technology realm, such as in relation to our policy or supervisory role.

2.2 Data Protection Technology Sonar: improvements during the second execution

For the second edition of TechSonar, we improved our "Data Protection Technology Sonar" **(presented last year)** by focusing on:

- improving the way we gather information on emerging technologies;
- improving the diversity of the team and the agility of execution.

To improve the way we gather information on emerging technologies, we decided to add a data analysis layer that supports the analysis of our experts.

Thanks to a fruitful collaboration with the **Competence Center on Text Mining and Analysis of the European Commission's Joint Research Center**, its JRC Tim Analytics Team supported us in the creation of a set of Dashboards for each trend. For the second iteration of TechSonar, the Tim Analytics tool analysed:

- 53 million of peer-reviewed scientific publications from **Scopus**;
- 27 million worldwide patent applications from **Patstat**;
- 87 thousand projects funded by the EU's framework programs for research and innovation (FP5 to Horizon 2020) from **Cordis**.

The information was extracted using semantic proximity techniques [1] and were presented in graphic formats to support the analysis by the EDPS team.

In particular, the following visualizations were used:

- geographical heatmap – revealing the countries in the world where a specific technology is mostly developed;
- organization graph – revealing the relationship between organizations in the world that deal with a specific technology;
- top 10 cited publications – revealing the most important publications to consider - through a dedicated application of the h-index [2] ;
- top 10 EU founded projects – revealing specific technologies or applications of it that will come in the future;
- triadic patents – revealing commercially-appealing technologies patented within the legislations of the European Union, United States and Japan.

Despite the additions describe above, the core methodology has not changed compared to the first edition of TechSonar.

The first phase (initial scouting) consists of a monitoring activity carried out by the Trend Coordinator with the goal of detecting a series of data from the wider technological landscape.

1. For more information about this technique see **https://en.wikipedia.org/wiki/Semantic_similarity**
2. For more information about h-index see **https://en.wikipedia.org/wiki/H-index**

2. METHODS

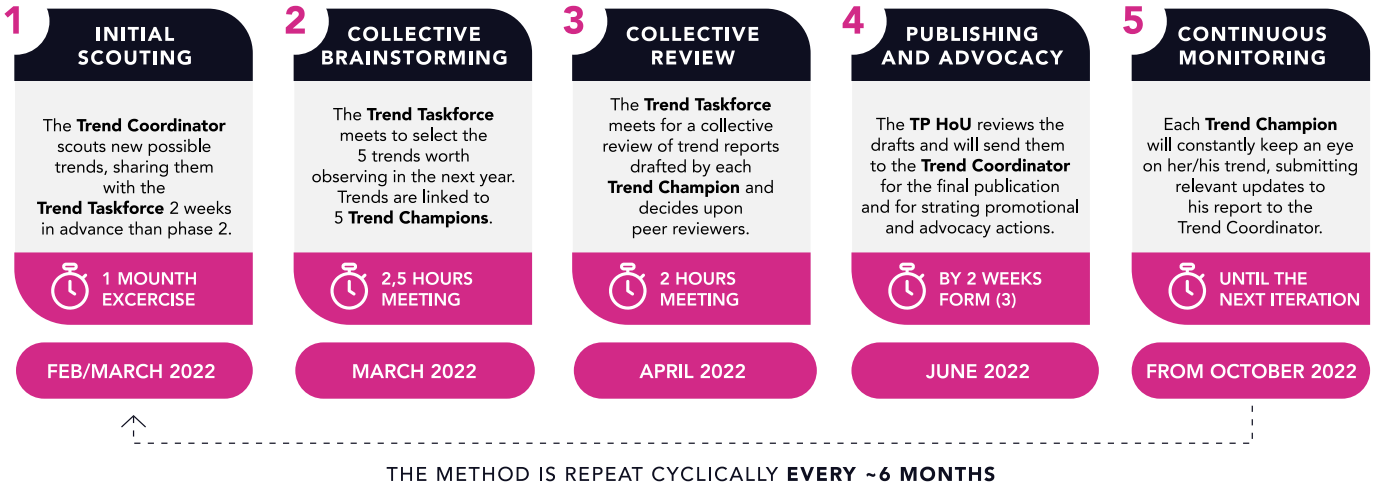

Figure 1 - Data Protection Technology Sonar methodological steps

The outputs provided by Tim Analytics were analysed together with a series of other supporting sources (newsletters, websites, trend reports, market analysis). Then, the Trend Coordinator short-listed the 15 technologies to be assessed in subsequent phases. The shortlisted technologies were then assessed on the basis of two indicators:

- the Privacy Risk Ratio, that helped the Team to understand the risk level of the short-listed technologies. It consists of a qualitative ratio and a quantitative ratio, and is grounded on the EDPB and EDPS guidelines;
- the Compounded Growth Rate, that was useful to understand the growth rate of each selected technology in the world market. It is a quantitative ratio, based on open data available in the web.

During the second phase (collective brainstorming) the Trend Coordinator presented the result of the first assessment (as shown in the Figure 2) to the Trend Taskforce, that now consists of technology, legal and policy experts across various units of the EDPS and of dedicated Trend Correspondents to help the team remain connected and to coordinate the work. This change to the composition of the task force has allowed us to increase the diversity of ideas and reduce internal bias.

The group firstly discussed and agreed on plausible future scenarios aiming at understanding driving forces of change, weak signals and their interactions[3]. According to the results, the group selects five technologies considered most impactful for the nearest future. These correspond to the technologies circled in figure 2.

The open consultation approach adopted during this phase is necessary to bring further value to the assessment of the previous phase, and to extract further meaning especially from the quantitative elements. Indeed, we are convinced to keep our methodology strongly grounded on human intelligence, and only used the data layer to provide for a more solid foundation upon which qualitative analysis can be built.

3. We define "forces of change" as trends that exists in the present and that will provide projections of the present into a future scenario. Moreover, a weak signal is an existing thing or phenomenon - limited in time and space - that can be interpreted as an indicator of potential greater change.

2. METHODS

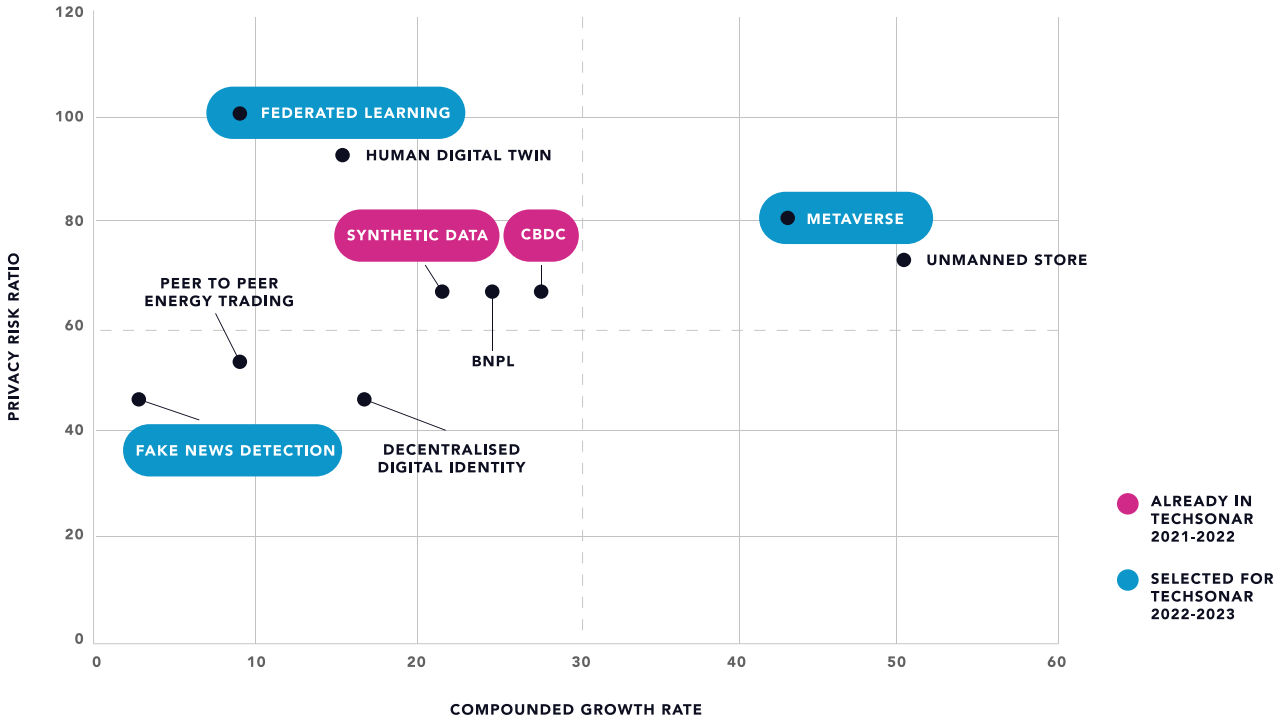

Figure 2 - The different technologies selected by the Trend Correspondent are arranged on the two axes for an easier analysis of the Trend Taskforce

Figure 2 shows an interesting outcome. Two technologies selected during the second execution of TechSonar were already selected during the first iteration. We consider this as an element of confirmation of the validity of the methodology.

Finally, the group assigned the five selected technologies to individual experts, the Tech Champions, who have technical expertise and who will follow the developments of the assigned technology. Before the publication of the final outcome of TechSonar, TechChampions will produce a brief dedicated report.

The agility of the exercise was another major improvement that we reached. In fact, the length of the reports was reduced. Positive and negative impacts were limited to a maximum of three bullet points, while a last section dedicated to EDPS related work was added.

In the third phase (collective review) each Tech Champion has been paired with another technology expert in order to improve the outcomes and avoid problems and bias that might arise in the process due to the inherent execution speed.

During the fourth phase, the Trend Coordinator performed a review of the contents, published the output on the EDPS website and launched a series of internal and external promotional and advocacy activities. The Trend Correspondents were also involved in this phase, providing the

TechChampions with questions and reviews. In the last phase (continuous monitoring) each Tech Champion will continue to monitor the developments of the assigned technologies and notifies any relevant updates. This way, a reference staff is appointed as a contact point and a centre as expertise for both internal and external stakeholders on each selected technology.

Tim Analytics will also be of huge help in this phase. In fact, it will be frequently updated and will allow the Tech Champion to monitor the evolution of technologies under analysis.

2.3 Conclusion

This section of the TechSonar Report 2022-2023 aimed at improving the transparency of the process through which we selected emerging technologies.

In the upcoming months, the team will further develop internal continuous monitoring processes.

The goal in the years to come is to build a linear and consistent process that goes from the identification of technology trends to the development and management of structured internal knowledge. The final outcome is to provide quality and timely background to inform different actions of the EDPS.

We trust that our efforts will be reused and improved by other stakeholders in the field and that we will be able to benefit from future co-operations and synergies.

Selected technologies for 2022-2023

Fake news detection systems
pag.9

Central bank digital currency
pag.18

Metaverse
pag.12

Federated learning
pag.16

Synthetic data
pag.14

Fake news detection systems

By Marco Anders

EXPLORE OUR DASHBOARD ON FAKE NEWS DETECTION SYSTEMS

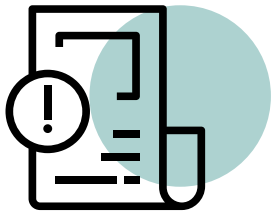

In recent years, the dissemination of fake news has been brought more and more into the spotlight as it has been massively used to disseminate political propaganda, influence the outcome of elections or harm a person or a group of people.

Highly sophisticated applications (bots) are organised in networks and massively spread to amplify fake news over social media in the form of text, images, audio or video files. Often, these bot nets happened to be organised by foreign state actors, trying to obscure the originator.

Fighting fake news is extremely challenging, as:

- in a democracy, freedom of speech is a fundamental right fostering media independence and pluralism; however, sometimes there is a very subtle line between separating unconventional personal views and claims of truth from fake news;
- fake news can be detected by checking consistency of the news with different domains, such as technical background to discover the real sender or social and/or judicial background (for example: what is the intention of the fake message, e.g. putting harm on a person or a group); therefore, fact-checking requires having awareness on different contexts and the availability of reliable sources;
- the sheer mass of fake news spread over social media cannot be handled manually.

Manual fact checking can address some of these challenges, for example when checking the consistency of news in different contexts. However, manual fact-checking is too slow to cover big information spreaders such as social media platforms. This is where automation comes into play.

Automated fact-checking tools often combine different methods, for example artificial intelligence, natural language processing (analysing the language used) and blockchain. As regards to fake news embedded in images and videos, the tools often combine metadata; social interactions; visual cues; the profile of the source; and other contextual information surrounding an imaze or video to increase accuracy.

Algorithms are trained to verify news content; detect amplification (excessive and/or targeted dissemination); spot fake accounts and detect campaigns. Often, the fake news analysis process applies several algorithms sequentially. However, effectiveness of these algorithms is yet to be improved.

3.1 REPORTS

Even if fake news is spread heavily on social media, research has found that human behaviour ("word of mouth" marketing) contributes more to the spread of fake news than automated bots do. This shows that fighting the fake news sender is not the only approach. It also makes sense to increase the resilience to fake news on the side of the recipient and our society. Therefore, another important pillar of fake news detection is to increase citizens' awareness and media literacy.

Positive foreseen impacts on data protection:

- **Awareness and media literacy will be raised at consumer level with an effect on data protection:** the European Union has already launched a number of projects to analyse the phenomenon of fake news and develop countermeasures. As a result, one pillar identified is to increase awareness and media literacy. Such awareness-raising initiatives may have a positive impact on data protection in general: media literate consumers are capable of reflecting on media messages and understand the power of information and communication. Therefore, these consumers will be more careful when disclosing their personal data thoughtlessly.
- **Effective fake news detection will reduce defamation of individuals:** A common practice to hide the source of the entity spreading fake news is to hijack other individuals' accounts. The owners of such accounts may be defamed, e.g. by the spread of fake news. At the same time, as it is common for fake news to be spread with the goal to harm individuals or groups of people, for example in political campaigns, technology for fake news detection would limit this kind of defamation.

Negative foreseen impacts on data protection:

- **Lack of transparency and need fora legal basis:** fake news detection algorithms combine different sets of information with each other among which there is also personal data (e.g. related to the source of the messages). Currently, it is not transparent to individuals what personal data is processed in the context of fake news detection, nor what the legal basis is for this processing. As a result, individuals cannot effectively exercise their rights to access, correction and deletion of their personal data.

- **Accuracy of the algorithms:** While technology can help to assess large numbers of fake news instances, its effectiveness is bound by the error rates of the applied algorithms (sometimes a set of different algorithms is applied sequentially). Given the contextual complexity, as well as cultural differences and the challenges of artificial intelligence, fake news detection may lead to biased results. This could lead to true information being blocked or categories of users/ opinions that marginalised.
- **Increase of automated decision-making:** Fake news detection technology consists mainly of automated detection tools for which effective human oversight should be applied. Often, human resources devoted to oversight are not sufficient and data subjects may not be able to exercise their rights for human oversight and/or access to their personal data.

Our three picks of suggested readings:

- C. Wardle, H. Derakhshan, **Information disorder: toward an interdisciplinary for research and policy making**, Council of Europe report DGI (2019)09, September 2017.
- European Parliamentary Research Service, **Automated tackling of disinformation**, 2019.
- The social observatory for disinformation and media analysis (SOMA), **https://www.disinfobservatory.org/**

EDPS related works:

- Opinion 3/2018 on **online manipulation and personal data**, March 2018.

Metaverse

By Christian Ivanov

EXPLORE OUR DASHBOARD ON METAVERSE

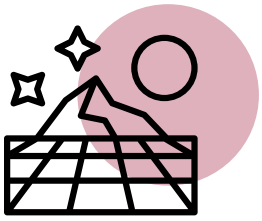

In 1992, Neal Stephenson defined the metaverse as a network of 3D virtual worlds focused on social connection. Today, the metaverse is a concept, which aims to determine the general digital transformation in every aspect of our physical lives.

The concept captures a mix of virtual reality and other technologies. It is a world of interconnected physical and virtual communities where users can develop professionally socialise, entertain, commerce and even trade with virtual properties. All of this is accessible from any place in the world, using different types of smart devices, such as virtual reality (VR) headsets, smart bracelets, smartphone apps, etc. Metaverse platforms are the next evolution of connectivity where the features and functionality of each individual application and device are coming together in one place.

There is no unified view about how exactly the metaverse platforms look like and how they will continue to grow. **Major companies** are already developing their own versions and each of them is adapting the idea of metaverse to their strengths. Examples are the leading **social media platforms, gaming companies, online shops etc**. Some brands are also claiming **virtual real estate** in the metaverse platforms with their own digital properties.

One can perceive metaverse as a digital version of our reality (representing cities, buildings, streets, individuals etc).

At the same time, the idea has also grown to building elements that do not exist in reality, such as virtual events and digital venues. Notwithstanding the way this will develop further, the metaverse aims to steer the direction of the world's evolution as it might completely transform the way individuals, communities, governments and corporations interact. Considering the vast quantities of personal data that may be collected on participating individuals, the metaverse platforms pose significant privacy-related challenges.

Positive foreseen impacts on data protection:

- **Difficulties in demonstrating positive impacts:** Since the metaverse is a concept under development and the design and configuration of the technology is still not specified, at this stage direct positive impacts cannot be demonstrated with concrete figures. Eventual privacy enhancing features could be implemented in the metaverse to obtain an enhanced level of privacy.
- **Anonymity in some processes:** Depending on a particular case, certain metaverse platforms could allow individuals to create avatars with entirely fictional characters that do not resemble the physical appearance or include any related information with the real person; or to create any other elements and objects relating to them having features different from the corresponding objects in reality, insofar as this might be considered fair and without negative implications for others. This could be used to enhance anonymity towards the other users/vendors within the entire interacting process in the platform.

Negative foreseen impacts on data protection:

- **Deeper profiling:** Profiling hides risks in each social media platform. However, compared to traditional social media, metaverse platforms can collect, store, and rely on more personal data than ever before in order to examine users' behaviour. This gives possibilities to the metaverse providers to classify people in precise profiles, even considering new categories of data.
- **Constant monitoring:** The metaverse makes technologies closer to every aspect of our physical lives, leading to constant observation of every aspects of it. This is connected to a constant privacy invasion, which becomes the normality. The usage of expanded amount of devices allows tracking through multiple channels, like wearable devices, motion sensors, microphones, heart and respiratory monitors, etc. This allows the surveillance of users' interactions to an even higher extent than traditional applications.
- **Interference of special categories of data:** The metaverse platforms allows for the monitoring of special categories of personal data like physiological responses, emotions and biometric data, such as a person's gait; facial expressions; eye movements; vocal inflections; and vital signs in real time. Considering the direct statements and actions in the platform (ex. visiting a special place in the platform), other special categories of data can also be easily reviewed, such as political beliefs, sexual orientation etc. Processing sensitive data and targeting users based on them creates high risks for the fundamental rights and freedoms of individuals.

Our three picks of suggested readings:

- M. O'Brian, K. Chan, **EXPLAINER: What is the metaverse and how will it work?**, 2021.
- L. Lee, T. Braud, P. Zhou, L. Wang, D. Xu, Z. Lin, A. Kumar, C. Bermejo, P. Hui, **All One Needs to Know about Metaverse: A Complete Survey on Technological Singularity, Virtual Ecosystem, and Research Agenda**, 2021.
- Y. Wang, Z. Su, N. Zhang, D. Liu, R. Xing, T. Luan, X. Shen, **A Survey on Metaverse: Fundamentals, Security, and Privacy**, 2022.

EDPS related works:

- Technology Report no. 1, **Smart glasses and data protection**, January 2019
- TechDispatch #1/2021, **Facial emotion recognition**, May 2021.

Synthetic data

By Robert Riemann

EXPLORE OUR DASHBOARD ON SYNTHETIC DATA

Synthetic data is artificial data that is generated from original data and a model that is trained to reproduce the characteristics and structure of the original data. This means that synthetic data and original data should deliver very similar results when undergoing the same statistical analysis. The degree to which synthetic data is an accurate proxy for the original data is a measure of the *utility* of the method and the model.

The generation process, also called synthesis, can be performed using different techniques, such as decision trees, or deep learning algorithms. Synthetic data can be classified with respect to the type of the original data: the first type employs real datasets, the second employs knowledge gathered by the analysts instead, and the third type is a combination of these two. Generative Adversarial Networks (GANs) were introduced recently and are commonly used in the field of image recognition. They are generally composed of two neural networks training each other iteratively. The generator network produces synthetic images that the discriminator network tries to identify as such in comparison to real images.

A privacy assurance assessment should be performed to ensure that the resulting synthetic data is not actual personal data. This privacy assurance evaluates the extent to which data subjects can be identified in the synthetic data and how much new data about those data subjects would be revealed upon successful identification.

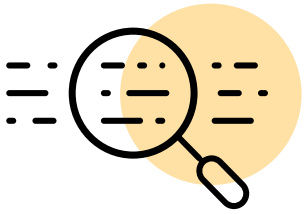

Synthetic data is gaining traction within the machine learning domain. It helps training machine learning algorithms that need an immense amount of labeled training data, which can be costly or come with data usage restrictions. Moreover, manufacturers can use synthetic data for software testing and quality assurance. Synthetic data can help companies and researchers build data repositories needed to train and even pre-train machine learning models, a technique referred to as transfer learning.

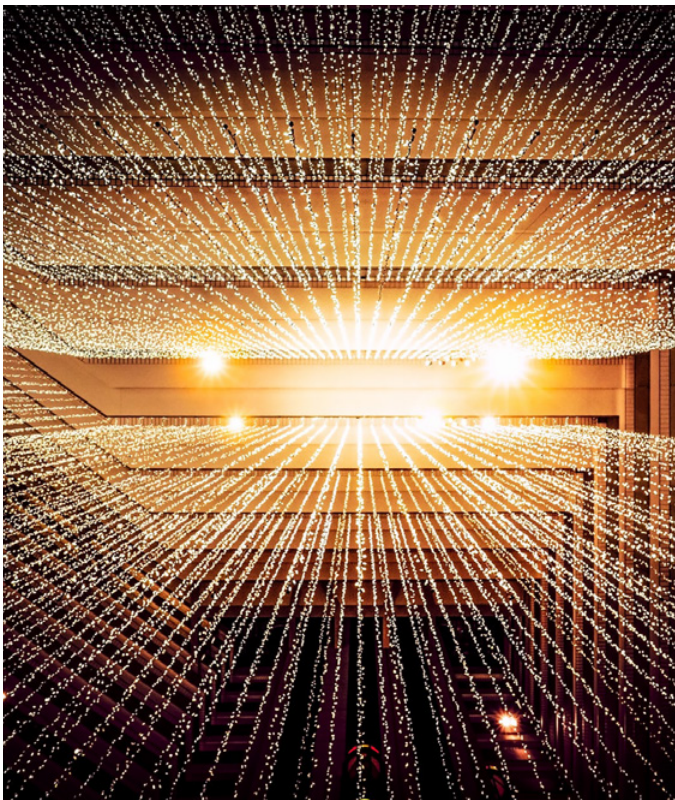

Positive foreseen impacts on data protection:

- **Enhancing privacy in technologies:** from a data protection by design approach, this technology could provide, upon a privacy assurance assessment, an added value for the privacy of individuals, whose personal data does not have to be disclosed.
- **Improved fairness:** synthetic data might contribute to mitigate bias by using fair synthetic datasets to train artificial intelligence models. These datasets are manipulated to have a better representativeness of the world (to be less as it is, and more as society would like it to be). For instance, without gender-based or racial discrimination.

Negative foreseen impacts on data protection:

- **Output control could be complex:** especially in complex datasets, the best way to ensure the output is accurate and consistent is by comparing synthetic data with original data, or human-annotated data. However, for this comparison again access to the original data is required.
- **Difficulty to map outliers:** synthetic data can only mimic real-world data; it is not a replica. Therefore, synthetic data may not cover some outliers that original data has. However, outliers in the data can be more important than regular data points for some applications.
- **Quality of the model depends on the data source:** the quality of synthetic data is highly correlated with the quality of the original data and the data generation model. Synthetic data may reflect the biases in original data. Also, the manipulation of datasets to create fair synthetic datasets might result in inaccurate data.

Our three picks of suggested readings:

- T. E. Raghunathan, Synthetic data, Annual Review of Statistics and Its Application, 8, 129-140, 2021.
- K. Dankar, I. Mahmoud. Fake it till you make it: guidelines for effective synthetic data generation, Applied Sciences 11.5 (2021): 2158, 2021.
- J. Hradec, M. Craglia, M. Di Leo, S. De Nigris, N. Ostlaender, N. Nicholson, **Multipurpose synthetic population for policy applications**, EUR 31116 EN, Publications Office of the European Union, Luxembourg, ISBN 978-92-76-53478-5 (online), doi:10.2760/50072 (online), JRC128595, 2022.

EDPS related works:

- **TechSonar Report 2021-2022**, December 2021.
- Internet Privacy Engineering Network (IPEN), **Synthetic data webinar**, June 2021.

Federated learning

By Xabier Lareo

EXPLORE OUR DASHBOARD ON FEDERATE LEARNING

Training, testing and validating machine-learning models require data. Data that sometimes is dispersed amongst many, even millions of, parties (devices). Federated learning is a relatively new way of developing machine-learning models where each federated device shares its local model parameters instead of sharing the whole dataset used to train it. The federated learning topology defines the way parameters are shared. In a centralised topology, the parties send their model parameters to a central server that uses them to train a central model which in turn sends back updated parameters to the parties. In other topologies, such as the peer-to-peer or hierarchical one, the parties share their parameters with a subset of their peers. Federated learning is a potential solution for developing machine-learning models that require huge or very disperse datasets. However, it is not a one-size-fits-all machine learning scenarios.

Federated learning still has open issues that scientists and engineers work hard to solve, some of which are detailed below.

- **Communication efficiency:** federated learning involves numerous data transfers. Consequently, the central server or parties receiving the parameters need to be resilient to communication failures and delays. Ensuring efficient communication and synchronisation amongst the federated devices remains a relevant issue.
- **Device heterogeneity:** computing capacities of the federated parties are often heterogeneous and sometimes unknown to the other parties or central server. It is still difficult to ensure the training tasks will work within a heterogeneous set of devices.
- **Data heterogeneity:** federated parties' datasets can be very heterogeneous in terms of data quantity, quality and diversity. It is difficult to measure beforehand the statistical heterogeneity of the training datasets and to mitigate the potential negative impacts such heterogeneity might have.
- **Privacy:** there is a need for efficient implementation of privacy enhancing technologies to avoid information leakages from shared model parameters.

Positive foreseen impacts on data protection:

- **Decentralisation:** by leveraging on distributed datasets, federated learning avoids data centralisation and allows the parties to have better control over the processing of their personal data.
- **Data minimisation:** federated learning reduces the amount of personal data transferred and processed by third parties for machine-learning model training.
- **International cooperation:** when the shared parameters are anonymous, federated learning facilitates the training of models with data coming from different jurisdictions.

Negative foreseen impacts on data protection:

- **Interpretability:** machine-learning developers often rely on the analysis of the training dataset to interpret the model behaviour. The developers using federated learning do not have access to the full training dataset, which can reduce the models' interpretability.
- **Fairness:** some federated learning settings may facilitate bias toward some parties, for example towards devices hosting the most common model types.
- **Security issues:** the distributed nature of federated learning facilitates some types of attacks (e.g. model poisoning). Classic defence mechanisms do not currently provide sufficient protection in a federated learning setup. Ad hoc defence methods still have to be developed and tested.

Our three picks of suggested readings:

- L. Tian, A. Kumar Sahu, A. S. Talwalkar and V. Smith, **Federated Learning: Challenges, Methods, and Future Directions**, IEEE Signal Processing Magazine 37, 2020.
- Q. Li, W. Zeyi, H. Bingsheng, **A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection**, ArXiv abs/1907.09693, 2021.
- P. Kairouz et al, **Advances and Open Problems in Federated Learning**, Foundations and Trends in Machine Learning Vol 4 Issue 1, 2021.

Central bank digital currency

By Stefano Leucci

EXPLORE OUR DASHBOARD ON CENTRAL BANK DIGITAL CURRENCY

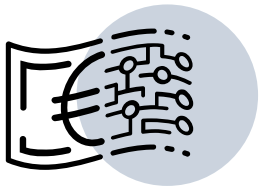

Central Bank Digital Currency (CBDC) is a new form of money that exists only in digital form. Instead of printing money, the central bank issues widely accessible digital coins so that digital transactions and transfers become simple.

Efforts towards CBDC grow all over the world for many reasons. First, the COVID-19 crisis induced a shift in payment habits towards **digital, contactless payments and e-commerce** due to a now refuted danger of **banknotes being way of transmitting infection**, which has accelerated the decline of cash use. Second, cryptocurrencies developed by private organisations or informal communities (e.g. Bitcoin) have seen significant developments and value gain. As a response, **87 countries** (representing over 90 percent of global GDP) are now exploring central bank digital currencies, while **9 of them have fully launched** a state-owned digital currency.

CBDC could be developed in a number of ways. In a centralised approach, transactions are recorded in ledgers managed by central banks that also provide user-facing services. In a decentralised approach, a central bank sets rules and requirements for the settlement of CBDC transactions that are then recorded by users and/or financial intermediaries.

The impact of CBDC depends also on the chosen implementation.

Conventional money requires many intermediaries in the payment chain, resulting in less efficient and secure payment experiences, **as we showed in our recent TechDispatch**. CBDC could find solutions to these issues, developing a more efficient, fast, secure and sovereign form of payment process.

The European Central Bank, after exploring possible **design scenarios** for launching a Digital Euro and **consulting** with stakeholders, decided to **launch a CBDC project** with an **investigation phase that will last from October 2021 to October 2023**.

Positive foreseen impacts on data protection:

- **Privacy is one of the most important design feature:** the **consultation launched by the ECB in October 2020** revealed that privacy is considered as the most important feature of a digital euro by both citizens and professionals; this was also confirmed by different **focus groups.** Design decisions are still open, and this situation results in a wide range of opportunities for configuring the product with an effective data protection by design approach.
- **More control over personal data and security:** assuming that the development of CBDC will follow a strict data-protection-by-design and by-default approach, a CBDC could increase data protection and security in digital payments and provide payers more control over their personal data.
- **Enhanced possibility to have anonymity in the payment process:** privacy-enhancing technologies could be used to enhance the way anonymity is wired within the entire payment process while allowing the auditing only in pre-determined lawful cases, such as preventing money laundering, counter terrorism financing and tax evasion.

Negative foreseen impacts on data protection:

- **Concentration of data in the hands of central banks could lead to increased privacy risks for citizens:** if payment data of all citizens were concentrated in the databases of a central bank, it would generate incentives for cyberattacks and a high systemic risk of individual or generalised surveillance in case of data breaches or, more in general, of unlawful access.
- **Wrong design choices might worsen data protection issues in digital payments:** payment data already reveals very sensitive aspects of a person. Wrong design choices in the underlying technological infrastructure might exacerbate the privacy and data protection issues that already exists in the digital payment landscape. For example, transactional data could be unlawfully used for credit evaluation and cross-selling initiatives.
- **Lack of security might turn into severe lack of trust from users:** security concerns in the CBDC infrastructure, whose security requirements and expectations are high, may turn into a significant loss of trust from users.

Our three picks of suggested readings:

- European Central Bank, **Report on a digital euro**, 2020
- Raskin et.al., **Digital currencies, decentralized ledgers, and the future of central banking,** National Bureau of Economic Research, 2016
- Institute and Faculty of Actuaries, **Understanding Central Bank Digital Currencies (CBDC),** March 2019

EDPS related works:

- TechDispatch #2/2021, Card-based payments, December 2021 - **https://edps.europa.eu/data-protection/our-work/publications/techdispatch/techdispatch-22021-card-based-payments_en**
- TechSonar Report 2021-2022, December 2021 - **https://edps.europa.eu/system/files/2021-12/techsonar_2021-2022_report_en.pdf**
- EDPB letter to the European institutions on the privacy and data protection aspects of a possible digital euro, June 2021 - **https://edpb.europa.eu/system/files/2021-07/edpb_letter_out_2021_0112-digitaleuro-toep_en.pdf**
- EDPB Statement 04/2022 on the design choices for a digital euro from the privacy and data protection perspective, 10th October 2022 - **https://edpb.europa.eu/system/files/2022-10/edpb_statement_20221010_digital_euro_en.pdf**

edps.europa.eu

@EU_EDPS

EDPS

European Data Protection Supervisor

EU Video

EU Voice

QT-03-22-131-EN-N

Publications Office
of the European Union

ISBN 978-92-9242-710-8
doi:10.2804/06518