State-of-the-Art Report

The potential that **Privacy Tech** brings for the German and European tech ecosystem and a special look at the growing field of computer vision

September 2021

KI BUNDESVERBAND

iconomy

Table of Contents

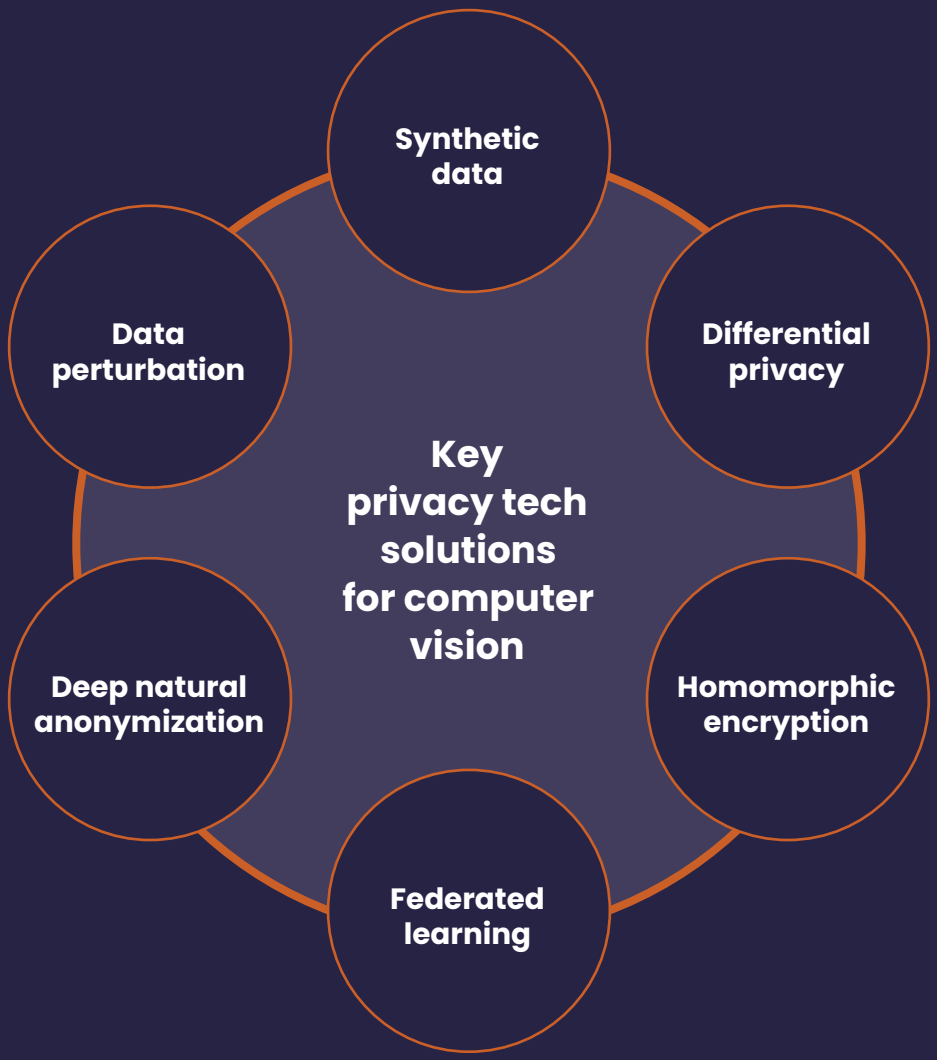

Foreword

After the EU's flagship privacy law - the General Data Protection Regulation - was enacted in 2018, it took a while until the public discourse shifted away from the burdens it imposed upon data-based products and services. Today, European privacy standards are increasingly regarded as a competitive advantage and a buzzing new tech vertical has emerged: privacy tech.

This is why the German AI Association, the Federal Office for Information Security (BSI) and iconomy teamed up in early 2020 to initiate an informal community that brings privacy tech leaders and supporters together on a regular basis. Since then, the community explores how the alleged conflict between high European privacy standards on the one hand and scalable data-based business models on the other can be - technologically - resolved and how they can support and develop the burgeoning new field of privacy tech.

For the first time, the community decided to produce a report in order to make best practices, narratives and knowhow accessible to a broader audience. We set out to shed light on the current state of privacy tech solutions and took a closer look at a unique case study at the center of the public discourse on privacy: the use of AI-powered computer vision. To create this paper, we interviewed community members, privacy lawyers and technologists to identify and examine the most promising privacy tech solutions that allow companies to comply with the GDPR and utilize computer vision (one of the fastest growing fields in artificial intelligence) simultaneously.

The findings are impressive, and we would like to thank everyone that contributed to our work, especially Sina Youn and the Privacy Working Group she heads at the AI Association, and we encourage those active in the field to get involved. We are convinced that collaboration and public support for one another is necessary to position privacy tech's technological and societal value, help relevant stakeholders understand the benefits and make best practices known to a broader audience.

But first, we wish you a pleasant read.

Vanessa Cann
– Managing Director,
German AI Association

Felix Styma
– Managing Partner,
iconomy

1. The (rocky) birth of Privacy Tech

In 2018, the European Union (EU) enacted its foundational General Data Protection Regulation (GDPR) and became an international leader in privacy.

Although privacy regulations have existed for a long time and affected companies, governments and individuals, the GDPR set a new, higher standard around the world. It introduced big fines that pushed companies to quickly implement strict privacy measures and the market principle, which took GDPR worldwide by including all companies that offer products and services in the EU in its jurisdiction.[1] Because of this, GDPR set a new higher bar for citizen and consumer trust. Since the introduction of GDPR, more than 60 jurisdictions around the world[2] have enacted or proposed privacy and data protection laws, and by 2023, Gartner predicts that 65% of the world's population will have their personal information covered under modern privacy regulations.[3]

However, as many know, the story of GDPR is not perfect. When it first entered into force across all European markets, most actors - including businesses, governments, public authorities and startups - were poorly prepared. They faced legal uncertainty and new terrain. To put it into context: even within a leading community of over 400,000 information security professionals (the Information Security Community), 60 percent of its members missed the GDPR's May 25th deadline, despite 80 percent of them calling GDPR a "top priority for their organization."[4] There was a lack of necessary expertise and budget to radically improve their security practices. So, consumer-facing companies and organizations scrambled. They started collecting user's consent and gradually adapted their business models to comply with GDPR.

In the meantime, a new set of GDPR-related technological approaches emerged. Its objective was to tackle various aspects of compliance and implementation, enable actors to maintain the ability to store and process data to the extent necessary for modern digital products and services, and allow companies to comply with GDPR - or other upcoming data privacy regulations - in a user-centric way. Thus, a novel vertical in Europe's tech ecosystem was born: privacy-preserving and privacy-enhancing technologies - "**privacy tech**".

What is 'Privacy Tech'?

Technologies that embody fundamental data protection principles and regulations in information systems, so that the privacy of personally identifiable information (PII) is preserved or enhanced without compromising system functionality or the value and quality of the data nor the accuracy of the model based on it.

This definition was developed by the 'privacy tech community' and may be regarded as an intermediary working definition.

[1] https://gdpr-info.eu/art-3-gdpr/
[2] This includes: Argentina, Australia, Brazil, Egypt, India, Indonesia, Japan, Kenya, Mexico, Nigeria, Panama, the U.S., Singapore and Thailand.
[3] https://www.gartner.com/smarterwithgartner/gartner-predicts-for-the-future-of-privacy-2020/
[4] https://www.businesswire.com/news/home/20180417005296/en/Most-Companies-Not-Prepared-for-GDPR-Data-Privacy-Compliance-Deadline-According-to-New-Study

1. The (rocky) birth of Privacy Tech

Privacy tech makes the insight and added-value from immense amounts of sensitive data possible by ensuring that no information can be related to an identified or identifiable natural person. It also helps companies and consumers address complex data privacy regulations in a user-centric way. One of the central aims of privacy tech is to not only enable data-based products and services to comply with GDPR, but shield companies and users from the complexity of data privacy regulation. It is a set of technologies and techniques that companies can use to comply with GDPR - or sector-specific data privacy regulations or upcoming regulations, like the EU's Data Governance Act - without compromising on data use and modern user experiences.

Privacy tech is a large and promising field that encompasses a wide range of services and data formats, including text data stored in the databases of common apps or websites and image data stored in smartphones or by companies. Yet, as this paper will outline, there is a unique subset of privacy tech that specifically utilizes visual data. This state of the art paper will introduce privacy tech in the context of visual data and provide a comparative analysis of the most important privacy tech solutions in the broader context of computer vision. The goal is to inform policymakers, authorities, businesses and emerging start-ups about the current state of privacy tech for visual data and inspire deeper discussion about where the German and European tech ecosystem ought to go.

"The GDPR is not the endpoint and data regulation will keep evolving. Privacy tech is playing a central role protecting rights and interests for consumers and the business."

Michael Bültmann
– Managing Director, HERE Technologies

2. Computer vision in the EU: Privacy Tech to the rescue?

Powered by deep learning and artificial neural networks, the field of computer vision has exponentially improved over the past five years.

Artificially intelligent systems can now "see" like humans, and decipher images to identify complex patterns. This technological leap is full of promise for our global economy.[5] The field of computer vision allows one to develop a new type of machines and technology; one that can now automate tasks requiring visual cognition.

Already, there are impressive outcomes: disease diagnosed from images of CT scans, autonomous vehicles along the road, train operators able to ensure safe social distancing in the global COVID pandemic.[6] Computer vision and video analytics are key innovations that kickstart a growing set of digital solutions and businesses worldwide. Yet, these AI techniques are arguably not living up to their potential in the European Union and Germany in particular. This is largely due to privacy rules, including GDPR. Computer vision requires a significant amount of image and video data[7] in order to train and optimize their models; data that is highly sensitive and for obvious reason: the tracking of human faces strikes a different chord and carries a severe risk--for individuals and societies writ large.

For example, in smart cities modern systems could manage and optimize traffic flow. However, these systems are often based on analytical insights from video and image data that can detect when cars enter cities -- the precise time and location, and where the cars are originally driving from. Autonomously driving vehicles rely on 10 cameras for their active driving assist systems, and use the video footage to "learn" how to reliably react to any object, person and event in their surroundings. Car manufacturers also use computer vision technology to test and train the autonomous vehicle decision-making systems, to ensure that the cars make the right decision at the right time. For all of these innovative systems and technologies based on image data to work and to even be developed within the EU, the GDPR and additional regulations have to be taken into account.

Privacy tech is poised to unlock the potential of computer vision and video analytics for European markets. It has the chance to resolve the supposed conflict of interest between high European privacy requirements on the one hand and scalable, innovative usage of computer and machine vision on the other. And it may even be the beginning of one of the European success stories, where an ambitious regulatory regime triggers the birth of a buzzing new tech vertical bringing together highest ethical standards and an enabling role for the development of scalable AI business on this foundation.

"Advanced image and video technology combined with other sensors will become pervasive very soon - it is one of the most powerful IT applications that will change our lives. All the more important to integrate privacy capabilities from the start."

Ansgar Baums
– Director Government Relations, Zoom Video Communications

[5] https://www.forbes.com/sites/cognitiveworld/2019/06/26/the-present-and-future-of-computer-vision/?sh=41c92fe3517d
[6] Although it should be noted that as the field of computer vision grows, it is also increasingly susceptible to adversarial attacks and manipulation.
[7] Throughout this paper, the authors will refer to video and image data as 'visual data' and 'video and image data' interchangeably.

3. The most important existing Privacy Tech Solutions

Although the following overview is not exhaustive, there are currently six buzzed-about privacy-preserving technologies that utilize visual data.[8]

There are: synthetic data, differential privacy, federated learning, deep natural anonymization, data perturbation (or 'data poisoning' and randomized smoothing), and homomorphic encryption. Each of these technologies are well regarded by privacy experts and currently being deployed in projects by some of the biggest governments and companies in the world.

Figure one below describes each privacy tech, how it works, its specific benefits--in addition to privacy, its unique challenges and highlights a few selected use case scenarios. It is worth noting that each privacy technology is often not used in isolation. Rather, it is frequently used in combination with other privacy-preserving technologies or techniques, which we outline in two longer use cases below.

Since each privacy tech is technical and can be difficult to comprehend, we have organized them into three clusters (though we do not claim perfection.) To start, two types of privacy tech can be conceived of as foundations of privacy: new 'fake' synthetic data (based on reality) and blurred data from differential privacy, with enough noise to lower the risk of data leaks. Both create robust forms of privacy since they, arguably, are closest to private from the onset.

Next, there are two types of privacy tech that, essentially, enable new methods of large-scale data sharing and storing: homomorphic encryption (HE) and federated learning (FL). Both HE and FL allow companies to scale and massively increase the size and diversity of their datasets while still keeping information confidential.

Finally, there are two privacy techniques which alter the data itself - not to the extent of synthetic data or differential privacy, but to the extent where all facial identifiers are covered and data can still be used for processing and analysis. These are Deep Natural Anonymization and Data Perturbation or 'Data Poisoning' and randomized smoothing.

"Developing products - developing projects - hand in hand with regulators and having their early stage buy-in is the most secure way to build privacy."

Dr. Felix Wittern
– Partner Technology, Outsourcing and Privacy, fieldfisher

[8] This list was put together with the help of the privacy tech community, including leading privacy experts, lawyers, start-ups, political stakeholders and industry representatives.

3. The most important existing Privacy Tech Solutions

Privacy Tech	How it works	Specific Benefits	Challenges	Use Case[9]
Synthetic Data (SD)	Synthetic data is data that computer simulations or algorithms generate as an alternative to or alteration of real-world data, partially to redact or minimize personal information in the data.[10] Created either: (a) from scratch based on real-live or fictional scenarios; (b) advanced data manipulation techniques to create more novel and diverse training examples.[11]	Increases diversity within insufficiently small datasets or limited datasets. Particularly useful for computer vision, which relies on visual data. SD can be shared publicly, which makes it easier for researchers or developers to accelerate and reproduce research.	Often, synthetic data does not look like real images, and developers need to invest more time and resources to make them look as photo realistic as possible, or developers must create additional technical means that help models transfer synthetic data sets to real test sets. Even synthetic or augmented data has its limits in terms of diversity because an artificially generated scenario will never fully represent reality. For augmented data and also fully synthetic data based on real-life reference, personal information might be processed.	Healthcare providers in fields such as medical imaging use synthetic data to train AI models while protecting patient privacy. For example, the startup Curai trained a diagnostic model on 400,000 simulated medical cases.[12]
Differential Privacy (DP)	Differential privacy is a field of data science used to mine user data while protecting individual user's privacy. It is a mathematical method of sharing information about a dataset in a privacy-compliant manner by describing the patterns of groups within the dataset while withholding information about individuals in the dataset. DP defines privacy not as a binary notion of "was the data of an individual exposed or not", but rather as a matter of accumulated risk. That is, every time a person's data is processed their risk of being exposed increases. To this end, the definition of differential privacy has parameters that quantify "privacy loss" -- the additional risk to an individual that results from her data being used, and factors it in.[13]	High level of accuracy Increased shareability of data and thus facilitate working with data without leaking personal information (here, adding a layer of noise adds more safety).	Difficult for small datasets: in order to maintain a precise enough model that is not too distorted by error from DP's 'noise,' the dataset needs to be large. Computationally intensive and requires significant resources and personnel to deploy it.[14] Risk of overstating privacy: since there is not one algorithm, there is no firm guarantee of privacy, just a statistical one. Risk of human error: during the process of identifying which is general and which is private information, human error could occur since differential privacy only guarantees to protect private information.	The United States Census Bureau used differential privacy to publish its 2020 results without revealing confidential information.[15] Apple employs differential privacy to accumulate anonymous usage insights (for example for Emoji suggestions or Health Type Usage) from devices like iPhones, iPads and Mac.[16]

[9] Since the utmost objective of this paper is to introduce the concept of existing privacy tech solutions and how they work, the listed use cases occasionally extend beyond image and video data.
[10] https://blogs.nvidia.com/blog/2021/06/08/what-is-synthetic-data/
[11] It should be noted that this is "augmented data", a sub-concept of synthetic data which equals "partially synthetic data" in which existing data is diversified, e.g. original is a street scene in sunny daylight, this can be altered through synthetic rain to also have a scene with rain.
[12] https://blogs.nvidia.com/blog/2020/08/21/curai-ai-healthcare-app/
[13] https://privacytools.seas.harvard.edu/differential-privacy
[14] https://www.brookings.edu/techstream/using-differential-privacy-to-harness-big-data-and-preserve-privacy/
[15] The United States Census Bureau used Differential Privacy to protect the identities of its 2020 census bureau. https://www.wsj.com/articles/census-data-change-to-protect-privacy-rattles-researchers-minority-groups-11627902000
[16] https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

3. The most important existing Privacy Tech Solutions

Privacy Tech	How it works	Specific Benefits	Challenges	Use Case[9]
Homomorphic encryption	The problem with classic encrypted data is that it must be decrypted to be "used", which makes it vulnerable. With homomorphic encryption, a computation can be done on encrypted data without ever decrypting it.[17] HE makes it possible for companies to outsource computation or storage of encrypted data; with HE, a cloud service can perform computations while protecting the customer's data with a state-of-the-art cryptographic security guarantee. The cloud only ever sees encrypted data, and only the customer can reveal the result of the computation.[18]	Protects the sensitive details of the actual data, but still allows the data to be analyzed and processed. As more companies and individuals switch to cloud storage and computing, HE allows for easily available secure computation technology.	Requires significant computational resources and is very slow - to an extent that it's not yet practical to use for many applications.	Facebook is currently working on using HE to analyze encrypted whatsapp data in its cloud to deliver users target advertisements while preserving end-to-end encryption.[20]
Federated Learning (FL)	It's a new decentralized way of machine learning that trains algorithms without exchanging the underlying data itself. In FL, the machine learning process happens locally at each participating institution (i.e. your phone,) and only the outcomes are transferred to the collaborative cloud-based ML model.[21] For example, your device downloads the then-current model, improves it by learning from data on your phone, summarizes the changes as a small update that is shared.	No personal data is transmitted, and data remains beneath firewalls and private. More opportunity to capture large data variability (i.e. diversity) and include different demographics into dataset Data does not have to be duplicated by each user for local model training but can scale naturally	There is a risk of FL models 'memorizing information.' For example, differential privacy is regarded as more secure. Because of its natural heterogeneity, FL means less traceability and accountability for researchers. They cannot investigate the training data nor their results as easily.	Google used Federated Learning in its Google Keyboard on Android. Meaning when Gboard shows you a suggested query, your phone locally stores information about the current context and whether you clicked the suggestion. [23]

[17] https://www.microsoft.com/en-us/research/project/homomorphic-encryption/
[18] https://www.forbes.com/sites/bernardmarr/2019/11/15/what-is-homomorphic-encryption-and-why-is-it-so-transformative/?sh=6eb6ff147e93
[20] https://www.slashgear.com/facebook-is-hunting-ways-to-push-targeted-ads-into-encrypted-whatsapp-chats-03685208/
[21] https://www.nature.com/articles/s41746-020-00323-1
[23] https://ai.googleblog.com/2017/04/federated-learning-collaborative.html

3. The most important existing Privacy Tech Solutions

Privacy Tech	How it works	Specific Benefits	Challenges	Use Case[9]
Deep Natural Anonymization	State-of-the-art anonymization solution for visual data. The technology automatically detects a personal identifier such as a face or license plate, and generates a synthetic replacement that reflects the original attributes. It then replaces the original personal identifier with the synthetic replacement in an irreversible way.	GDPR compliant as anonymized data is not subject to this regulation Protects identities while keeping necessary information (gaze direction, sociodemographic information, emotion) for analytics or machine learning. Builds upon established privacy solutions but adds an intelligence layer which makes it more robust and usable for a variety of use cases.	In comparison to classic anonymization (e.g. blurring for visual data), it requires more computational resources due to the technology's ML-component.	Deutsche Bahn and brighter AI's uses edge-anonymization and -analytics for automated passenger density analyses.[24]
Data Perturbation or 'Data Poisoning' and randomized smoothing	Adding 'pixel-level perturbations' ('cloaks') to images that stop facial recognition from identifying the natural persons in the image. The perturbations are typically so minuscule that they are imperceptible to the naked eye. Data perturbation relies on real or synthetic images - or a combination of the two - from a third party to generate additional 'cloaks'.[25]	High protection against user recognition, including from Microsoft's Azure Face API, Amazon Rekognition and Face++ Notably different from adding 'AI-spoof face paint' since the changes are smaller and less perceptible.	Does not protect user from existing or previous systems trained on images scraped of you on the internet and social media sites Always a cat-and-mouse game as recognition algorithms improve as well and take into account state-of-the-art methods to inhibit them.[26]	The SAND Lab at University of Chicago has developed Fawkes1, an algorithm and software tool that gives individuals the ability to limit how unknown third parties can track them by building facial recognition models out of their publicly available photos.[27]

A helpful way to understand privacy tech is to consider each a safeguard. Each privacy-preserving technology, whether it is an application of data (a verb) like homomorphic encryption or a new form of the data (a noun) like synthetic data, is a safeguard for user privacy in and of itself.

However, like other complexities in life, there is no black and white solution.[28] Instead, the privacy techs are often stacked and combined for more robust forms of privacy, with additional trade-offs and complexities at each step of the way. And it is important to note how, as technology improves exponentially, as do the privacy preserving techniques--in size and quality each year. For a deeper understanding of existing privacy tech and how it all compares, here are two use case scenarios from the German tech ecosystem that illustrate its impact.

[24] https://ecapital.vc/news/pilot-project-of-deutsche-bahn-and-brighter-ai-anonymized-passenger-analysis-to-comply-with-corona-distance-regulations/?cookie-state-change=1628254639131
[25] https://arxiv.org/pdf/2002.08327.pdf
[26] https://bdtechtalks.com/2021/04/05/machine-learning-data-poisoning-2/
[27] https://lowkey.umiacs.umd.edu/
[28] Consider the trade-offs we make while driving a car: the trade-offs between convenience and safety. For instance, we have three safeguards that increase our safety: an airbag, a seat belt and a stop-sign, but our safety is highest - and risk lowest - when we use all three. Though this requires the most time and investment.

4. Two Real world use cases from our community

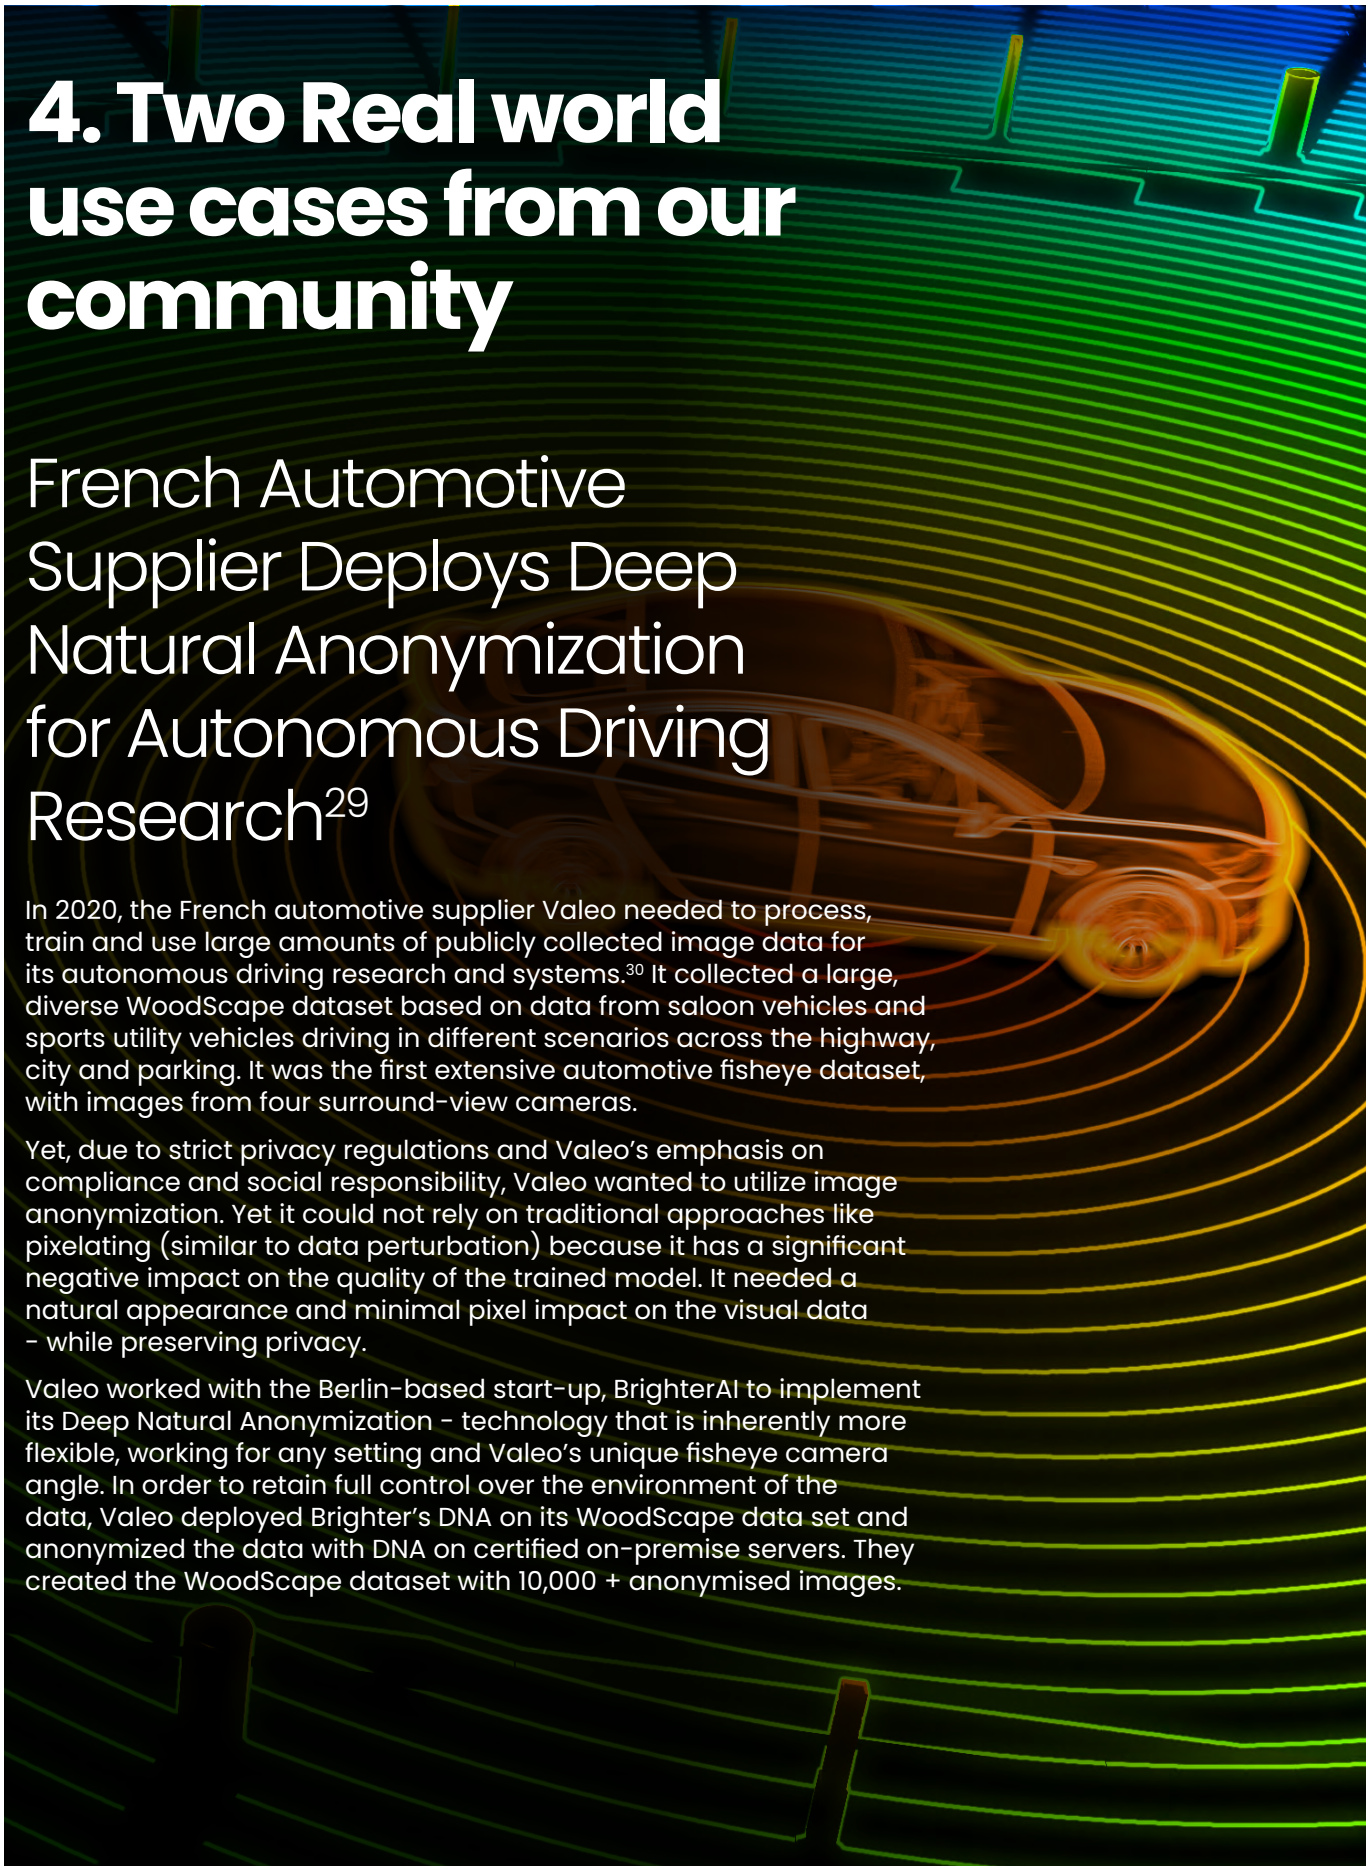

French Automotive Supplier Deploys Deep Natural Anonymization for Autonomous Driving Research[29]

In 2020, the French automotive supplier Valeo needed to process, train and use large amounts of publicly collected image data for its autonomous driving research and systems.[30] It collected a large, diverse WoodScape dataset based on data from saloon vehicles and sports utility vehicles driving in different scenarios across the highway, city and parking. It was the first extensive automotive fisheye dataset, with images from four surround-view cameras.

Yet, due to strict privacy regulations and Valeo's emphasis on compliance and social responsibility, Valeo wanted to utilize image anonymization. Yet it could not rely on traditional approaches like pixelating (similar to data perturbation) because it has a significant negative impact on the quality of the trained model. It needed a natural appearance and minimal pixel impact on the visual data - while preserving privacy.

Valeo worked with the Berlin-based start-up, BrighterAI to implement its Deep Natural Anonymization - technology that is inherently more flexible, working for any setting and Valeo's unique fisheye camera angle. In order to retain full control over the environment of the data, Valeo deployed Brighter's DNA on its WoodScape data set and anonymized the data with DNA on certified on-premise servers. They created the WoodScape dataset with 10,000 + anonymised images.

[29] Cross-chain middleware is software that acts as a bridge between different blockchains (on-chain data) and also non-blockchain (off-chain data) networks in this case. Chainlink's widely adopted blockchain oracle solution serves as middleware that helps on-chain applications securely access off-chain data and computation to achieve a wider range of functionalities.
[30] https://woodscape.valeo.com/dataset

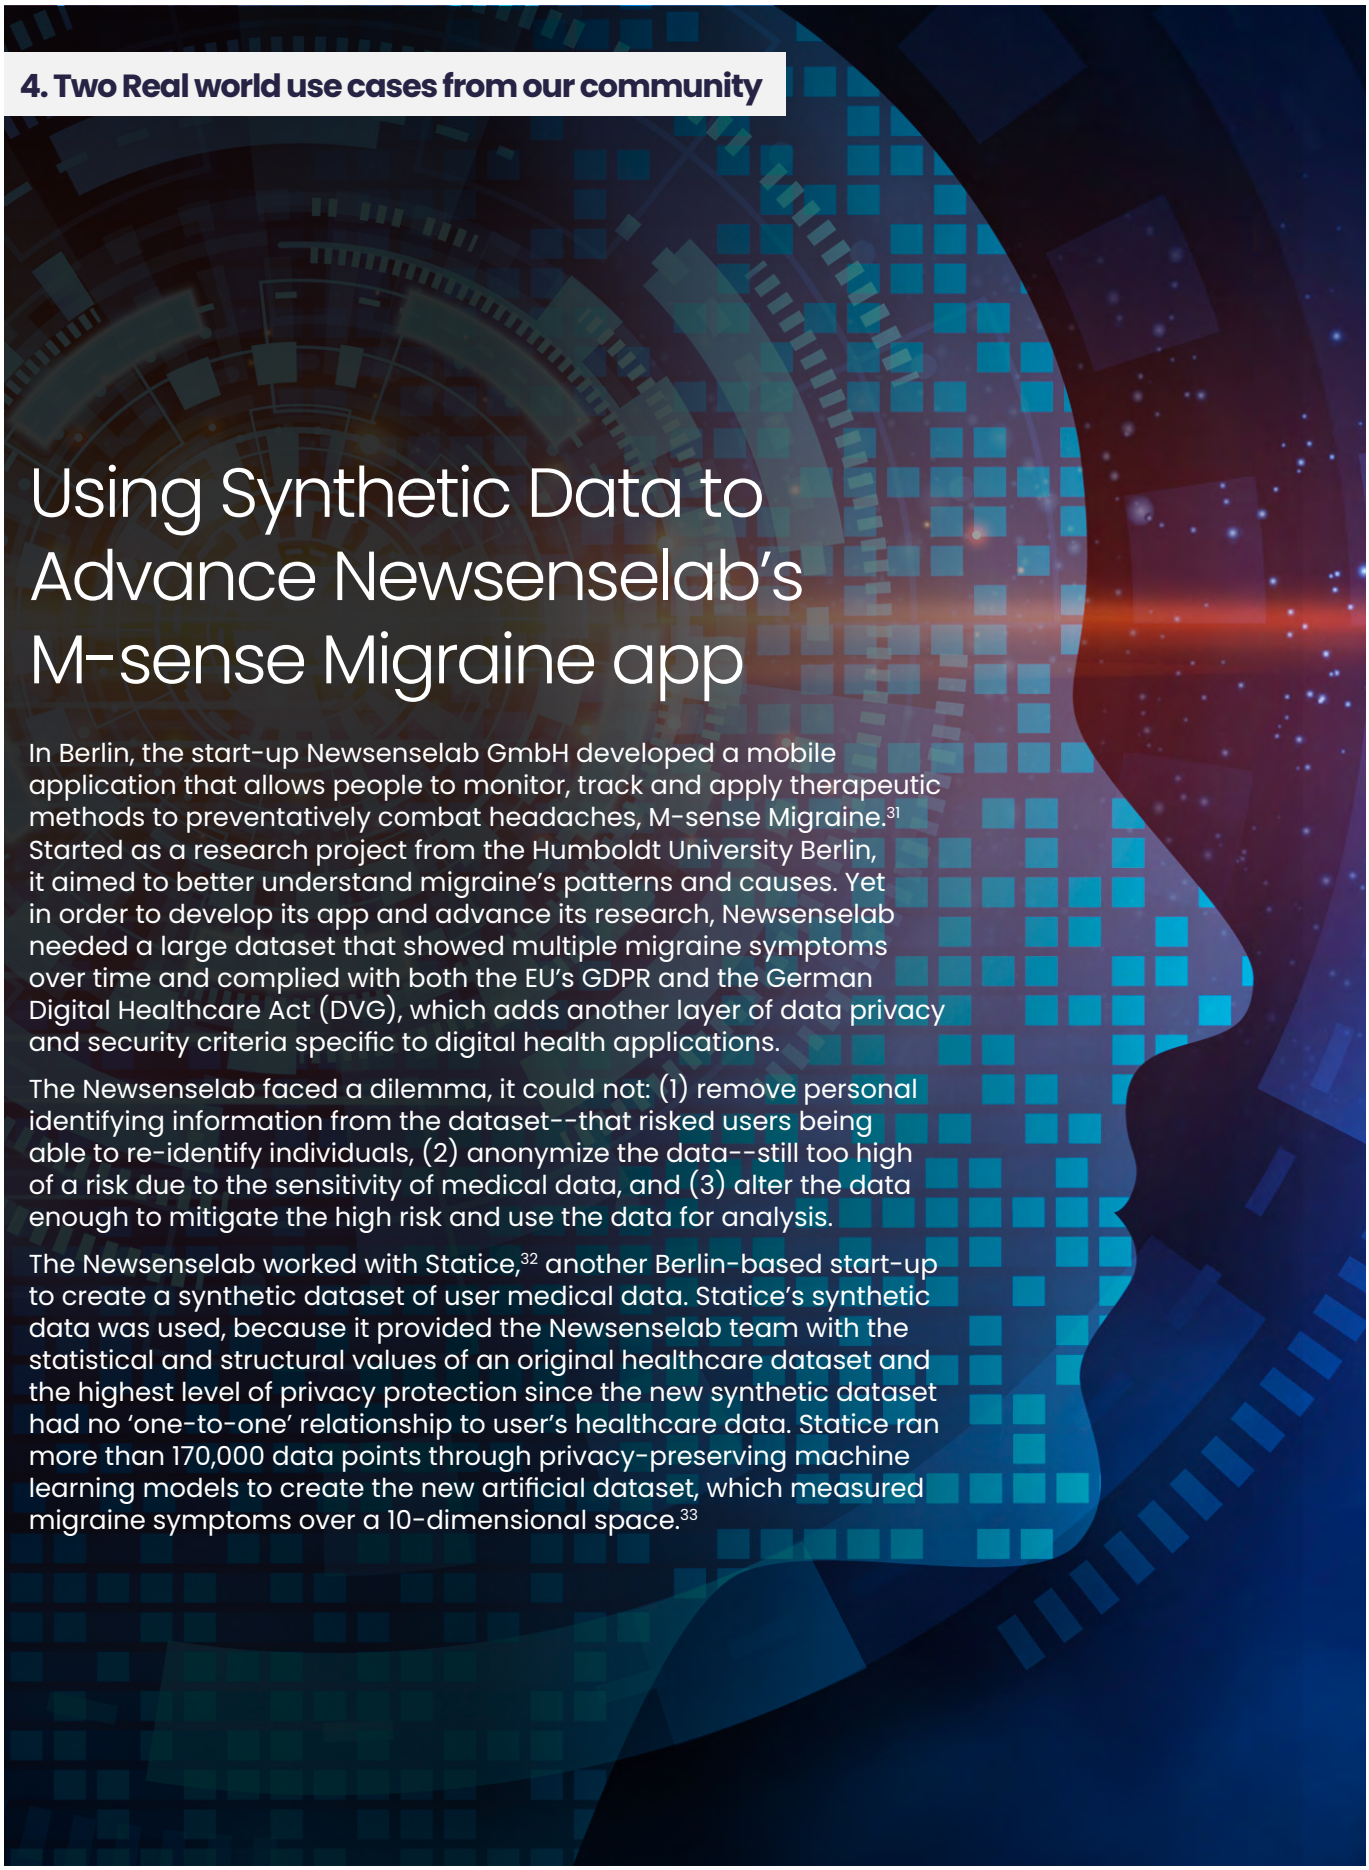

Using Synthetic Data to Advance Newsenselab's M-sense Migraine app

In Berlin, the start-up Newsenselab GmbH developed a mobile application that allows people to monitor, track and apply therapeutic methods to preventatively combat headaches, M-sense Migraine.[31] Started as a research project from the Humboldt University Berlin, it aimed to better understand migraine's patterns and causes. Yet in order to develop its app and advance its research, Newsenselab needed a large dataset that showed multiple migraine symptoms over time and complied with both the EU's GDPR and the German Digital Healthcare Act (DVG), which adds another layer of data privacy and security criteria specific to digital health applications.

The Newsenselab faced a dilemma, it could not: (1) remove personal identifying information from the dataset--that risked users being able to re-identify individuals, (2) anonymize the data--still too high of a risk due to the sensitivity of medical data, and (3) alter the data enough to mitigate the high risk and use the data for analysis.

The Newsenselab worked with Statice,[32] another Berlin-based start-up to create a synthetic dataset of user medical data. Statice's synthetic data was used, because it provided the Newsenselab team with the statistical and structural values of an original healthcare dataset and the highest level of privacy protection since the new synthetic dataset had no 'one-to-one' relationship to user's healthcare data. Statice ran more than 170,000 data points through privacy-preserving machine learning models to create the new artificial dataset, which measured migraine symptoms over a 10-dimensional space.[33]

[31] https://www.m-sense.de/en
[32] https://www.statice.ai/
[33] https://synthetic-data.statice.ai/hubfs/Resources/case_studies/CS_Newsenselab_x_Statice.pdf

5. The takeaway

This report attempted to understand, conceptualize and provide a first overview of privacy tech, because the privacy rules GDPR introduced mark the beginning of modern data-related legislation.

The European Union is in the midst of its Digital Decade, and as regulators progress and introduce more privacy regulations in the next five years, privacy tech will only become more relevant. This report has shown that there are already strong privacy tech use case scenarios--within Germany and the EU, and best practices. It has also illustrated how some players - especially from the startup ecosystem - are pioneers in this new technological field.

Even though the entire privacy tech market seems to be just taking shape, slowly finding its value propositions and its place in the greater digital economy, the technological advancements have a significant potential to change what 'made in Europe' means profoundly and sustainably. Privacy tech is poised to help reap the tremendous economic value of data and the benefits of digitalization without compromising an individual's important right to privacy. The privacy tech stack is full of solutions that make a zero-sum-game between data-based innovation and privacy obsolete.

However, in order for young and quintessentially European companies in the field to flourish, there are important key developments that need to happen:

#1 First, **legislators** developing new data policies or looking into amendments of existing privacy-related legislation must provide current privacy tech start-ups and scale-ups with legal certainty and publicly acknowledge the benefits they deliver for a higher overall level of data protection in our everyday lives. In particular, this extends to European legislators, who are in the midst of their digital decade and currently rolling out: the EU's Strategy for Data, the Data Governance Act, the Digital Services Act and the Digital Markets Act.

Bundesamt
für Sicherheit in der
Informationstechnik

"Predefined criteria, methods and tools are essential for the assessment of individual products and services. In the field of IT Security, the BSI already develops these in close cooperation with industrial and academic partners. These should be expanded to the privacy tech realm in order to evaluate the effectiveness of products."

Dr. Arthur Schmidt
– Head of the Artificial Intelligence Division, Federal Office for Information Security (BSI)

5. The takeaway

#2 Secondly, **public enforcement authorities** like the German Data Protection Authorities (DPA) are strongly encouraged to learn about the new and growing privacy tech stack available, and to support organizations exploring privacy tech as a way to establish compliant forms of data-use.

#3 Third, **established businesses** are encouraged to adopt and engage with privacy tech, and to understand privacy tech for what it is: a catalyst for digitization processes and services and the foundation for the development of entirely new data-based innovation(s).

#4 Fourth, **public sector employees** need to deploy privacy tech in lighthouse projects, especially concerning smart cities and other citizen-facing activities where digital solutions and a modern state is past due and public trust in robust privacy is needed. That way, the state can simultaneously help the new privacy tech vertical reach broad market maturity and develop compliant public digital infrastructures.

#5 Finally, **privacy tech companies and supporters** need to collaborate and publicly support one another to help position privacy tech's technological and societal value, help relevant economic, regulatory and social stakeholders understand the benefits and make best practices known to a broader audience.

We would like to thank everyone who supported this report. Last, we would like to thank all who already support this new field and encourage everyone who wants to engage to reach out to us and join the privacy tech community. We will keep observing market developments and technological advancements closely, and we are happy to share that this report is likely the first of many.

"Technological solutions that incorporate privacy and data protection fundamentals are a cornerstone for creating the modern web. This means placing the user in the driving seat and allowing them to decide who can get access to their data and on what terms, while always reflecting both the legal requirements and the spirit of GDPR."

Cornelius Witt
– Group Data Privacy Officer, eyeo

Imprint

Publishers
iconomy GmbH, Uhlandstraße 175, 10719 Berlin

E hello@iconomy.partners,

w iconomy.partners

KI Bundesverband e.V., Im Haus der Bundespressekonferenz, Schiffbauerdamm 40, 10117 Berlin,

E info@ki-verband.de

w www.ki-verband.de

Authors
German AI Association: Sina Youn

iconomy: Alexandra Magaard, Felix Styma

Publication Date
September 2021

Disclaimer
This publication has been prepared for general guidance only. The reader should not act according to any information provided in this publication without receiving specific professional advice. iconomy GmbH shall not be liable for any damages resulting from any use of the information contained in the publication.

© 2021 iconomy GmbH. ALL RIGHTS RESERVED.

KI BUNDESVERBAND

iconomy