Intelligent Environments 2019
A. Muñoz et al. (Eds.)

doi:10.3233/AISE190044

Operability of the GDPR's Consent Rule in Intelligent Systems: Evaluating the Transparency Rule and the Right to Be Forgotten

Gizem GULTEKIN VARKONYI [a, 1]
[a] *Faculty of Law and Political Sciences, University of Szeged, Hungary*

Abstract. This paper discusses the difficulties to obtain valid consent for data processing activities executed by Artificial Intelligence (AI) systems. Although the European Union's General Data Protection Regulation (GDPR) is one of the most updated and most comprehensive legal instruments ensuring the right to data protection, the so-called consent obligation is challenged by several technical and practical issues in the case of AI systems. Data controllers obligation to demonstrate transparent information and to ensure the right to be forgotten is being challenged by the technical capabilities of AI taken together with some opaque statements in the GDPR. More detailed explanations should be delivered by the EU Institutions on the implementation of the GDPR for data controllers offering AI systems.

Keywords. Data protection, Artificial Intelligence, consent, transparency, Right to be Forgotten

1. Introduction

Current AI applications have many abilities that could not have been realized without data and hardware availability and advancements in AI engineering knowledge. Human beings interact and share several personal issues with machines on a daily basis, such as their pictures, videos, political opinions, emails, text messages, call logs, personal documents, browser history, financial data, location data and more. People seem very generous when sharing such data without assessing the consequences of reaching indefinite places and persons within seconds. Possibly this is what Mark Zuckerberg meant when he said almost 10 years ago that the era of privacy is over, it is no longer the social norm. On the other hand, Internet and Social Media grow every day with the help of personal data and become a treasure chest for the development of AI technologies. IDC analysts predict that in 2025 175 zettabytes of data will be available in data storages such as clouds, smart phones, Internet of Things (IoT) devices, or cell towers [1]. Data availability, therefore, is certainly one key factor for AI systems contributing their ability to see, hear, understand [2], learn, plan, reason, negotiate [3] solve problems, recognize

[1] E-mail: gizemgv@juris.u-szeged.hu.

voices and faces, process languages, make or support decisions, guide an interaction with human in a personalized way or [4] even in a social and an emotional way.

People often make decisions based on different criteria: (i) personal experience or knowledge, (ii) based on an analysis of several external conditions, (iii) processing and extracting meaning out of data is at the core of human decision-making processes. AI obviously simulates human decisional patterns in many aspects also in decision-making: it is able to exhibit signs of rational thinking, it is capable to adapt to a detected change in circumstances and it is able to engage in autonomous actions based on the above. AI's data collection and processing ability is based either on the past data that is used for initial training purposes, or on new data that it collects and analyzes itself as an outcome of its learning curve [5].

When it comes to defining what AI is, the instinctive answer is that there is no single definition. However, ones understanding of what an AI is reflect their approach to this topic. This is demonstrated by the example of software engineers who define technical terms very differently from lawyers, but the same is true vice versa for legal terms in this realm (e.g. liability, damages, notions of self, the human concept, rights and obligations). According to the EU [6], AI is a system that displays intelligent behaviour by analysing its environment and taking action, with some degree of autonomy, to achieve specific goals (emphasis added). Although the definition was later expanded enormously [7], the new definition was adopted after the GDPR entered into force, and is rather an ethical concept, not necessarily a legal one (with the only legal element of the notion being degree of autonomy). This means that at the time of drafting the GDPR, EU lawmakers did not give enough emphasis to the AIs undeniable relationship with data and Machine Learning, but rather focused on the intelligence and autonomy aspects of the AI systems. Furthermore, this definition excludes AIs learning ability, but in turn focuses on the outcome (the action) and its goal-driven behavior to generate this outcome. Such a generic definition reflecting the EUs position towards AI, excluding the two most important aspects of AI applications (data processing and ML) might be the reason why the GDPR would fall short regarding its applicability to AI systems, as we will present in the following sections.

Certainly, more personal(ized) services mean sharing more personal data required for AI systems and people will not be afraid of sharing their data with a machine in exchange for said services [8] customized to them. In this paper, we intend to cover AI systems designed for personal(ized) use to manage peoples monotonous daily activities. We first establish the possible legal basis for operating such systems in the GDPR. Then we discuss how difficult it would be to comply with the GDPRs certain principles and provisions since no specific guidelines are designed for the GDPRs implementation on AI technologies.

2. Consent as a Legal Basis

What might be the applicable legal basis enabling an AI system to process data and make predictions? The answer is quite easy. As we focus on the EU legal framework in this paper, the GDPR is the only applicable law in this domain. The first question is, however, which legal bases within the GDPR could permit data controllers to operate their AI system?

Besides giving data subjects effective control over their own data, the GDPRs essence is about ensuring the legality of processing: Article 6 lists in what cases processing qualifies as legal. At least one of the six clauses mentioned therein should be valid, to constitute legal processing. Proving the legality of data processing traditionally carried out e.g. by hospitals, banks, or businesses (not offering AI-based services) is already problematic due to technological issues and to automation of daily tasks. As the Cambridge Analytica case taught everyone well, the absence of consent makes it difficult to identify breaches or incidents, and to interpret the case in the broadest possible meaning under the current legislation. The simple example of mobile apps installed on personal mobile phones could be mention as case studies for AI-based processing services when looking at the legal basis of processing under the GDPR. An application set up on the phone by default might be an essential part of the phone, for example, its operating system. If the components of the application essential to make the phone work require personal data processing, then the legal basis for such data processing would be most probably based on performance of a contract (Art. 6/b). If there is another app aiming to offer personal(ized) services to the users, then the data subjects consent (Art. 6/a) will be the legal basis for processing. Data processing on mobile phones in personal use are neither legal obligations for the data controller (Art. 6/c), nor are they necessary to protect vital interests of any person (if we exclude the extreme possibly theoretical cases (Art. 6/d). When a data subject uses a mobile phone, legal persons behind the mobile phone, e.g. manufacturers, or software developers, do not process data to execute tasks related to their public interest (Art. 6/e) if the app offers only personal(ized) services. The legitimate interest rule, as well as performing a task carried out for public interest do not apply unless the mobile phone is part of a public service. Referring back to the performance of a contract and consent rules, even if the application is essential to operate the mobile phone, once the user starts using it (by entering into the sales contract), then (s)he will immediately come across pages of consent language often disregarded (as part of the contractual fine print), but still continues to use the application because of its personalized components and relevant advantages (We could call this the trade-off effect). Applications want to know about people to offer them better personalized services. A personal health assistant app, e.g., asks peoples consent to reach their messages, contacts, pictures, calendars and sometimes to their social media accounts. However, it is usually unclear why the app reaches such wide range of personal data, and the issue of relevance surfaces. What possible consequences could be drawn from a persons social media presence regarding their health? If not from the photos taken of their meals before posting them to Instagram With this example in mind, we will further present how obtaining valid consent is a challenge in the case of AI-based processing services, since presenting transparent information and ensuring data erasure might not be as easy as the GDPR and its guidelines make it seem.

2.1. Giving Consent

Pursuant to Article 6 of the GDPR, consent is a legal basis and the expression of data subjects' will which safeguards their freedom to control their personal data [9]. Article 7 of the GDPR sets forth the conditions of said consent, such as data subjects right to manage and withdraw it. This constitutes informational self-determination, which refers to the right of the data subjects to freely share their personal data while giving them

"control of the use that is being made of his data[10], or competence to prevent unwanted situations/interference. The principle of data subjects' free will plays a crucial role since it informs their intention to permit or not to permit the further use of their data by the data controller operating and intelligent (AI) system[11].

For data subjects to be able to give consent, they must be aware of almost everything related to the fate of their data. It is the data controller, surely, who shall inform them prior to obtaining consent (point 3 of Article 5 of the GDPR), so they could assess the risks properly and make a decision. Data controllers information obligation is placed in several parts of the GDPR, but we will focus only on Article 12 in following analysis. If data controllers fail to present true and complete information about the data processing activities, they obviously fail in fulfilling their legal obligations, firstly, the transparency principle.

2.1.1. Transparency Rule

Pursuant to Article 12 of the GDPR explaining the conditions of transparency, data subjects are given the right to obtain information, for example, on the purposes of the processing; to request erasure (better known as right to be forgotten), and meaningful information about the logic involved with automated decision making, including profiling activities that data controller processes data for. Presenting meaningful information is an important case since it orders data controllers to somehow explain to the data subjects how the AI system works. Although there are many issues related to purpose limitation by data controllers and data subjects right manage personal data within AI systems, we exclude both instances in order not to extend the scope of this paper.

2.1.2. Meaningful Information

Articles 13, 14 and 15 of GDPR entrust data controllers with presenting meaningful information related to data processing activities they carry out including such activities in which the decision is made algorithmically. What constitutes meaningful information has been discussed in literature from several points of view. Wachter et. al. firstly argues that right to be informed within the GDPR is an ex post right which would contravene the the essence of consent, and further stresses that the right to explanation should be inserted in the GDPR to make the rule more consistent and clear [12]. Selbst and Powles [13], on the other hand, strongly emphasize that explanation of meaningful information already is the right to explanation, and meaningful information refers here to any information regarding system functionality. Although both views cannot be easily and clearly understood neither from the related articles, nor Recitals, and nor from the EDPS opinions, we think that the GDPR is practically unclear on explaining the concept of meaningful information. We question, firstly, who should determine the meaningfulness of information the data subject in a personalized (subject-specific) fashion, or other standards should be applied to assess meaningfulness in the eyes of the general public. [14] (This again brings about the issue of relevance interpreted as meaningful what is meaningful information in a particular context, it is also relevant to it.) If the answer is subject-specific meaningfulness, then we argue that data subjects have no interest in any complex technical terms specific to the applied AI, and would prefer the simplest and clearest explanation (laymens terms); while other data subjects might prefer more detailed information in line with their level of AI knowledge. (This necessitates a context and subject-dependent as-

sessment of meaningfulness.) Furthermore, transparency, as what the developers understand under it and present information (to the data controller) may not be understood the same way, which generates more complications especially for people without technical background knowledge [15]. On the other hand, a generic explanation may not be understood clearly by every data subject in a similar way (as it might leave room for ambiguity). The question of how to measure data subjects' understanding especially when they interact with AI systems only via a screen or during a natural talk, erects another obstacle to assess the operability of the meaningfulness concept in practice.

Whether it is a duty of data controllers to ensure each data subjects' understanding, which is obviously not the case according to the GDPR, carries the discussion to another dimension. Based on this loophole, data controllers like the tech-giants (e.g. Google, Facebook, Amazon) which provide their services based on algorithmic calculations, do not pay attention to whether the users would be able to easily understand the information provided and track and control their data within the system. This problem is related to the existence of insufficient regulations and difficulties to regulate diverse populations that AI systems serve [16]. For example, data controllers may tend to circumvent the stress of fulfilling their legal obligations and as a result, provide explanations that are not accurate. Data controllers fearing the loss of their users trust or unwilling to disclose shortcomings to the competitors may prefer not to reveal privacy losses (data breaches) within the system to the users transparently, even if they implement PETs or other technological solutions such as differential privacy which also has its own technical shortcomings in the implementation. [17]

2.1.3. Intelligible Form

One may claim that the EU lawmaker already took the possibility into account and repeated in the GDPR the intelligible form requirement for data con- trollers to better fulfill transparency and consent principles. The Court of Justice of the European Union (CJEU) received several questions regarding the form of the explana- tion that would meet the transparency requirement at the time when Directive 95/46 was in force. Articles 12 and 7 of the GDPR, just as Article 12 of Directive 95/46, further put obligations on data controllers to provide information to the data subjects about processing in an intelligible form, which as the CJEU states is a form which allows [them] to become aware of those data and to check that they are accurate and processed in compliance with that directive, so that [they] may, where relevant, exercise [their] rights [18]. This statement is particularly related to data subjects' right to obtain information on what data is being processed about them, and then right to request an update in case it is inaccurate. In another case [19], CJEU refers to specific rights which data subjects should be able to exercise in line with the right to access data concerning them. The Court stated that the data subject has a right to have the data communicated to him in an intelligible form, so that he is able, to exercise his rights to rectification, erasure and blocking of the data. In the GDPR, Articles 13 and 14 seem complementary to these statements and may give a clue on what an intelligible form is since types of information to be delivered by data controllers to data subjects are listed. However, none of the listed information orders data controllers to ensure understandability of the information they present.

The requirement that information be intelligible means that it should be understood by an average member of the intended audience. Intelligibility is closely linked to the requirement to use clear and plain language. An accountable data controller will have

knowledge about the people they collect information about and it can use this knowledge to determine what that audience would likely understand (calculated intelligibility). For example, a controller collecting the personal data of working professionals can assume its audience has a higher level of understanding than a controller that obtains the personal data of children. If controllers are uncertain about the level of intelligibility and transparency of the information and effectiveness of user interfaces/notices/policies etc., they can test these, for example, through mechanisms such as user panels, readability testing, formal and informal interactions and dialogue with industry groups, consumer advocacy groups and regulatory bodies, where appropriate.

If intelligible form would mean to ensure data subjects understanding of the technology, we then turn to the difficulty of fully understanding the technology which is already complex in and of itself. [20] The famous black box algorithms may prevent even data controllers from understanding what the algorithm is exactly is doing with the personal data and how it evaluates it. The AI system may receive so many personal data that it may cause fundamental changes in the way of the algorithms decision-making system which is not predictable by its creators and in this case, data controllers are somehow bound with explaining something that they do not even know technically. Not surprisingly, this is the very nature of the AI, and it is "not a bug" [21]. Which personal data, from what source, and in what way it was considered by an algorithm is still a question for many researchers. Research on creating explainable (transparent) AI and accountable algorithms [2]are on-going, but until finding a universal solution, data controllers may make up stories [22] and feed them to data subjects who cannot verify any of the information they provide. A study measuring Android apps behaviors and their potential non-compliance level with the companys own privacy statement reveals that almost half of the studied apps were found potentially inconsistent with the policy they presented and only a small portion of the apps were found completely consistent with it[23].

The updated guidelines of the Article 29 Working Party on transparency [25] actually give a clue about preparing information tailored to different audiences, so that the information could be understandable by each. According to that, data controllers first should identify the audience, including the factor or age, especially minority, then present the information. In connection with that, intelligible information means that it should be understood by the average of the target groups as assessed by the data controller, not by each of them or not by all of them. This statement remains vague, if the service to be offered is a personalized one developed based on an algorithm learning from personal data. If the condition is to first evaluate the groups based on criteria such as age, there could be quite big differences between the understanding level of people even within the same group. (However, recent experience might show that younger people understand specific terminology much better than older ones do.) The document also suggests that the level of intelligibility, not the level of users' understanding, could be tested with several methods which still may not ensure every single data subject's personal characteristics. This explanation, in our view, should further be revised in line with the characteristics of the specific AI services.

[2]Interestingly, accountability has never before been an issue in technological, only in legal terms in light of institutions, decision-makers. It, however, strongly applies to algorithms as decision-makers, without the AI being actually qualified as a person in a legal sense. However, the EU introduced the idea of giving them an electronic personality, and the scientific community has already started assigning principles to AI systems that have so far only been used or applied to persons.

Since the data controllers must present information also about the existence of the right to be forgotten (RTBF), we further evaluate it to prove our statement regarding the impracticability of the transparency rule in the frame of the current GDPR rules.

2.2. Withdrawing Consent

Microsoft once stated that computers are very good at remembering things. Absent a system failure, computers never forget [2]. According to this statement, by asking AI developers to delete data, as Article 17 of the GDPR orders, the law sets up these systems for failure. Article 17 of the GDPR defines the conditions of the right to erasure, for example, in case the data is being processed outside of the scope of initially indicated purposes. If there is no other legitimate basis available for the data controller to continue data processing, then data subjects erasure request must be fulfilled.

Why would anyone ever want a system to forget something related to them? Considering the AI systems, data subjects might not like to consider themselves as part of a particular group decided on by an algorithm or might not want to disclose the entirety of their private choices to others. As a result, they might want to remove themselves from the algorithmic model which evolves dynamically as long as new data is being contributed, fed into it. However, both technically and legally, exercising RTBF is especially hard in todays data-driven AI systems. RTBF has evolved in its legal perspective in a way that data subjects may request search engines to hide information related to their past which is no longer public interest information. It was interpreted as right not to be found or right not to be seen in different jurisdictions, such as Italy, because being fully forgotten is technically not possible and promising such a result could mislead the data subjects as well as the courts [26]. The Italian interpretation as well as the Google v. Spain [27] rationale, from where GDPRs RTBF originates, in fact, state repetitively that balancing this right against other fundamental rights such as right to obtain information, or freedom of expression is based mostly on public interest. If a court believes that a particular information related to a data subject is in the area of public information, RTBF cannot be exercised (conditionality).

Practically, the decision to remove or not remove a particular data (from a search query) is left first up to the data controllers who created the data as part of their business investment or use the AI system to produce something new. In these cases, most often, they might find that deleting data is against their intellectual property rights (as proprietary information) and business interests. They may refuse to delete such data based on Article 17 of the GDPR which paves the way for data controllers not to delete, only remove data from all publicly available sources, because such data might be useful for future purposes such as law enforcement [28] or developing new business solutions. Removing data from a database (which is technologically more feasible than data erasure) and data erasure are not the same. In their analysis, Villaronga Kieseberg and Li [29] state that the current law appears to treat human and machine memory alike but they obviously are not the same. Ability to forget something is exclusively human (or relevant to any other biologically existing creatures), but the question whether a robot could actually forget something is not a philosophical or biological one but a technical one. Starting from this point, the above-mentioned paper, proves that it is technically impossible to delete data from databases since each data added to them is stored in various points throughout a network of databases (in real life). Logs and backups are inseparable parts of a sys-

tem monitoring such databases and as insurance, copies of almost each data pinned to a system file are a necessity for the correct functioning of a database. Erasure, in this case, appears as either going around those points where the intended data for erasure is stored and a query is made upon, and deleting something from logs or backups is almost impossible (also not recommended). Moreover, there is no real erasure, data is only earmarked as removed, without actually being erased. In practice, in these cases, the query will no longer look for that specific data in the search index because of this earmark. In AI systems where each data is evaluated and fed back into the Machine Learning (ML) process as training data, we think that it is not possible to pull a single data from such a structure. Authors come to such a conclusion that the GDPR may affect the development of ML techniques in Europe because of legal obligations misfit to technological realities.

By assuming the existence of AI systems ability to forget data, the other side of the coin makes us assume that they may have an ability to remember. Carlini et. al.[30] admits that data memorization starts at the early phase of training the Neural Networks and it is hardly an unavoidable common issue even if it is unintended (meaning where the data is not useful for learning task and neither for the accuracy of the model). The researchers prove that there are solutions for avoiding the unintended memorization, although they reach such conclusion by testing white-box algorithms at a small cost. Applying differential privacy methods may make the algorithm learn slower than standard and may cause utility loss. We are not sure whether such losses would also affect the costs for AI development, but if so, these safeguards should be implemented in the development phase anyhow. The GDPR is quite well prepared in this sense, since Data Protection by Design and by Default (Art. 25) was inserted in the legal text. However, as the European Data Protection Supervisor also refers to it in its preliminary opinion[31], the wording of this article does not include the developers, only the data controllers who may not necessarily get involved with the early phase of the AI systems development. Moreover, general recommendations made for specific technologies like AI may not always cover all possible scenarios, so more specific recommendations and explanations should be delivered by the EDPS regarding implementation of the GDPRs Article 25 on AI technologies.

3. Conclusion

In this paper, we focused on the practicability of the consent rule in connection with many rights and principles regarding data protection, but especially with the principle of transparency and RTBF. Data controllers whose obligation is to deliver transparent information to the data subjects may find themselves in a difficult position to present universally understandable (intelligible) information regarding their AI systems. The GDPR, however, does not order data controllers to verify the understandability of the information they present to data subjects. Consequently, some data controllers may use the transparency rule to trick data subjects, some of them may undertake unrealistic commitments not matching their actual capabilities and practices. Consent is also interrelated with RTBF, as the lack of consent obliges data controllers to make the AI system forget upon the request of data subject. However, as we illustrated, AI systems may not be able to forget in a way as the GDPR intends RTBF to function. Technical and practical problems implementing the GDPR on AI systems together with legal uncertainties may con-

fuse data controllers' mind, therefore more detailed explanations should be delivered by the EU regarding implementation of the GDPR on certain technologies, such as AI. As we presented in the example of applying differential privacy methods ensuring GDPR's principles at some level could be one path to the solution of some of the issues, but the effect of this would be to decrease the efficiency of the AI system. Finally, we suggest that the EU Institutions may put more weight on assessing the streamlined and practical applicability of the GDPR to the AI systems.

Acknowledgments

This research was supported by the project nr. EFOP-3.6.2-16-2017-00007, titled Aspects on the development of intelligent, sustainable and inclusive society: social, technological, innovation networks in employment and digital economy. The project has been supported by the European Union, co-financed by the European Social Fund and the budget of Hungary.

References

[1] Reinsel, D., Gantz, J., Rydning, J., Data Age 2025: The Digitization of the World From Edge to Core, *IDC* (2018).
[2] *The Future Computed: Artificial Intelligence and its role in society* Microsoft Corporation Redmond, Washington, 2018.
[3] M. Lewis, D. Yarats, Y. N. Dauphin, D. Parikh, D. Batra, Deal or no deal? Training AI bots to negotiate, *Facebook Code,* (2017) https://code.fb.com/ml-applications/deal-or-no-deal-training-ai-bots-to-negotiate/.
[4] D. Kamarinou, C. Millard, J. Singh, Machine Learning with Personal Data *Queen Mary School of Law Legal Studies Research Paper,* (2016), 247, https://papers.ssrn.com/sol3/papers.cfm?abstractid142865811.
[5] M. Taddy, The Technological Elements of Artificial Intelligence, in *The Economics of Artificial Intelligence: An Agenda* , A. K. Agrawal, J. Gans, and A. Goldfarb, eds, University of Chicago Press, 2018 http://www.nber.org/chapters/c14021.
[6] COM (2018) 795 final.
[7] European Commission High-Level Expert Group on Artificial Intelligence, A definition of AI Main capabilities and scientific disciplines, (2019).
[8] K.P.L. Coopamootoo, T. Gro, Why Privacy is All but Forgotten: An Empirical Study of Privacy Sharing Attitude, *Proceedings on Privacy Enhancing Technologies,* 2017 **4,** 3960.
[9] Le Metayer, S. Monteleone, Automated consent through privacy agents: Legal requirements and technical architecture, *Computer LawSecurity Review* (2009) 136-144.
[10] Article 29 Data Protection Working Party Opinion 15/2011 on the definition of consent.
[11] C. Mulligan, Revenge Against Robots, *S.C.L Rev.,* **69** (3), (2018), 579-596.
[12] Wachter, B. Mittelstadt, L. Floridi, Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation, *International Data Privacy Law,* **7** (2), (2017), 76-99.
[13] A. D. Selbst, J. Powles, Meaningful information and the right to explanation *International Data Privacy Law,* **7**(4), (2017), 233242.
[14] N. Laoutaris, Data Transparency: Concerns and Prospects, *Proceedinds of the IEEE,* **106** (11) (2018), 1867 - 1871.
[15] T. Kim, P. Hinds, Who Should I Blame? Effects of Autonomy and Transparency on Attributions in Human-Robot Interaction, *The 15th IEEE International Symposium on Robot and Human Interactive Communication (ROMAN 2006),* 2007, 80-85.
[16] AI Now Report, 2018, https://ainowinstitute.org/AINow2018Report.pdf.

[17] J. Tang, A. Korolova, X. Bai, X. Wang, X. Wang, Privacy Loss in Apples Implementation of Differential Privacy on MacOS 10.12, (2017), arxiv preprint:1709.02753v2.
[18] Joined Cases C141/12 and C372/12 YS (C141/12) v Minister voor Immigratie, Integratie en Asiel, Minister voor Immigratie, Integratie en Asiel (C372/12) [2014] Judgement of the Court, ECLI:EU:C:2014:2081
[19] Case C486/12, X [2013], Judgement of the Court ECLI:EU:C:2013:836
[20] M. Karyda, S. Gritzalis, J. H. Park, S. Kokolakis, Privacy and fair information practices in ubiquitous environments: Research challenges and future directions, *Internet Research,* **19**(2), (2009), 194-208, doi.org/10.1108/10662240910952346.
[21] J. Millar, I. Kerr, Delegation, relinquishment, and responsibility: The prospect of expert robots, in *Robot Law,* R. Calo, A. Froomkin, I. Kerr, eds, 102-127 Cheltenham: Edward Elgar, 2016.
[22] D. Monroe, AI, Explain Yourself, *Communications of the ACM*, **61**(11), (2018), 11-13 doi:10.1145/3276742.
[23] S. Zimmeck, Z. Wang, L. Zou, R. Iyengar, B. Liu, F. Schaub, S. Wilson, N. Sadeh, S. M. Bellovin, J. Reidenberg, Automated Analysis of Privacy Requirements for Mobile Apps, in *NDSS 2017*, Reston, VA, Internet Society, 2017, dx.doi.org/10.14722/ndss.2017.23034.
[24] Veale M., Binns R., Edwards L., Algorithms that remember: model inversion attacks and data protection law, *Phil. Trans. R. Soc.,* 2018 dx.doi.org/10.1098/rsta.2018.0083
[25] Article 29 Working Party Guidelines on transparency under Regulation, 2016/679, (2018)
[26] G. Tiberi, The right to be forgotten as the right to remove inconvenient journalism? An Italian perspective on the balancing between the right to be forgotten and the freedom of expression, *e-conference on the Right to be Forgotten in Europe and Beyond,*, Blogdroiteuropeen, 2017, 49-61. http://wp.me/p6OBGR-27k.
[27] Case C131/12, Google Spain SL and Google Inc. v Agencia Espaola de Proteccin de Datos (AEPD) and Mario Costeja Gonzlez, Judgement of the Court, [2014], ECLI:EU:C:2014:317
[28] P. Van Cleynenbreugel, What Should be Forgotten? Time to Make Sense of Article 17 GDPR From the Point of View of Data Controllers, *e-conference on the Right to be Forgotten in Europe and Beyond,* Blogdroiteuropeen, 2017, 93-94, http://wp.me/p6OBGR-27k.
[29] E.F. Villaronga, P. Kieseberg, T. Li, Humans forget, machines remember: Artificial intelligence and the Right to Be Forgotten, *Computer Law Security Review,* **34**, (2018), 304313.
[30] N. Carlini, C. Liu, J. Kos, . Erlingsson, D. Song, The Secret Sharer: Measuring Unintended Neural Network Memorization Extracting Secrets, (2019), arXiv:1802.08232v2.
[31] EDPS Opinion 5/2018, Preliminary Opinion on privacy by design, 2018.