Post Hoc Interventions and the General Data Protection Regulation

Martin L. Jönsson and Jonas Ledendal

In Gunnemyr, Mattias & Jönsson, Martin L. (2023) *Post Hoc Interventions: Prospects and Problems.* Lund: Department of Philosophy, Lund University. https://doi.org/10.37852/oblu.184

ISBN: 978-91-89415-60-7 (print)
978-91-89415-61-4 (digital – pdf)
978-91-89415-62-1 (digital – html)

DOI: https://doi.org/10.37852/oblu.184.c506

Published by the Department of Philosophy, Lund University.
Edited by: Mattias Gunnemyr and Martin L. Jönsson
Cover layout by Cecilia von Arnold, Pufendorf Institute for Advanced Studies

6

Post Hoc Interventions and the General Data Protection Regulation

Martin L. Jönsson and Jonas Ledendal[1]

Abstract. Post hoc interventions rely on having access to certain personal data – such as the gender, age, ethnicity, and sexual orientation of the persons being evaluated – in order to detect and correct for prejudice. This brings these interventions into possible tension with pertinent data protection legislation, which might restrict the processing of said data. We discuss the compatibility of post hoc interventions, more specifically the Generalized Informed Interval Scale Update (GIIU), and the General Data Protection Regulation (GDPR). In particular, we investigate the legality of applying GIIU to datasets which haven't been collected with consent from the data subjects that their data is to be processed by GIIU. We conclude that many such applications are in compliance with the GDPR, but others, specifically those where the processing includes special categories of personal data that is considered sensitive, might not be.

1. Introduction

Post hoc interventions (Jönsson and Sjödahl 2017; Jönsson 2022; Jönsson and Bergman 2022; Bergman and Jönsson in preparation) embody the idea that prejudiced evaluations (competence scores, grades, performance reviews etc.) can sometimes be made more accurate after they have been produced. The most worked out such intervention, GIIU (Generalized Informed Interval Scale

[1] Martin L. Jönsson, Senior Lecturer in Theoretical Philosophy, Department of Philosophy, Lund University. Jonas Ledendal, Senior Lecturer in Business Law, School of Economics and Management, Lund University.

Update), relies on statistically identifying patterns of prejudiced (quantitative) evaluations in the history of evaluations of a particular evaluator, and then correcting for these patterns in future evaluations produced by the same evaluator.

Post hoc interventions rely on having access to certain personal data – such as the gender, age, ethnicity, and sexual orientation of the persons being evaluated – in order to detect and correct for prejudice. This brings these interventions into possible tension with pertinent data protection legislation, which might restrict the processing of said data. The following article is concerned with investigating this tension in the context of the European union, by investigating the compatibility of GIIU and the General Data Protection Regulation (GDPR).[2]

To illustrate the tension and to make the discussion below more vivid, the article will discuss the legislation in conjunction with two fictitious cases, corresponding to two types of situations that GIIU was designed to handle.

In the first case, imagine an upper secondary school math teacher – Matt – who consistently awards significantly lower grades to female students than what is to be expected from the national average for these students.[3] And imagine further that there is no reason to believe that the students in Matt's class are not representative of the populations to which they belong.

In the second case, consider a recruiter for a private care unit – Phyllis – who evaluates black applicants for positions as physicians at a significantly lower level than her fellow recruiters.[4] And imagine that there is no reason to believe that the applicants handled by Phyllis should stand out from the norm in the way they do.

---

[2] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ L 119, 4.5.2016, p. 1–88.

[3] For instance, we might be in a situation where we know that there is little difference between boys and girls on national pseudonymized math tests. Cf. the methodology used by the Swedish National Agency for Education, e.g. Skolverket (2019; 2020).

[4] GIIU essentially corrects for deviations from expectations based on population means. Since these are seldomly directly available they must be estimated, and this can be done in different ways. This is illustrated by our two cases. In the first one we use independently obtained information about the national averages for the students (see previous footnote), and in the second we use the evaluations of Phyllis' colleagues to estimate what non-prejudiced assessment looks like (cf. Jönsson 2022: fn. 12).

An advocate of GIIU might recommend that Matt's and Phyllis's future evaluations be modified in order to compensate for the detected incongruities. This can either be done automatically or by way of a recommendation of a decision support system. Either way, such a procedure requires that we know – in Matt's case – the gender of the students that Matt has evaluated in the past, and – in Phyllis's case – the skin color of the applicants that Phyllis has evaluated in the past. If we know this, we can calculate the average score members of the relevant social groups have received by Matt and Phyllis, and thus measure the size of the prejudice we are looking to correct. This presupposes, of course, that we can legally process the required data which is dependent on pertinent data protection legislation.

2. The General Data Protection Regulation

The General Data Protection Regulation is a regulation in EU law on data protection which was adopted in April 2016, and which has been directly applicable in all member states since May 25th 2018. The EU data protection framework does, however, also to a large degree rely on legal acts in the form of guidelines form the European Data Protection Board (EDPB), formerly the Article 29 Working Party. Such guidelines are adopted by the board under the GDPR to ensure consistent application and interpretation of the regulation.[5] Although, non-binding EDPB guidelines have a high de facto impact on how GDPR is applied by supervisory authorities and courts.[6]

The primary aim of the regulation is to protect the fundamental rights of individuals (data subjects) with regard to the processing of their personal data, mainly by making the processing more transparent and enhancing an individual's control over his or her personal data. The regulation also has a second aim of safeguarding the free movement of personal data within the union by ensuring that data protection legislation is uniform.[7] The GDPR lists

[5] See Article 70 of Regulation (EU) 2016/679 (GDPR). The consistent application is also ensured by the consistency mechanism (Article 63 of GDPR), which enables the EDPB to resolve disputes between national data protection supervisory authorities. The decision of the board is binding on the member states.

[6] Article 288 of the Treaty on the Functioning of the European Union (OJ C 326, 26.10.2012, p. 47–390). See also Craig & de Búrca 2020 on how the admixture of formal and informal law is a common feature of the legal order but can nonetheless give rise to problems.

[7] Article 1 of Regulation (EU) 2016/679 (GDPR).

detailed and fairly restrictive rules for how to process personal data. Although it is formally applicable only to a restricted region of the world, it has since its adoption become a model for similar legislation in many other parts of the world as well (Bradford 2021).

A noteworthy aspect of the regulation is its preventive nature. A person processing personal data is responsible in various ways (described below) for how the data is processed. It is, however, not enough that the person responsible implements safeguards to manage risks arising from its own processing but must also account for risks related to how the data can be used by others, such as potential malicious actors.

3. The Processing of Personal Data Required by Post Hoc Interventions

The GDPR lays down rules relating to the protection of natural persons (i.e., humans) with regard to the processing of personal data (Art. 1, GDPR). A first natural question then – to determine the applicability of the GDPR to post hoc interventions like GIIU – is to ask whether post hoc interventions involve (1) *the processing* of (2) *personal data*? These questions can in turn be fruitfully subdivided into five sub-questions, each relating to one of the following steps or processing operations required by GIIU:[8]

Collection – The step in which an evaluator passes an evaluative judgment concerning someone. For instance, the math teacher Matt, deciding to give one of his students, Molly, a particular grade in math.[9]

Recording – The step in which the evaluator records his judgment. For instance, Matt entering the grade he has decided on into a grade reporting system on his computer.

[8] Although the exact division of a process into steps, and the granularity of such a division, can be important from the perspective of the GDPR, we don't see a need for a more fine-grained division in the present context.

[9] The concept of data "collection" in data protection law is not limited to the act of obtaining data through literal collection, but also encompass passively receiving data and creating data. Hence, setting a grade for an assignment or passing a judgement is equated with collection.

Storing – The retention of the recorded judgment so that a history of evaluation (i.e. a set of such judgements which is big enough for statistical analysis) can be constructed. For instance, the retention of Molly's and her peers' grades on a server.

Analysis – The statistical analysis of differences in means between members of different social groups. For instance, the comparison of the means for the female students and the male students Matt has graded from the perspective of a particular assumption about how these means relate on the population level.

Modification – The potential modification of newly passed evaluative judgments in light of found incongruities in the history of evaluations. For instance, increasing the grades of Mikes newly evaluated female students in light of past female student having received biased grades.[10]

So in order to determine whether the GDPR is applicable to post hoc interventions we need to ask, for each of these steps, whether it requires (1) *the processing* of (2) *personal data*.

Processing is quite broadly defined as follows:

> 'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction;
>
> (Article 4(2) of GDPR)

From this it is clear that each of the five steps involves processing, at least to the extent that they involve personal data: steps 1, 2 and 3 are all explicitly mentioned in the definition, step 4 involves retrieval and use, which is mentioned in the definition, and step 5 involves alteration which is also mentioned in the definition. In addition, it is clear from the language of the definition ("such as") that this list is non-exhaustive and intended to be illustrative.

[10] If GIIU is used as a decision support system, this step does not involve actually modifying data, but only suggesting to the evaluator that data should be modified. This is an important difference in the present context. See Section 6.

Personal data is also broadly defined in the following way:

> 'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person;
>
> (Article 4(1) of GDPR)

Hence, personal data is any data that both "relates to" a natural person and makes it possible to identify him or her.[11] In the context of post hoc interventions, data "relates to" humans in the sense that they are statements about humans (e.g., their school or work performance). However, since non-identifiable data is outside the scope of the GDPR, it is possible to anonymise personal data to make further processing steps compliant with the regulation.[12] This requires an appropriate anonymisation method which makes the risk of re-identification practically impossible or at least insignificant due to that it would require a disproportionate effort in terms of time, cost and man-power.[13]

The pieces of data that post hoc interventions process have the following relational form (illustrated by our first example):

> (M1) Matt Berry has given Molly Sinclair the grade 3.

[11] See further on the concept of personal data Article 29 Working Party, Opinion 4/2007 on the concept of personal data, WP136, adopted on 20 June 2007. See also Judgment of the Court of Justice of the European Union of 20 December 2017 in Case C-434/16, Nowak (ECLI:EU:C:2017:994).

[12] It is here worth noting that data pseudonymisation is not the same as anonymisation (see Article 4(5) of GDPR, which defines "pseudonymisation" as "means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person"). Whereas anonymous data is non-personal data, pseudonymised data remains personal data and must comply with GDPR.

[13] Judgment of the Court of Justice of the European Union of 19 October 2016 in Case C-582/14, Breyer (ECLI:EU:C:2016:779), para. 46. See also Article 29 Working Party, Opinion 5/2014 on Anonymisation Techniques, WP216, adopted on 10 April 2014. The European Data Protection Board is developing new guidelines, which have not been made public at the time of writing.

M1 is clearly a piece of personal data on account of it featuring two names, which makes their bearers easily identifiable. Step 1 ('collecting') in particular thus seems inescapable to involve the processing of personal data since there is no way to anonymize at this point – particular students need to be assigned their grades – and is thus subject to the GDPR. The recording and storage of pieces of data like M1 typically involves the recording and storage of them in an unaltered state and this would mean that Steps 2 (Recording) and 3 (Storing) would also be subject to the GDPR. It should be noted though that this is not needed from the perspective of Step 4 (Analysis). In particular, what is needed from the perspective of this step (and thus what this step needs from steps 2 and 3) is instead something like the following:

(M2) Matt Berry has given a female student the grade 3.

The application of post hoc interventions does not require us to retain any identifiers relating to the people being evaluated in the past. However, M2 would still count as personal data from the perspective of the previous definition since it features Matt's name. M1 featured identifiers for two data subjects and one still remains in M2. Moreover, since it is Matt's future evaluations that we are looking to update, it seems that we must retain his identifier. It seems highly unlikely that any anonymisation technique can be applied to the data that would break the link to the evaluator in a useful way that would make him or her unidentifiable in the manner required by the GDPR.[14] We can thus conclude that each of the aforementioned five steps involve the processing of at least some personal data, as defined by the GDPR.

4. Controllers, Processors and Territorial Scope

In order to determine the applicability of the GDPR to any particular processing of personal data, certain roles described by the GDPR must be identified, in particular a controller – "the natural or legal person, public authority, agency or other body which, alone or jointly with others,

[14] This is so because GDPR is not limited to data that is directly identifiable, i.e., data that is contained in the same dataset or otherwise held by the data controller. It is enough that the data subject is indirectly identifiable, e.g., by combining the data with other data which might or might not be held by the controller. See Judgment of the Court of Justice of the European Union of 19 October 2016 in Case C-582/14, Breyer (ECLI:EU:C:2016:779).

determines the purposes and means of the processing of personal data" (Article 4(7) of GDPR, our emphasis) – and one or more processors – "a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller (Article 4(8) of GDPR).[15] The processing must also fall within the territorial scope of the regulation. Whether the GDPR is applicable to the relevant processing is hence a matter of whether it is carried out "…in the context of the activities of an establishment of a controller or a processor in the Union…". (Article 3(1) of GDPR) It is not significant whether that processing actually takes place in the union (Ibid.).[16]

Usually in cases like the first example, the controller is the public authority or similar entity in charge of the school, and in the second the employer. Although natural persons can also be controllers, the evaluators, i.e. Mike and Phyllis, who only access and process the personal data under the authority of their respective employers, are not considered controllers. Like other employees they are not directly responsible for the processing, but can themselves be data subjects, since their personal data is also processed during the intervention. The same is true for the post hoc intervener – the person administering GIIU – which is involved in the last two processing steps. Given that the school and the company are located within the European Union – which we will stipulate – the GDPR is applicable.

Another important role assignment in what follows is that of the data subject.[17] As was mentioned above, M1 featured identifiers for two different data subjects: the evaluator – which we will refer to using "data subject^{EV}" – and the evaluated – which we will refer to using "data subject^{ev}". This will become important in Section 5.

[15] See further European Data Protection Board, Guidelines 7/2020 on the concepts of controller and processor in the GDPR, adopted on 7 July 2021. Joint controllership is also possible when the purposes and means of the processing have been jointly determined by two or more controllers (Article 26 of GDPR).

[16] See further European Data Protection Board, Guidelines 3/2018 on the territorial scope of the GDPR (Article 3), adopted on 12 November 2019.

[17] The data subject is always a living human (natural person) and is protected regardless of nationality or residence. Deceased persons or legal persons are not protected (Recitals 14 and 27 of GDPR). See further on the concept of data subject Article 29 Working Party, Opinion 4/2007 on the concept of personal data, WP136, adopted on 20 June 2007.

5. Principles of Personal Data Processing

Now that we have determined that the GDPR is applicable to our two cases, we need to determine whether the corresponding post hoc interventions can be carried out in compliance with the principles of personal data processing stipulated by the GDPR. These are as follows:

> "Personal data shall be:
>
> a) processed lawfully, fairly and in a transparent manner in relation to the data subject ('lawfulness, fairness and transparency');
>
> b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; … ('purpose limitation')
>
> c) adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation');
>
> d) accurate and, where necessary, kept up to date; … ('accuracy');
>
> e) kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; … ('storage limitation');
>
> f) processed in a manner that ensures appropriate security of the personal data, … ('integrity and confidentiality')."
>
> (Article 5(1) of GDPR)

Of these principles, the last two are of little importance when it comes to whether post hoc interventions can be in compliance with the GDPR and they will not be discussed further.[18]

5.1 Lawfulness and Purpose Limitation

The first point of tension between GIIU and the GDPR comes from the fact that post hoc interventions are fairly novel (first described by Jönsson and Sjödahl 2017). This means that there are few, if any, data sets ('histories of evaluations') that have been collected with the express purpose of applying post hoc interventions to them. The legality of applying GIIU to extant datasets

[18] The controller would have to ensure that these requirements are fulfilled, but in our opinion post hoc intervention does not pose any burden, legal uncertainties or complications that would go beyond what is required for any other processing of personal data.

without collecting consent is thus important to determine, because being able to do so could increase GIIU's scope of applicability. This depends on at least two different conditions.

First, it can be noted that the concept of lawfulness in Article 5(1) lit. a of the GDPR, also known as the requirement of legal basis, is expounded in a later article as follows:

> "Processing shall be lawful only if and to the extent that at least one of the following applies:
>
> a) the data subject has given consent to the processing of his or her personal data for one or more specific purposes;
>
> b) processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract;
>
> c) processing is necessary for compliance with a legal obligation to which the controller is subject;
>
> d) processing is necessary in order to protect the vital interests of the data subject or of another natural person;
>
> e) processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller;
>
> f) 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."
>
> (Article 6(1) of GDPR)

This makes it clear that in cases where the data subject has not given its consent to the processing required by the last two steps of post hoc interventions (analysis and modification), lit. f (legitimate interest) can be used as the legal basis in the relevant processing context. Thus, the lawfulness of the analysis and modification steps depends on whether that 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" (ibid.)

Second, lit. b of Article 5(1) stipulates that processing that goes beyond the purposes for which the data was originally collected must not be incompatible

with those purposes. It is worth noting that pursuant to lit. b such purposes must also have been explicitly specified by the controller at the point of collection.

Given the aforementioned, we must thus determine (1) whether post hoc analysis and modification is in the legitimate interests of the controller and not overridden by the interests of the data subject, as well as (2) whether this processing is compatible with the purposes for which the personal data was originally collected.

With respect to the first question, the answer seems to be affirmative for the two cases we are working with. The purpose for which the data was collected is likely something that can be paraphrased as "attempting to accurately measure a student's math proficiency" and "accurately determine an applicant's competence as a physician" respectively. It is difficult to see then how attempts to increase the accuracy of the relevant evaluations could be illegitimate (or in conflict with "interests or fundamental rights and freedoms of the data subject").

With respect to the second question, the answer also seems to be affirmative. Since the aim of the intervener is to increase accuracy, it appears to be an aim very much in line with the original purpose (as paraphrased in the preceding paragraph). Article in 6(4) of the GDPR seems to support this conclusion.

> "Where the processing for a purpose other than that for which the personal data have been collected is not based on the data subject's consent or on a Union or Member State law which constitutes a necessary and proportionate measure in a democratic society to safeguard the objectives referred to in Article 23(1), the controller shall, in order to ascertain whether processing for another purpose is compatible with the purpose for which the personal data are initially collected, take into account, inter alia:
>
> a) any link between the purposes for which the personal data have been collected and the purposes of the intended further processing;
>
> b) the context in which the personal data have been collected, in particular regarding the relationship between data subjects and the controller;
>
> c) the nature of the personal data, in particular whether special categories of personal data are processed, pursuant to Article 9, or whether personal data related to criminal convictions and offences are processed, pursuant to Article 10;

d) the possible consequences of the intended further processing for data subjects;

e) the existence of appropriate safeguards, which may include encryption or pseudonymisation."

Let's go through these in order. There is a clear link between the purposes for which the data was initially collected and the purpose of the further processing (cf. lit. a). The relevant context is one of processors attempting to give accurate measures on behalf of the controller (cf. lit. b). Hence, the further processing would meet the reasonable expectations of the data subject (Recital 50 of GDPR). [19]

The personal data that is being processed might belong to special categories and might relate to criminal convictions and offences but we will discuss such cases separately below so we will assume for now that it doesn't (cf. lit. c) as is the case in our first example.

The consequences for the data subject^{ev} in post hoc *analysis* is likely negligible, but they might be quite dramatic in post hoc *modification*. Still, the consequences are not more severe than they were in the original processing (one might get a poor grade in math, or be graded as a poor physician). Post hoc *analysis* might, however, have real consequences for the data subject^{EV} since it might reveal him or her to be biased with the possible effect of stigmatization if this becomes known. It should be noted though that such consequences are already possible for the data subject^{EV} if their evaluation is overtly biased. Still, careful analysis makes detection of bias more likely. It thus becomes important that the result of the analysis is safeguarded appropriately (e.g. with encryption). The consequence of post hoc *modification* for the data subject^{EV} might be either that they are informed that a bias has been detected and it is suggested to them they update, or that their evaluation is overridden (Jönsson forthcoming). Both might of course cause the data subject^{EV} concern, but since GIIU attempts to help the data subject^{EV} make more accurate decisions, this concern shouldn't lead us to think that post hoc

[19] At least for post hoc analysis, it seems plausible to say that the kind of statistical analysis carried out there falls within the reasonable expectations of an average data subject (both Mike and his students in the first example for instance). Post hoc modification is more questionable since the kind of modification being carried out there is likely not expected by most data subjects. However, if the modification is not automated but merely suggested to the evaluator, it seems more plausible that it falls within the subjects' reasonable expectations.

modification goes against the purpose for which the data was originally collected. (cf. lit. d)

For the purposes of post hoc analysis anonymization or pseudonymization is possible with respect to the data subject who is being evaluated (i.e. moving from M1 to M2), although this is not possible from the perspective of automatic post hoc modification (we have to know who should be updated). For both kinds of processing, we can use encryption in the processing. (cf. lit. e).

Jointly, these five considerations seem to point towards the further processing being compatible with the purposes of the original processing. Both the legitimacy and the compatibility of the processing can, however, depend on appropriate safeguards to ensure that the personal data is not misused for other purposes. For example, as have been mentioned above, that new knowledge about data subject^{EV} is not used by the employer to evaluate him or her, since such further processing would go beyond the original purpose and would require a separate analysis to determine whether it is lawful under the GDPR.

It thus seems to us that applying post hoc analysis and modification to the personal data like that in our two cases can be compatible with the GDPR even if the data has not been collected explicitly for that purpose.

5.2 The Risk of Inducing Error

According to the principle of accuracy, personal data shall be ”accurate and, where necessary, kept up to date” (Article 5(1) lit. d of GDPR). Although GIIU aims to improve accuracy, it is statistically possible, although unlikely, that it will instead decrease accuracy. This is another point of tension between GIIU and the GDPR.

The GDPR mandates in the present context, that the controller must ensure that appropriate safeguard measures are in place to minimize the risk that GIIU induce errors when personal data (e.g., grades) are modified. This is even more important if modifications are automated as the impact of batch processing might potentially affect a much larger number of data subjects (see also below Section 5.3). Safeguards could include mechanisms that detect deviations from the conditions under which GIIU will work as intended (cf. Jönsson and Bergman 2022) or other appropriate statistical measures. The transparency of the processing in relation to the data subject is also imperative, since such transparency makes it possible for him or her to review the accuracy of the processing and request that inaccurate results are rectified pursuant to Article

16 of the GDPR. The controller can, however, not solely rely on such data subject review, but must regularly perform its own audits to detect and rectify inaccurate data. Personal data is considered inaccurate when it is unfit for the purpose of the processing.[20] Hence, in relation to historical data, it must be considered whether later changes need to be reflected in the data set that is used for the post hoc intervention assessment. Provided that such safeguards are in place to detect and rectify inaccuracies both in the data that is used as the basis for the assessment and the resulting modifications, such interventions should be compatible with the GDPR.

5.3 Automated Decision-Making and Profiling

GIIU can be implemented in different ways that range from manual to fully automated processing of data. In practice, the processing is likely to use some automated processing, but, as was mentioned in the introduction, the final decision to correct the grade or evaluation could be left to the evaluator. GIIU would then only act as a decision support tool. It would, however, also be technically possible to fully automate the procedure in such a way that it would not involve any human intervention. Such processing would bring GIIU into tension with the GDPR in a further way.

This kind of processing would have to comply with Article 22 of GDPR, which contains rules on automated individual decision-making. The provision is applicable when a decision based solely on automated processing produces legal effects concerning an individual or similarly significantly affects him or her. The latter includes decisions affecting someone's employment opportunity.[21]

The interpretation of the rule is disputed, but it either generally prohibits automated decisions falling within its scope or at least gives the data subject a right to object to such processing (Drożdż 2020). It is also disputed whether the rule only covers profiling or any automated decision-making.[22] In this

[20] See Article 5(1) lit. d and Article 16 of GDPR, which both state that the accuracy of personal data should be determined having regard to the purposes for which they are processed.

[21] Article 29 Working Party, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, WP251, adopted 6 February 2018.

[22] See Article 4(4) of GDPR, which defines "profiling" as "any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural

context, it is relevant to note that GIIU does not intend to make any analysis or predictions about data subject[ev], but instead assess the potential bias of data subject[EV]. Hence, the person that is potentially being profiled is not the person affected by the automated decision (i.e., the post hoc modification). With regard to the legal uncertainties surrounding Article 22 of the GDPR, controllers should carefully assess whether a fully automated GIIU procedure is permitted and where required acquire consent from the relevant data subjects and implement appropriate safeguards (see also Section 6 on the need for Data Protection Risk Assessment).[23] Since this depends on the context the assessment has to be made on a case-by-case basis.

6. The Processing of Special Categories of Personal Data

From the above discussion we can gather that post hoc processing of personal data can be in compliance with the GDPR even if the data subjects haven't given their consent to this processing. However, this doesn't take into account the possibility of the data containing 'special categories' of personal data (also known as "sensitive personal data"). Concerning such categories the GDPR maintains the following:

> Processing of personal 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 shall be prohibited.
>
> (Article 9 (1) of GDPR)

What this means is that the processing of such sensitive personal data is prohibited unless the processing is subject to one of the exceptions stipulated

person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements".

[23] See Article 22(2) lit. c of GDPR, which stipulates that such consent must be explicit. Consent is required unless the automated decision-making is permitted by law or necessary to conclude or perform a contract between the controller or the data subject. Article 22 para. (2) and (3) also stipulates that suitable safeguards must be implemented. This includes but is not limited to the rights to require human intervention.

in point 2 of the same article. The first thing to note is that gender and age are not special categories (as described by Article 9(1)), and our first case thus remains unproblematic.[24] Our second case, however featured skin color which likely falls under "revealing racial or ethnic origin". An evaluation of the exceptions available in Article 9(2) of the GDPR shows that most are usually not relevant for our case, although some might be applicable in certain special cases. It is for example possible that some post hoc interventions could be viewed as being 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 law (cf. lit. b). For instance, in situations where pro-active work is required to avoid discrimination or improve diversity. In general, the controller would, however, have to rely on the explicit consent of the data subject (cf. lit a).

Does the post hoc processing involving sensitive personal data thus require explicit consent from the data subjects (both the evaluator and the evaluated)? To answer this we need to treat post hoc analysis and post hoc modification separately.

As we saw above, post hoc analysis only requires personal data of the following form (here illustrated with an example from our second case).

(P2) Phyllis Berry has given a black applicant the competence assessment 3.

Not the more inclusive P1

(P1) Phyliss Berry has given Donald Glover the competence assessment 3.

The name of the applicants (i.e. the evaluated) are not needed in order carry out this processing. But since skin color is a property of the data subject^{ev} (and not the data subject^{EV}, i.e. the evaluator), P2 does not fall under the requirements of Article 9. The reason for this is that although the sensitive data is part of the data set being processed it does not "relate to" the data subject, i.e., data subject^{EV} (cf. Article 4(1) of GDPR). The GDPR aims to protect the rights of the data subject, not the data as such. Hence, the data must be sensitive to the data subject to fall under Article 9 and not merely sensitive in nature, since it would otherwise not

[24] This can, however, depend on the context of the processing. See Judgment of the Court of Justice of the European Union of 1 August 2022 in Case C-184/20, Vyriausioji tarnybinės etikos komisija (ECLI:EU:C:2022:601). The Court found that name-specific data relating to someone's spouse, cohabitee or partner can reveal sexual orientation and fall under Article 9 of the GDPR even where that was not the intention of the processing.

be able to create the kind of special risks for the data subject that is the object of the prohibition in Article 9. This means that consent is still not required for post hoc analysis even if processes data that are derived from personal data relating to special categories. The exception to this would be cases where the black applicant in P2 could be identified indirectly, e.g. through being one of very few black applicants in the history of evaluations of Phyllis.

However, in the fully automated post hoc *modification* we need to know who we should update and we cannot anonymize (or even pseudonymize). There thus seem to be no escaping the need to ask for explicit consent pursuant to Article 9 of the GDPR from the evaluated persons for this step to be in compliance with the GDPR. However, if we consider the variant of GIIU which acts as an advisory decision support system, one can envision it only generating general recommendations concerning members of certain social categories to the evaluator, e.g. "It looks like your grades for female students are 1 point lower than what is to be expected". This would not require the processing of personal data concerning data subject^{ev} and thus not of sensitive personal data, and hence there is no need to ask for explicit consent.

In the above we have assumed that no sensitive data relating to data subject^{EV} is processed during the intervention. Since such data (e.g. grades) would constitute personal data (also relating to the evaluator) and cannot be anonymized, we also need to assess whether evaluative judgements can constitute sensitive data under Article 9(1) of the GDPR. At first glance, this does not seem to be the case. It should, however, be noted that the CJEU has held that Article 9 should be given a fairly wide interpretation.[25] It could for instance not be ruled out that processing of data concerning bias or discriminatory practices – given an extensive interpretation – could be considered personal data "revealing … political opinions". In our view this interpretation is too extensive since the processing merely creates an abstract risk of revealing someone's political opinions.[26] It would be another matter if this was the purpose of the processing, but that is not the aim of post hoc interventions.

---

[25] See Judgement of the Court of Justice of the European Union of 6 November 2003 in Case C-101/01, Lindqvist (ECLI:EU:C:2003:596), para. 50. See also the Judgement of 1 August 2022 in Case C-184/20, above footnote n.

[26] For a similar argument see Judgement of the German Administrative Court of Mainz of 20 February 2020 case no. 1 K 467/19.MZ, ECLI:DE:VGMAINZ:2020:0220.1K467.19.00. The court held that merely the abstract risk of the transmission of a (zoonotic) disease from an animal to a human did not mean that such data in general should be viewed as "data concerning health" relating to the animal's owner under Article 9(1) of the GDPR.

7. Risk Assessment and Data Protection Impact Assessment

The GDPR requires that controllers assess the risk for data subjects' fundamental rights prior to processing their personal data. When there is an indication that this processing is likely to result in high risks, pursuant to Article 35 the controller must conduct a formal Data Protection Impact Assessment (DPIA). This is particularly so when the processing involves new technologies. Post hoc interventions have a well-documented basis in the literature, but different implementations are still to be tested and applied in practical decision-making. Guidance on when processing is likely to result in a high risk can be found in the Article 29 Working Party's Guidelines on Data Protection Impact Assessments.[27] The guidelines have been endorsed by its successor the European Data Protection Board. Such guidance has also been adopted by the national data protection supervisory authorities. Since implementations of GIIU might involve new technologies and involve processing of both sensitive personal (e.g. ethnic origin) data and data concerning children (e.g. grades of schoolchildren), it is likely that such processing in the light of these guidelines falls within the scope of Article 35. This is so in particular when the processing consists of automated decision-making, including profiling (see above Section 5.3). Hence, organizations that intend to implement automated post hoc intervention procedures must perform a DPIA and when the assessment indicates that the processing would result in a high risk in the absence of measures taken by the controller to mitigate the risk also consult the competent supervisory authority (Article 36 of GDPR).

8. Conclusion

Even though restrictive on the face of it, the GDPR seems to be compatible with post hoc interventions being applied in cases like the first of our two examples (featuring Mike) even without the consent of the data subjects (either evaluators or the evaluated) specifically for this processing. Similar considerations apply in legally similar contexts such as government agencies.

[27] Article 29 Working Party, Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679, WP248, adopted on 4 October 2017.

Exceptions arise only if the personal data is of sensitive nature or if the processing is automated. So, for instance, if our second example features a version of GIIU where processing is fully automated, or features sensitive data, post hoc *modification* would require explicit consent, even though post hoc analysis would not.

Acknowledgments

The research was funded by a research grant from the Swedish Research Council (Dnr. 2017-02193) and research funding provided by the Pufendorf IAS in Lund. The text has benefitted from discussion with and/or careful reading by members of Post Hoc Interventions Pufendorf theme, as well as the participants at the conference, Post Hoc Interventions: Prospects and Problems, organized in Lund in October 2022.

References

Article 29 Working Party, Opinion 4/2007 on the concept of personal data, WP136, adopted on 20 June 2007.

Article 29 Working Party, Opinion 5/2014 on Anonymisation Techniques, WP216, adopted on 10 April 2014.

Article 29 Working Party, Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679, WP248, adopted on 4 October 2017.

Article 29 Working Party, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, WP251, adopted 6 February 2018.

Bergman, J. and Jönsson, M. L. (submitted) “Gender Bias in Grant Applications. Inquiry and the Potential for a Post Hoc Remedy”. Manuscript.

Bradford, A. (2021) *The Brussels Effect: How the European Union Rules the World.* New York, NY.: Oxford University Press.

Craig P. and de Búrca, G. (2020) *EU Law: Text, Cases and Materials (7th ed.).* Oxford: Oxford University Press.

Drożdż, A (2020) *Protection of Natural Persons with Regard to Automated Individual Decision-Making in the GDPR.* Alphen aan den Rijn: Kluwer Law International.

European Data Protection Board, Guidelines 3/2018 on the territorial scope of the GDPR (Article 3), adopted on 12 November 2019.

European Data Protection Board, Guidelines 7/2020 on the concepts of controller and processor in the GDPR, adopted on 7 July 2021.

Jönsson, M. L. (2022) "On the Prerequisites for Improving Prejudiced Ranking(s) with Individual and Post Hoc Interventions" *Erkenntnis*.

Jönsson, M. L. and Sjödahl, J. (2017) "Increasing the veracity of implicitly biased rankings". *Episteme 14*(4), 499–517.

Jönsson, M. L. and Bergman, J. (2022) "Improving misrepresentations amid unwavering misrepresenters", *Synthese*, *200*.

Skolverket (2019) *Analyser av likvärdig betygssättning mellan elevgrupper och skolor.*

Skolverket (2020) *Analyser av likvärdig betygssättning i gymnasieskolan.*