2021 ISCA Symposium on Security and Privacy in Speech Communication
10-12 November 2021, Online

Brave New World?
Processing of personal data about employees under Art. 9 of GDPR in the context of human-robot interaction

Ivo Emanuilov[1], *Katerina Yordanova*[2]

[1]KU Leuven Centre for IT & IP Law
[2]KU Leuven Centre for IT & IP Law
ivo.emanuilov@kuleuven.be, katerina.yordanova@kuleuven.be

Abstract

Ensuring safe and secure collaboration between robots and human workers is essential for the successful deployment of artificial intelligence on the factory shop floor. Any such interaction depends on actions such as perception, sensing and action on the part of the robot which are, essentially, enabled by the real-time processing of personal data concerning the factory workers. The majority of these data would easily fall into the special category of personal data under article 9 GDPR, e.g. as biometric data. This means that their processing would in principle be prohibited unless allowed by one of the explicit exceptions. In this paper, we analyse which of these grounds may be applicable, taking into account the specifics of these interactions on the shop floor, that is, in an employment context with high level of safety risks. We explore the problem focusing on selected scenarios which are inspired from real or planned deployments of human-robot collaborative manufacturing technologies in the industries of aerospace, maritime and automotive manufacturing.

Index Terms: Human-robot interaction, Manufacturing, Data protection, Sensitive personal data, Biometric data, Employment, Safety, Security.

1. Introduction

Deployment of robotic systems on the shop floor in factories is not a new phenomenon. Robotic arms and other highly automated systems have already been supporting the manufacturing process for quite some time [1]. However, with the recent progress made in machine learning, such as enhanced sensing and perception, object recognition and even emotion detection, industries have realised the potential of human-robot collaboration in many new areas. For instance, an industrial robot could improve workplace safety by conducting tasks that could otherwise carry a high risk for the involved persons, such as heavy lifting, handling dangerous substances or operating high-speed machines with significant kinetic force. Furthermore, new adaptive capabilities of industrial robots now enable human workers to 'teach' them simple new tasks by example. Equally, an industrial robot on the shop floor configured to perform a particular task could also observe the actions of a worker who is alone in a particular area of the factory and perform rescue activities in case of an accident. We refer to such new modes of interaction between humans and collaborative robots (cobots) as to collaborative robotics.

This research is funded by the European Union's Horizon 2020 research and innovation programme under the Secure Collaborative Intelligent Industrial Automation (SeCoIIA) project, grant agreement No 871967.

The domain of autonomous mobile robots (AMRs) holds the most promising applications. AMRs can easily learn to navigate areas using downloadable maps to improve their situational awareness. Unlike automated ground vehicles, they are highly configurable which makes them a good fit for collaborative tasks that are typically of low added value but are repetitive and time-consuming. As Fragapane et al. explain [2]:

> *Compared to an automated guided vehicle (AGV) system in which a central unit takes control of scheduling, routing, and dispatching decisions for all AGVs, AMRs can communicate and negotiate independently with other resources like machines and systems and thus decentralize the decision-making process.*

While delegating such tasks to robots is by no means new, new scenarios would require higher a degree of collaboration with human workers. In turn, this would require high level of situational awareness which depends on the processing of data about the workers, often in real time. For example, audio and video feeds could be used to extract features about voice patterns, posture, body movements etc. which can then be used to train a robot to recognise them in humans. Thus, if a robot 'hears' a sudden loud bang or 'sees' a person lying on the floor, it can perform rescue operations, call emergency or communicate this event to a human operator for further action. Essentially, all of these actions would necessitate that the robot performs extensive, continuous processing of data concerning, in most cases, an identified or, more rarely, identifiable natural person. In other words, the robot would become a means of processing personal data under the EU General Data Protection Regulation (GDPR) triggering corresponding obligations for the data controller(s).

As the European Data Protection Board (EDPB) clarified in its Guidelines 07/2020 on the concepts of controller and processor, the term "means" does not only refer to the technical ways of processing personal data, but also to the 'how' of processing, which includes questions like which data shall be processed, which third parties shall have access to this data, when data shall data be deleted etc.[3] In other words, in an employment context, by determining the means and purposes of the processing of personal data, the employer who has decided to deploy cobots on the shop floor would undoubtedly qualify as a data controller.

As the majority of data processed in this context could easily fall under the heading of special categories of data (e.g., biometric data used for unique identification or data concerning health) under Article 9 GDPR, their processing would have to be justified by one of the exceptions to the general prohibition on processing of such data. However, the fact that such processing would take place in an employment context may limit

10.21437/SPSC.2021-14

significantly the available legal grounds and therefore the feasibility of long-term deployment.

This contribution analyses which of the exceptions in Article 9 GDPR could be relied upon by data controllers to justify processing of special categories of data using collaborative robots in an employment context. It is grounded in use case scenarios inspired by an ongoing research and innovation project in the field of secure intelligent collaborative industrial assets. Our objective is to demonstrate whether and when data controllers could justify the processing of data concerning shop floor workers by cobots for safety, security or business operations purposes. The main research question is whether the legal framework for processing of biometric data under Article 9 GDPR strikes a fair balance between legitimate economic interests of companies to optimise their operations for safety, security or economy purposes by deploying interactive robot systems on the shop floor and the rights and interests of workers as data subjects. To this end, the paper is structured in two main parts. First, we set the context by specifying the use and misuse cases which are considered relevant from an industrial perspective (section 2). Second, we introduce the content and purpose of the general prohibition under Article 9 (1) GDPR, sketch out the relevant categories of data for the use and misuse cases, and analyse the applicability of the exceptions under Article 9(2) GDPR to these scenarios (section 3).

2. The (mis)use cases of human-robot collaboration

In an industrial context, the collaboration between humans and robots may lead to both desired and undesired consequences. We classify these, respectively, in the groups of use and misuse cases. A use case is a typical scenario where human-robot collaboration has added value for a business process, such as optimisation of a logistic operation. A misuse case, on the other hand, concerns a scenario where human-robot collaboration may unleash undesired consequences which are the result of either intentional or unintentional actions by an insider or an external party and where, for the purpose of mitigation of further damages, additional personal data may need to be processed. These undesired consequences may vary from data loss or corruption to damages to property or even loss of human life.

2.1. Use cases

Our use case explores three related shop floor activities, namely logistics and inventory automation, AI learning by demonstration from a human worker, and lone worker monitoring and rescue. These activities rely on a set of technical capabilities such as trackless navigation, avoidance of obstacles, recognition of shapes and voices, tablet plotting and learning by demonstration. The automated logistic and inventory operations involve activities where the robot, for example, visits certain locations, carries tools for human workers, performs measurements, monitors supplies and keeps an inventory of materials to prevent mistakes or theft. Learning by demonstration refers to the robot's ability to learn from a human worker without the need of additional interfaces. There are two main modes of learning by demonstration: passive observation and kinesthetic teaching. In passive observation mode, the robot learns from a human worker by observing their actions, that is, without direct interaction, e.g., by following the path of the worker or analysing their gestures and mimics. In kinesthetic teaching mode, there is a direct interaction between the robot and the human worker. In this case, the human worker demonstrates by physically interacting with the robot, e.g. by moving it in space whereby the robot learns by recording the actions through onboard sensors resulting in the generation of training data. Finally, the activity of lone worker protection concerns a situation where a human operator might be working alone in a shop floor area when an accident or an incident occurs. The robot would then be able to detect such events, report them and assist first responders, thereby reducing the response time. Usually, this activity would be supported by the robot's 3D thermal camera which ensures that the identified object is indeed a person in trouble.

2.2. Misuse cases

Our misuse case explores a layered scenario where an adversary gains access to the robot's control system to change its behaviour or perform evasion, poisoning or even model stealing attacks. The modalities of this scenario include network attacks (e.g., an adversary gaining access via a wireless interface), offensive AI attacks (e.g., an adversary develops an AI capability that manipulates the robot's models during either training or operations phase) and insider's attack (e.g., where a human worker or operator deliberately fools the robot's sensing, learning or decision-making functions, for example, by performing learning by demonstration in unsafe scenarios).

3. Data protection and human-robot collaboration

3.1. Prohibition of processing of special categories of personal data in an employment context

The processing of special categories of personal data is in principle prohibited by Article 9, para 1 GDPR, subject to explicitly provided derogations. Special categories of data include *data revealing racial* or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, *biometric data for the purpose of uniquely identifying a natural person*, *data concerning health* or data concerning a natural person's sex life or sexual orientation. Processing of these special categories of data is only allowed under the derogations listed in the second paragraph of article 9, as modified by conditions or limitations introduced in national law pursuant to article 9, paragraph 4 GDPR.

The GDPR is relatively laconic when it comes to the employment context as a source of special requirements or conditions. More specifically, recital 155 reads that:

> *Member State law or collective agreements, including 'works agreements', may provide for specific rules on the processing of employees' personal data in the employment context, in particular for the conditions under which personal data in the employment context may be processed on the basis of the consent of the employee, the purposes of the recruitment, the performance of the contract of employment (...)* .

The provision of Article 88 GDPR offers some general guidance on the processing of personal data in the context of employment. Thus, paragraph 1 stipulates that it is within the discretion of Member States to adopt more specific rules to ensure the protection of the rights and freedoms in respect of the processing of employees' personal data in the employment context. These rules may take the form of either a legislative instrument or a collective agreement. Importantly, in the third

paragraph of Article 88, the GDPR specifies that such rules

> *include suitable and specific measures to safeguard the data subject's human dignity, legitimate interests and fundamental rights, with particular regard to the transparency of processing, the transfer of personal data within a group of undertakings, or a group of enterprises engaged in a joint economic activity and monitoring systems at the work place*

The regulation also provides for some special conditions in Article 9, paragraph 2, lit. b, which concern a derogation from the general prohibition on processing special categories of data where

> *processing is necessary for the purposes of carrying out the obligations and exercising specific rights of the controller or of the data subject in the field of employment and social security and social protection law in so far as it is authorised by Union or Member State law or a collective agreement pursuant to Member State law providing for appropriate safeguards for the fundamental rights and the interests of the data subject*.

The regulation provides yet another option for controllers to derogate from the prohibition by allowing processing of special categories of data in article 9, para 2, lit. h where

> *processing is necessary for the purposes of preventive or occupational medicine, for the assessment of the working capacity of the employee, medical diagnosis (...)*.

This brief review of the legal framework shows that the main concerns in the context of employment refer to the reliance on consent as a lawful ground for processing of employee data and processing for specific purposes (e.g., safety and health or protection of employee's or customer's property). Importantly, the regulation puts certain requirements to any rules that govern one or more of these matters. Thus, any such rules must safeguard the data subject's human dignity, legitimate interests and fundamental rights, with particular regard to the transparency of processing and workplace monitoring.

3.2. Categories of relevant personal data

For the purposes of our analysis, we deem the categories of data revealing racial origin, biometric data for the purpose of uniquely identifying a natural person and data concerning health as three most likely candidates of special categories of data. These legal categories, however, beg certain clarifications as there may be overlaps under certain circumstances, leading to terminological and practical confusion.

3.2.1. Data revealing racial origin

The regulation does not define 'data revealing racial origin'. The text of recital 51 provides only limited guidance in this respect, by specifying that

> *the term 'racial origin' in this Regulation does not imply an acceptance by the Union of theories which attempt to determine the existence of separate human races*.

The problem of defining 'data revealing racial origin' becomes obvious when one tries to distinguish between data revealing racial origin and biometric data. For example, a simple facial image obviously reveals whether someone is a person of colour or not. The distinction is important in the context of our use case as the robot would have to be trained to recognise people of colour as well as white people and it is only through diversity in the training data set that this objective could be achieved. However, when does the processing of facial images trigger the prohibition for processing of 'data revealing racial origin'? In other words, when is a facial image considered 'data revealing racial origin'? Indeed, the Article 29 Working Party acknowledged the connection between the different categories of sensitive data already in its Working document on biometrics of 2003 by saying that '[s]ome biometric data could be considered as sensitive (...) and in particular, data revealing racial or ethnic origin or data concerning health' [4]. This intrinsic link was also confirmed in the group's follow-up opinion of 2012, suggesting that '[i]n order to assess the sensitivity of data processed by a biometric system the context of the processing should also be taken into account' [5]. Furthermore, the working party recognises that the 'processing of biometric data could be used to determine [other] sensitive data, in particular those with visual cues such as race, ethnic group or perhaps a medical condition.' [5]. Therefore, it could be concluded that by processing biometric data of a worker, a data controller is very likely to be processing also other special categories of data simultaneously. In this case, to the extent any additional conditions exist in either national or EU law regarding the processing of such special categories of data, they would apply concurrently with the mainline requirements of GDPR regarding processing of biometric data. The case seems clear when a data controller knows for sure that they are processing biometric data. Things are less clear, however, on the question of whether and when a photograph, which is not *per se* biometric data, would fall under the prohibition of Article 9 on grounds that it reveals data about one's racial origin, and when it would not. Essentially, the problem boils down to the different metamorphoses of the definition of biometric data in the GDPR.

3.2.2. The metamorposes of biometric data

The category of biometric data processed for the purpose of uniquely identifying a natural person is a new category of sensitive data in the GDPR [6]. Therefore, it is a new category besides data revealing racial or ethic origin, religious or philosophical beliefs etc. This is seemingly the only category of biometric data whose processing is considered sensitive [6]. The differences between identification and verification do not seem to have been caught by the regulation in this respect. As Kindt observed, it is only the provision of article 4(14) that mentions the difference by referring to the construct 'allow or confirm the unique identification'.

The narrow definition of biometric data in the GDPR implies that the collection, storing and processing of 'physical, physiological or behavioural characteristics of a natural person does not fall under any specific regulation or benefit from any specific protection, other than the general data protection regime under the GDPR' [6]. Indeed, the mere creation of a 'database with facial images or fingerprints without biometric processing may hence not be considered a database with biometric data or a biometric database' [6]. However, as Kindt rightly points out, compiling databases is a prerequisite for biometric identification and, as such, the risks inherent in the creation of such databases should play a role in the evaluation of the risks to the fundamental rights and freedoms of data subjects [6].

In the context of the use case involving a cobot on the shop floor, the assessment would depend on whether the robot performs specific technical processing of the real-time data feed allowing or confirming the unique identification of the data subject. If the answer is negative, then the processing would qualify as mere processing of personal data under the general rules of the GDPR.

At first sight, such an assessment sounds relatively simple. However, it is precisely the 'specific technical processing' requirement that causes a lot of confusion in the definition and its application to practical use cases. As Kindt points out, it seems that the criterion here is the use of data, not the data itself [6]. However, she continues, the progress in biometric technology has largely obviated the differences between so-called 'specific' and 'ordinary' data. She highlights that it is not clear 'which characteristics are to be considered 'ordinary': while photographs of faces seem to be considered ordinary' and asks if this applies equally to photographs of fingerprints (fingerprint images) [6]. It is to note that even if the data results from 'specific technical processing' relating to physical, physiological or behavioral characteristics of a natural person, as long as they are not used for unique identification, it would be the general rules of the GDPR alone that would apply to it [6]. The use case and the technical nature of the processing would therefore be key to adequately assess the scope of the controller's obligations regarding the processing of biometric data.

3.2.3. Health data

The provision of recital 35 GDPR reads as follows:

> *Personal data concerning health should include all data pertaining to the health status of a data subject which reveal information relating to the past, current or future physical or mental health status of the data subject.*

The recital provides several types of data which usually would qualify as health data. These include a number, symbol or particular assigned to a natural person to uniquely identify the natural person for health purposes; information derived from the testing or examination of a body part or bodily substance; and any information on, for example, a disease, disability, disease risk, medical history, clinical treatment or the physiological or biomedical state of the data subject independent of its source. The provision of Article 4(15) GDPR lays down the legal definition of 'data concerning health' which includes

> *personal data related to the physical or mental health of a natural person, including the provision of health care services, which reveal information about his or her health status*

Much like biometric data for the purpose of unique identification, data concerning health fall within the ambit of the prohibition under Article 9(1), subject to the exceptions under Article 9(2) and national measures under Article 9(4).

What is of special interest in this context is the processing of sensitive personal data under Article 9(2)(c) when justified as necessary to protect the vital interests of the data subject or of another person where the data subject is physically or legally incapable of giving their consent. In a working document of 2007, the Article 29 Working Party highlighted that

> *[Any such] processing must relate to essential individual interests of the data subject or of another person and it must – in the medical context – be necessary for a life-saving treatment in a situation where the data subject is not able to express his intentions. [7].*

In the context of our use case, one plausible scenario is where the processing of data concerning health, such as information about a disease (e.g., diabetes, low blood pressure etc.), is necessary so that the robot can identify whether a particular behaviour might endanger a worker's safety when collaborating with it. In case of an accident where the robot might have to call in first responders, it is likely that such processing might be justified on grounds of vital interest of the data subject. However, the general rule is that this is an exception to the prohibition and as such it must be construed narrowly.

4. Processing in an employment context: data protection beyond (explicit) consent

The provision of Article 9(2)(1) specifies that the general prohibition for processing special categories of data does not apply where the data subject has given explicit consent to the processing of those personal data. Such consent must always be given for one or more specified purposes. Furthermore, there may be cases where Union or Member State law provide that the general prohibition may not be lifted by the data subject anyhow. Reliance on consent is particularly problematic in an employment context. As the EDPB clarified in its Guidelines 05/2020 on consent [8]:

> *An imbalance of power also occurs in the employment context. Given the dependency that results from the employer/employee relationship, it is unlikely that the data subject is able to deny his/her employer consent to data processing without experiencing the fear or real risk of detrimental effects as a result of a refusal. It is unlikely that an employee would be able to respond freely to a request for consent from his/her employer to, for example, activate monitoring systems such as camera observation in a workplace, or to fill out assessment forms, without feeling any pressure to consent. (...) For the majority of such data processing at work, the lawful basis cannot and should not be the consent of the employees due to the nature of the relationship between employer and employee.*

This view more or less summarizes the big issue with reliance on consent in an employment context. The power imbalance is even more obvious in the collaborative robotics use case since the worker would have little choice but to 'cooperate' with a machine that is part of their working environment. Furthermore, such a robot would not be merely a piece of machinery that the human worker can operate but rather a partner relying on active cooperation on the part of the human and, therefore, inevitably on processing of personal data.

In its Opinion 2/2017 on data processing at work, the Article 29 Working Party provided several scenarios that are typical of today's workplace. Of relevance to this analysis is the case of processing operations resulting from monitoring information technology usage at the workplace [9]. In the majority of these scenarios, an employer would seek to justify processing of personal data on grounds of its legitimate interest. Thus, Article 6(1)(f) stipulates that the processing is lawful when:

> *processing is necessary for the purposes of the legitimate interests pursued by the controller or by*

> *a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.*

The use cases discussed in the Article 29 Working Party (WP29) opinion are organised around wearable devices, security applications or tracking via unseen software. These are considerably different compared to a cobot, but some of the general guidelines would apply nonetheless. Thus, the WP29 suggests that:

> *Firstly, employers utilising these products and applications must consider the proportionality of the measures they are implementing, and whether any additional actions can be taken to mitigate or reduce the scale and impact of the data processing. (...) Secondly, employers must implement and communicate acceptable use policies alongside privacy policies, outlining the permissible use of the organisation's network and equipment, and strictly detailing the processing taking place. [9]*

Furthermore, the WP29 acknowledges that in certain cases the monitoring of employees is the result of using specific applications and recommends measures, as follows:

> *It should be ensured that employees can designate certain private spaces to which the employer may not gain access unless under exceptional circumstances. [9]*

A similar approach could be applied to a certain extent to cobots. Thus, an employer could designate areas on the shop floor which employees can use as private spaces and where cobots would not be allowed at all. Specifically concerning processing operations involving video monitoring, the WP29 provides guidance in the following sense:

> *With the capabilities given by video analytics, it is possible for an employer to monitor the worker's facial expressions by automated means, to identify deviations from predefined movement patterns (e.g. factory context), and more. This would be disproportionate to the rights and freedoms of employees, and therefore, generally unlawful. (...) employers should refrain from the use of facial recognition technologies. There may be some fringe exceptions to this rule, but such scenarios cannot be used to invoke a general legitimation of the use of such technology. [9]*

These suggestions are also in line with the general stance towards prohibition of facial recognition systems that has been gaining ground in the EU recently. It is likely that extensive and obtrusive monitoring using cobots would hardly meet the threshold of the legitimate interest impact assessment prescribed by Article 6(1)(f) GDPR. Indeed, the WP29 highlights in one its conclusions that:

> *The legitimate interest of employers can sometimes be invoked as a legal ground, but only if the processing is strictly necessary for a legitimate purpose and the processing complies with the principles of proportionality and subsidiarity. [9]*

Employers must also be very clear towards their employees about any such monitoring. Concrete transparency measures highlighted by the WP29 include:

> *Effective communication should be provided to employees concerning any monitoring that takes place, the purposes for this monitoring and the circumstances, as well as possibilities for employees to prevent their data being captured by monitoring technologies. [9]*

Finally, data processing in an employment context must be proportionate to the risks faced by an employer. The principle of data minimisation should guide employers when deploying any technologies that may involve monitoring of their employees [9]. In other words, an employer should always seek the least intrusive of solutions that can achieve their legitimate objectives; no more and no less. Importantly, also in line with the case law of the European Court on Human Rights, employers must be aware that:

> *The fact that an employer has the ownership of the electronic means does not rule out the right of employees to secrecy of their communications, related location data and correspondence. [9]*

5. Conclusion

The deployment of cobots on the shop floor is an attractive opportunity for many industrial actors. However, the risks to the fundamental rights and freedoms of data subjects should also come first in the list of decisive factors for and against deployment. The employment context is dominated by the inherent power imbalance between an employer and an employee. This imbalance reverberates in the status of employees as data subjects. Data protection law equips data subject with tools to keep in check the (legitimate) commercial interests of companies, even more so in an employment context. The result is a limited number of lawful grounds which may be relied upon by an employer and the need for continuous assessment of the risks to the rights and freedoms of data subjects. While early-stage experimental deployment of cobots on shop floors could rely on lawful grounds such as consent, it is highly unlikely that this would be the case in large-scale deployments. Furthermore, collaborative manufacturing would involve a growing number of actors, meaning that the data controller's responsibilities would become even more diffuse. The lack of clarity around some legal categories, such as biometric data, could deter the deployment of such technologies even further. The fact that a great deal of manufacturers would contract such cobots from companies that are external to the manufacturing process implies that new trust relationships would have to be built around the robustness of the technologies underpinning cobots, namely artificial intelligence and robotics. Finally, the room for manoeuvre given to EU Member States to maintain or introduce further conditions, including limitations, with regard to the processing of biometric data could create fragmentation in the EU, especially if states choose to legislate on a case-by-case basis.

6. References

[1] Federal Ministry of Labour and Social Affairs, "Work 4.0: White Paper." [Online]. Available: https://www.bmas.de/SharedDocs/Downloads/EN/PDF-Publikationen/a883-white-paper.pdf

[2] G. Fragapane, R. de Koster, F. Sgarbossa, and J. O. Strandhagen, "Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda," *European journal of operational research*, vol. 294, no. 2, pp. 405–426, 2021.

[3] EDPB, “Guidelines 07/2020 on the concepts of controller and processor in the GDPR,” 2020. [Online]. Available: https://edpb.europa.eu

[4] WP29, “Working document on biometrics,” 2003.

[5] ——, “Opinion 3/2012 on developments in biometric technologies,” 2012.

[6] E. J. Kindt, “Having yes, using no? About the new legal regime for biometric data,” vol. 34, no. 3, pp. 523–538.

[7] WP29, “Working Document on the processing of personal data relating to health in electronic health records (EHR),” 2007.

[8] EDPB, “Guidelines 05/2020 on consent under Regulation 2016/679.” [Online]. Available: https://edpb.europa.eu/

[9] WP29, “Opinion 2/2017 on data processing at work,” 2017.