

An Analysis of the Consequences of the General Data Protection Regulation (GDPR) on Social Network Research

Andreas Kotsios* Matteo Magnani† Luca Rossi‡
 Irina Shklovski§ Davide Vega¶

March 11, 2019

Abstract

This article examines the principles outlined in the General Data Protection Regulation (GDPR) in the context of social network data. We provide both a practical guide to GDPR-compliant social network data processing, covering aspects such as data collection, consent, anonymization and data analysis, and a broader discussion of the problems emerging when the general principles on which the regulation is based are instantiated to this research area.

1 Introduction

In the last decade online social networks have become a major source of data to study human and social behaviour. The availability of persistent and searchable traces of human communication on a large scale [11] provided new, previously inconceivable opportunities for unobtrusive research, but also raised new questions related to the potential misuse of personal information. Following awareness-increasing events attracting worldwide attention such as the Cambridge Analytica scandal, and related restrictions on data access established by large social media companies, several Internet researchers recently highlighted the necessity and complexity of ensuring that "independent, critical research in the public interest can be conducted while protecting ordinary users' privacy"¹.

*Faculty of Law, Uppsala University

†InfoLab, Dept. of Information Technology, Uppsala University

‡IT University Copenhagen

§IT University Copenhagen

¶InfoLab, Dept. of Information Technology, Uppsala University

¹<https://medium.com/@Snurb/facebook-research-data-18662cf2cacb>

In this context, a very important event happened on the 25th of May 2018, when the General Data Protection Regulation² (GDPR) came into force as the new and improved piece of legislation regulating how natural persons should be protected with regards to the processing of their personal data. The GDPR applies to the data collection performed by private marketing companies as well as to research by private companies or public universities, in very similar ways in all Member States, all sectors (public or private) and all purposes (commercial and non-commercial), and it has been welcomed as a progressive step towards rectifying the glaring power imbalance in current mass digital data collection by entities that develop, maintain and control access to digital infrastructures. In particular, the GDPR has two main goals³ that are very similar to the ones mentioned above and invoked by many Internet researchers. The first goal is to protect the fundamental rights and freedoms of the data subjects by creating a protective regiment with regards to the processing of personal data⁴. This is because new technologies and organisational models both in the private and public sector have made it easy to gather, use, combine, aggregate or otherwise process a vast amount of personal data without sufficient controls or oversight. The second goal is to create the optimal conditions so that the free flow of personal data – in parallel to the free movement of goods and services – can take place within the EU, supporting the creation of the European Single Market.

However, the impact of the new rules on the practice of research is unclear and this is especially relevant to researchers studying social network data (both on- and offline), where for example the subjects participating in a study may provide information about non-participants and the collected data is more difficult to effectively anonymize than in other research fields. This article considers the impact and implications of the GDPR and of the research exemptions built into the law on the activities of researchers engaging in social network analysis, both in general and in the specific case of online social networks.

The regulation replaces the earlier Data Protection Directive⁵ which, as a directive, was adopted and implemented through national laws by every EU Member State, which resulted in at times confusing patchwork of national regulations. More importantly, the earlier directive did not have any specific focus on research. Rather, the main regulatory mechanisms were codes of conduct and ethical guidelines advocating good practices but rarely systematically codifying these. In contrast, the GDPR explicitly recognizes the particularities of data processing in research through a series of formally specified research exemptions, which have important consequences on the feasibility and lawfulness of social

²Regulation (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

³Rec 1-7 GDPR and SOU 2017:50, p78

⁴In the context of GDPR 'personal data' means any information relating to an identified or identifiable natural person; 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

⁵Directive 95/46/EC

network research projects in practice. This includes the ability to limit and even avoid restrictions on secondary processing and the processing of sensitive categories of data⁶, to override the subjects’ right to object to processing and erasure as long as relevant safeguards are implemented⁷, and to collect some types of data without consent for some types of processing⁸.

The text of the GDPR is complex and not specifically targeted to researchers but its content will impact research practices in significant ways that depend on the specific research field. On the one hand, universities and other research institutions are of course providing general information about the GDPR to their employees, and have a Data Protection Officer who can be contacted for specific matters. On the other hand, GDPR-compliant strategies need to be instantiated to the specific research problem, and many of these are difficult to interpret without domain-specific knowledge. Therefore, researchers should themselves be aware of the implications of the GDPR and have reflected about it. In fact, several papers have already been published to address this issue and to examine the impact of the GDPR on research in general [9, 18, 21] and on specific research fields [20]. However, none have addressed the implications of the GDPR for social network analysis, where data processing differs from other quantitative approaches. For example, as discussed by Borgatti and Molina [3], respondent anonymity is not an option if we want to know who is talking about whom, which is necessary to define edges in the network. In addition, subjects providing information about their social relations may generate data about individuals not included in the study: a participant mentioning that she often performs some activity with someone may reveal a lot about this other person depending on the type of activity. Another issue is the fact that in social networks it is often possible to identify specific roles based on the network structure, with a limited number of individuals in each role, examples being high-degree and high-betweenness⁹ nodes as well as other special network configurations. Once these few nodes have been identified, it becomes very simple to connect them to specific individuals using some basic knowledge of the studied organization. More in general, network data is often impossible to fully anonymize. In this article we will also identify specific issues related to data protection emerging when social network analysis is applied to less traditional contexts, such as the analysis of large-scale networks of social relations derived from social media data.

In the next section we present a brief overview of the GDPR, including the terminology used in the rest of the article. This short section is necessary to make this article self-contained, and it can be skipped by the reader who is already familiar with the main actors, concepts and principles introduced by the regulation. The following section is organized along the main steps and problems of a typical social network research process. We start by discussing

⁶Art 6(4) and rec 50 GDPR

⁷Art 89 GDPR

⁸Art 6(1)(e) and (f); rec 47 and 157 GDPR

⁹Degree and betweenness are so-called centrality measures, that can be used to identify important actors in a network

approaches to data collection, also highlighting the differences between data collected directly from the data subjects or indirectly, such as through social media Application Programming Interfaces (APIs). We also discuss topics such as consent, data anonymization, profiling, and storage of the networks. We conclude the article with more general considerations about the implications of the GDPR for commercial data controllers as well as for the future of network data repositories which represent important teaching and training tools. We also suggest that the GDPR should be a new important element to be considered in the ongoing discussion about the establishment of a code of conduct for social network research.

Please notice that the GDPR applies only to the processing of personal data by entities in the European Union, or when the data subjects are in the Union. Our presentation will often take the perspective of a European public university, and we will extend the discussion to other cases regulated by this law when relevant. However, the principles defined in the GDPR are worthy of consideration even for researchers outside the Union processing data from non-EU subjects, as these principles highlight general fundamental issues to be considered when processing personal data.

2 The GDPR ecosystem

Before we proceed to analyze how the GDPR will affect social network research in practice it is important to make clear the fundamental terms and ideas of this piece of law. The regulation is complex in term of length (88 pages, 173 recitals, 99 articles), breadth of coverage and depth. Here we present the GDPR concepts and principles that are the most relevant for social network analysis research. In the next section we will discuss the (often unclear) role that these concepts and principles play in the various phases of a social network analysis process.

Figure1 exemplifies the main concepts described in this section in the context of a typical academic research project. The data subject is defined as every individual (natural person) who is identified or may be identified by the controller or third parties, directly or indirectly, by the act of processing her personal data. The natural or legal persons who decide how and why the personal data will be processed are the data controllers while the ones that process the data on behalf of the controller are the data processors¹⁰.

What constitutes personal data is defined quite broadly as any information that does or may lead to the identification of a natural person¹¹. The term processing is defined similarly broadly as "any operation or set of operations on personal data or sets of personal data"¹², including data collection. A special type of processing is profiling. With an equally broad definition profiling is defined as "any form of automated processing of personal data evaluating the

¹⁰Art 4(8) and 4(9) GDPR

¹¹Art 4(1) GDPR

¹²Art 4(2) GDPR

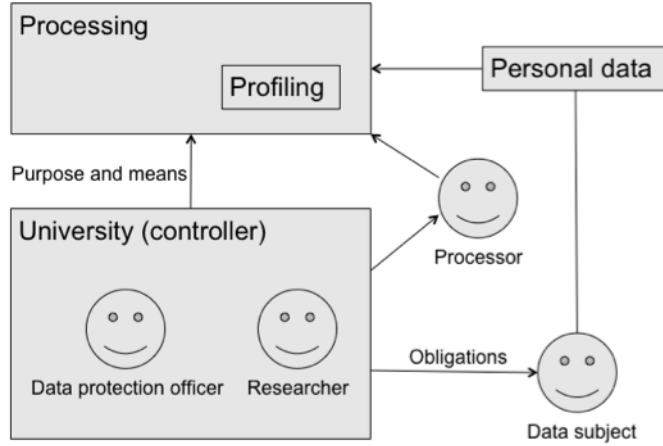


Figure 1: A typical configuration of the GDPR ecosystem.

personal aspects relating to a natural person, in particular to analyse or predict aspects concerning the data subject’s performance at work, economic situation, health, personal preferences or interests, reliability or behaviour, location or movements”¹³.

A data protection officer is the person designated by the controller and/or the processor in case the processing a) is carried out by a public authority, b) contains systematic monitoring of data subjects, or c) consists of large scale processing of special categories of data¹⁴.

These definitions, when applied to research, make clear that any research-based processing of social network data that not only directly identify but also possibly may identify (by the same researchers or third parties) individuals will be regulated under the GDPR.

Typically, in cases where research is conducted under the auspices of a university, the university is considered to be the data controller. While the universities are supposed to have organisational measures with regards to the GDPR, the researchers, as employees of the university, who one way or another process personal data as part of their role, are also expected to have an understanding of the GDPR as they design data collection and analysis protocols.

It may happen that other entities assist with the processing, but do not decide the purposes and the manner of it. For example, a researcher may pay individuals who are not employees of the university to perform a data collection.

¹³<https://gdpr-info.eu/recitals/no-71/>

¹⁴Art 37 GDPR

P1:	lawfulness, fairness and transparency
P2:	purpose limitation
P3:	data minimization
P4:	accuracy
P5:	storage limitation
P6:	integrity and confidentiality
P7:	accountability

Table 1: The seven basic principles in the GDPR

Where these entities are only following the guidelines of the controller, then these external parties can be seen as data processors, where the "processing by a processor shall be governed by a contract or other legal act"¹⁵. If the research is conducted by a private actor, such as a company, then it is the company that is the controller, and possible external sub-contractors (including researchers) may constitute the processors. It can also be so that a private party and a university can both jointly be regarded as controllers, depending on what agreement exists between these two parties. These complicated distinctions are important to consider and discuss with relevant internal data protection officers because their particular specifications can have an impact on the obligations of the researcher, or alternatively said, on how to comply with the GDPR. As a simple example, it is the controller that is responsible for providing specific information to the data subjects.

Given the diversity of research approaches it is important that researchers understand the particular aspects of the regulation that apply to them. This also means that any collaborative research project must consider what institutional agreements must be made with respect to data processing: a process that may take additional time and must be planned for.

The regulation also introduces seven principles, listed in Table1, to be followed when processing personal data¹⁶. (P1) The data must be processed in a lawful, fair and transparent way. (P2) Personal data may only be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes. (P3) The data may be processed only if they are adequate, relevant and limited to what is necessary with regards to the purpose of processing. (P4) Only data that are accurate and up to date, to the level that it is possible, may be processed. (P5) Personal data may only be processed for a period that is necessary for the processing and therefore the controllers must create criteria to determine what retention periods are suitable for their purposes. (P6) The controllers must apply technical and organisational measures in order to protect personal data they control against unauthorised and unlawful processing as well as accidental loss, destruction or damage. (P7) The data controllers have the responsibility to be compliant and to be able to demonstrate compliance when needed, which implies that written records must

¹⁵Art 28(3) GDPR

¹⁶Art 5 GDPR

be kept on whether and how the controller is compliant. These principles need to be instantiated to the specific cases. In the following sections we will highlight when the fact that data is processed in the context of research has an implication on the practical enactment of the principle, and we will discuss the meaning of these principles when they regulate the processing of social network data, emphasizing the cases where ambiguities arise.

3 The GDPR and Social Network Analysis

3.1 Lawful bases for data processing

According to the first principle, in order for data to be processed there has to be some lawful basis. The GDPR lists six lawful bases for processing of personal data¹⁷: a) the data subject has given her consent, b) it is necessary for the performance of a contract, c) it is necessary in order for the controller to comply with a legal obligation, d) it is necessary in order to protect individuals' (the data subject's and/or other natural persons') vital interests, e) it is necessary for the performance of a task carried out in the public interest and f) it is necessary for the purposes of legitimate interests pursued by the controller as long as these interests are not overridden by interests and fundamental rights and freedoms of the data subjects. Even though there are no specific lawful bases that are a priori dedicated to research, the three most relevant are the consent of the data subject, the task carried out in the public interest and the legitimate interests of the controller.

3.1.1 Consent vs. public interest

With a long history starting in medical science the practice of informed consent has been for long time the central pillar of research practices involving human subjects [7]. A key element of the GDPR is that, addressing a growing lack of satisfaction towards the efficacy of informed consent practices [17], it provides well-defined research exemptions.

Within the new scenario described by the GDPR the first issue is to understand whether and when consent is a legitimate legal basis. In particular, consent is not needed if the processing is necessary "for the performance of a task carried out in the public interest"¹⁸. Member State laws may provide specific provisions with regards to the processing in this case while national or Union law shall also define which actors may perform such acts. However, in case where no national legislation is introduced it should be accepted that public actors, such as universities, may use this lawful basis for processing of personal data¹⁹. National authorities are still in the process of developing the local leg-

¹⁷Art 6 GDPR

¹⁸Art 6(1)(e)

¹⁹SOU 2017:50, p.18. There remains of course the question which research can be regarded as a task of public interest, but this is left to the public authorities that have given by law the power to carry out tasks in the public interest, e.g. the universities

isolations, so we cannot generalize this point, but some relevant examples that are already available (i.e. Sweden) seem to grant a significant level of freedom to research done by public institutions.

In contrast, the basis of public interest is not available at all to commercial organizations and research labs - at least as long as no law provides for that - who must rely firstly on consent or otherwise on the legitimate interest legal basis to process personal data. The reason that we state here that private actors should only alternatively base their processing on the legitimate interest legal basis is because as it seems, art 6.1(f) is a last resort provision, meaning that if the controller cannot legitimise the processing on one of the 6.1(a)-(e) lawful bases only then it should be possible to lawfully process the data on the ground that the processing is necessary for the purposes of some legitimate interest (as long as these interests do not override interests and fundamental rights and freedoms of the data subject that require the protection of personal data).

Similarly if the personal data processed are of sensitive character, a public actor, such as a university, may primarily base the lawful processing of such data on the fact that the processing is necessary for scientific research purposes as long as appropriate measures are deployed according to art 89.1 and the research is based on a law "which shall be proportionate to the aim pursued, respect the essence of the right to data protection and provide for suitable and specific measures to safeguard the fundamental rights and the interests of the data subject" according to art 9.2(j) GDPR²⁰. Following the same argumentation as above we could, however, claim that if the processing is not necessary or if there is still no specific legislation with regards to processing for research purposes, consent could also be used as a lawful ground for such processing, according to art 9.2(a) GDPR²¹.

3.2 Data collection

Social networks can be obtained through a wide range of data collection strategies. Below we detail different approaches to data collection for social network analysis and consider the corresponding consequences of the GDPR. It is worth noting that we focus on networks where nodes represent natural persons: the GDPR does not apply when nodes represent companies, or animals, or even deceased persons (even though in this last case Member States may provide for specific rules²²).

²⁰Art 9.2(g), that the processing is necessary for reasons of "substantial public interest" could also be the basis for lawful processing of sensitive personal data but since art 9.2(j) specifically refers to scientific research purposes, processing that takes place for scientific purposes should be based on the legal ground of art 9.2(j)

²¹In many countries such processing by a university, even if consent is given by the data subject, could take place only after an ethics committee permits it. See also SOU 2017:50 s. 160.

²²Rec 27 GDPR

3.2.1 Primary vs. secondary data collection and the principle of transparency

An important conceptual and legal distinction resides in the selection of methods for data collection. For example, there is a significant difference between data that is collected directly from the data subject (e.g. small-medium scale data obtained through surveys) and data that is collected through a third actor (e.g. online social networks obtained from APIs). The difference here is not only in the scale or the nature of the data but in the relation between the data subject and the data controller: two different articles are concerned with providing information to the data subject when the data are collected from them²³ and when they have not been obtained from them²⁴.

In essence, these articles detail some of the ways that the principle of transparency must be put into action. Transparency addresses the right of the data subject to know and understand how the data are being used; it "requires that any information addressed to the public or to the data subject be concise, easily accessible and easy to understand, and [in] clear and plain language [in particular] in situations where the proliferation of actors and the technological complexity of practice make it difficult for the data subject to know and understand whether, by whom and for what purpose personal data relating to him or her are being collected [...]" If personal data are collected, the data subjects should be informed about the collection and its purposes in order to enable them to exercise their rights. Note that this is different from consent (explained in Section 3.2) but instead refers to the information that must be made available about data processing activities.

While making the data subjects aware of the processing and of their rights may seem straightforward when data are collected directly from them, this can become very difficult to accomplish when large networks are obtained from APIs. The potential difficulties to provide information under specific circumstances are acknowledged in the GDPR, where some exceptions are introduced. Article 14 states that providing information is not necessary if 1) "the data subject already has the information"; or 2) "the provision of such information proves impossible or would involve a disproportionate effort, in particular for [...] scientific or historical research purposes", subject to some safeguards²⁵, if providing information "is likely to render impossible or seriously impair the achievement of the objectives of that processing". Article 14 then continues stating that "[i]n such cases the controller shall take appropriate measures to protect the data subject's rights and freedoms and legitimate interests, including making the information publicly available" .

These are some examples of the kinds of research exemptions that are embedded in the GDPR, codifying and specifying research conduct. Both those exemptions apply to social network research based on online data collected from social media platforms assuming that social media platforms have already in-

²³Art 13 GDPR

²⁴Art 14 GDPR

²⁵Art 89 GDPR

formed their users through appropriate Terms of Services that their data will be shared with third parties (eg. through APIs) or assuming that the large scale of collected data will require a disproportionate effort to inform all affected data subjects. This is an example of balancing research needs against the derogation of the rights of the data subject. Technically termed "proportionality of the effort", this is a relatively vague concept. The controller, in order to determine whether it is going to be disproportionately difficult to provide the information, must take into consideration the number of data subjects, the age of the data and if there are any appropriate safeguards already adopted²⁶. If, after this assessment, the controller finds that the effort will be disproportionate, then she has to assess once again whether the effort involved to provide the information to the data subject exceeds the impact and effects on the data subject in the case where the information is not provided. This assessment has to be documented and depending on the outcome the controller may have to take extra measures (such as pseudonymisation or anonymization if possible and appropriate).

As an example, this means that although the research exceptions may not technically require that every single Twitter user of the millions involved in any large-scale Twitter network research be notified that their data are used for research, the logic involved in deciding to collect data and to skip the notification must be formally documented and appropriate storage, security and pseudonymization techniques must be considered. In addition, it is unclear whether providing information to these users should be considered an impossible or very difficult task. The concept of transparency is particularly relevant in the context of social network research, as previously highlighted e.g. by Borgatti and Molina [3], and as such it requires a more extensive discussion. In particular, some additional details should provide a better description of the obligations of the data controller with regard to the provision of information. There are three points that are important here.

First, the data controller must always provide information at the time the data is obtained (if obtained directly from the data subject) or within a reasonable period after the data is obtained and no later than a month (if the data is obtained indirectly) as long as this is possible given the appropriate adherence to the research exemptions detailed above.

Second, if the controller changes the purpose of the processing, she must provide the information to the data subject prior to this processing²⁷. For example, research data may have been collected for one purpose but the research question has shifted in the course of the data analysis and these data will now be used for a different purpose. This then speaks to how precisely the information about processing must be specified. Looking at rec 33, even though referring to consent, we can conclude that the specification in case of research can be a bit more general (such as the general research area or part of the project, not

²⁶Rec 62

²⁷Rec 61. See also Opinion where it is stated that in case the change is related to an incompatible further processing informing about the change does not "whitewash" other obligations of the controller, such as finding another lawful basis for the changed processing or asking for new consent

specific analytical task). Therefore changing data analysis approaches and even research questions may not require informing the data subject anew.

Related to the above is the fact that if the change leads to further processing that is incompatible to the initial purposes, mere information of the change does not "whitewash" other obligations of the controller. According to art 5.1(b) GDPR processing should comply to the purpose limitation principle. That means that as soon as the new processing is incompatible to the initial, the controller should either avoid the new processing or find a new lawful basis for it. There is, however, an exception with regards to research purposes, since in such case the further processing for such a purpose is considered to be compatible to the initial purpose.

Third, the general principle does not assume that the methods and the analysis are known in details at the moment of the data collection. However, the common practice in many areas of research where data is often collected with no specific hypothesis/evaluation framework becomes somewhat problematic because at least a limited explanation for the purposes of data processing is always necessary. The GDPR recognizes that it is not always possible to know from the beginning the entire scope of the research until the data is collected and used. Rec 33 (in case of consent) states that data subjects should be able to "consent only to certain areas of research or parts of research projects to the extent allowed by the intended purpose". Thus some specification of the intended purpose is necessary, limiting but not entirely eradicating exploratory forms of data collection.

3.2.2 The depth of online social network data and the principle of data minimization

Where some network data can be collected directly in the form of network information, that is, nodes and edges, many network datasets are obtained through processing of other types of data. For example this is often the case in research based on social media such as Twitter. Network studies of Twitter can be based on the user-articulated following/followers structure, that can be considered direct network information, or on networks mapping the interactions (tweets, retweets or replies) between users. To build this second type of network, researchers typically collect the content of users' posts and then extract and infer connection information. The problem arises if we consider the implications of collecting the content of the posts to build the network. Depending on the topic of posts, the type of content that is likely collected may vary but could include data revealing information that is not only identifying of natural persons but also includes sensitive data such as political affiliation, religious belief, etc.

The GDPR makes a distinction between different types of personal data, such as data with regards to ethnicity and sexual preferences (the so-called sensitive personal data²⁸), and in order for the processing to be considered

²⁸In the context of GDPR sensitive personal data is defined as "Personal data which are, by their nature, particularly sensitive in relation to fundamental rights and freedoms merit specific protection as the context of their processing could create significant risks to the fundamental

lawful the controller must respect the essence of data protection rights and follow suitable safeguards²⁹. Notice that data which in combination with other data can lead to revealing sensitive data may also be considered as sensitive data. For example name in combination with phone number, where each piece of data is not sensitive, may constitute sensitive data together if they probably reveal the ethnicity of a person. It is easy to see how the average stream of messages written by an average user might easily contain sensitive personal data or data that can be combined to reveal sensitive personal data about the data subject. Further, such data can be derived about persons simply from information produced by their connections. For example, it may be possible to ascertain a person's political affiliation if the majority of his connections explicitly communicate theirs.

Handling sensitive data is not forbidden, but before starting the data collection researchers need to plan some safeguards. Under the GDPR, controllers may not process sensitive personal data except if the subject has provided her "explicit consent"³⁰ or the data "was manifestly made public by the data subject"³¹, or in case of research purposes³². While one may consider using the concept of "manifestly made public" for special cases such as online social networks, where the information is publicly posted online by the users, we advise against this interpretation. In fact, in the context of social media, as a consolidated body of literature has made clear, assuming when something is "manifestly public" is problematic [5] and a potentially serious breach of standard ethical research practices. On the contrary, the exemption in case of research purposes can be used, but only if processing is necessary, in accordance to Article 89(1), based on Union or Member State law which shall be "proportionate to the aim pursued, respect the essence of the data protection and provide for suitable and specific measures to safeguard the fundamental rights and the interests of the data subject." Moreover, it seems that profiling on the basis of personal data is forbidden unless there are "suitable safeguards"³³. For example, in Sweden, it was recommended that one such security measure can be considered the decisions of the relevant ethics committee³⁴.

Finally, even if the data are not sensitive, the data minimization principle should still apply. Using again Twitter data as an example, when researchers collect information based on a hashtag they can fetch data using the hashtag with another meaning, and so not related to our study, or data using the hashtag as was intended, but still including additional unwanted information. This means that researchers must put in place mechanisms that will effectively strip out unwanted data and delete it as soon as possible, acknowledging and documenting the process.

rights and freedoms."

²⁹Art 9 GDPR

³⁰Art 9(2)(a) GDPR

³¹Art 9(2)(e) GDPR

³²Art 9(2)(j) GDPR

³³Rec 51 GDPR

³⁴SOU 2017:50

3.3 Data analysis and profiling

Social network analysis includes a wide range of data analysis tasks. Sometimes whole-network statistics are important, for example to correlate the communication/interaction structure of a team or organization to its performance. Sometimes meso-level structures are of interest, for example if we want to identify communities [12, 10, 4] or other relevant sub-structures such as online conversations [15, 24] inside a larger network. The identified groups can then also be used to classify individual actors, for example assigning them to a given community or role. Other types of micro-level analysis involve the characterization of single actors, for example when the most central or prestigious actors are identified [25]. When individuals are the object of the analysis, which is the case for most of the tasks listed above, an important concept to be considered is profiling.

The GDPR puts a special emphasis on the concept of profiling by specifying the definition and codifying acceptable practices. Accordingly, in the GDPR profiling is composed of three main stages "a) collection of personal data; b) automated analysis to identify correlations; c) applying the correlation [the result of b)] to an individual to identify characteristics of present or future behaviour"³⁵.

Note that the notion of "automated analysis" is used in the GDPR in opposition to "manual". Although both types of processing are under the purview of the GDPR, profiling is necessarily automated. However, automated here would mean both the use of a statistical software for conducting any form of data analysis as well as the use of more complex approaches such as machine learning algorithms. Thus any data analysis that includes computational assistance from software falls under automated analysis and thus can be classified as forms of profiling.

Given the above, many (but not all) social network analysis tasks can be classified as profiling. All centrality measures are clear examples, as they associate results of the network analysis to specific individuals. Any analysis that singles out individuals based on the identification of positions, roles and communities is similarly a form of profiling.

What is the researcher to do if their activities constitute profiling of the data subject? This does not mean that the particular data analysis is disallowed. However, this may require the performance of a data protection impact assessment (DPIA), typically conducted by the data protection officer. Although the GDPR states that profiling has to be systematic and extensive to require a DPIA, many authorities have made a broader implementation and if profiling may affect individuals in general (e.g. it provides custom access to services, it includes sensitive data, is related to vulnerable individuals, and in general the processing can lead to a high risk to the rights and freedoms of the data subject) and if it is conducted in a large scale combining sensitive data, then a DPIA is in general necessary. The question of whether a DPIA is necessary

³⁵Art 29 Data Protection Working Party, WP251rev.01, "Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679"

is clearly a very important one, because a very strict approach leading to an assessment for every possible case of social network analysis can become practically problematic for the researchers. While we wait for more guidelines³⁶ and other legal specifications, the role of the researchers together with the DPOs deciding on whether an assessment is needed or not (following the law but also being practical) is of even higher importance.

Alongside profiling, DPIAs are also applicable to systematic monitoring of individuals and locations. An interesting question arises with respect to what constitutes locations and public spaces. For example, the GDPR mentions a "systematic monitoring of a publicly accessible area on a large scale" as a reason for a DPIA³⁷. We are not aware of existing legal interpretations of whether e.g. Twitter is a publicly accessible area, but the WP29 interprets "publicly accessible area" as being any place open to any member of the public, for example a piazza, a shopping centre, a street or a public library. Clearly these are examples of physical places but Twitter is also a place that is open to any member of the public provided they have the means to access it (an internet connection and access to an email address). Such questions will likely be decided later on as the regulation stands the test of time and litigation, but it is an important item to consider for researchers conducting large-scale collection and processing of ostensibly "public" data.

3.4 Data storage

In this section we discuss what happens after the research is concluded, in case the researchers want to store the collected networks. If the data are still personal, e.g., they still contain identifiers or have been pseudonymized, then the data controller must guarantee some rights to the data subjects if she wants to keep the network data. On a general level we can organize these rights along three lines: a) temporal duration of personal data storage, b) the accessibility of the stored data to the data subject, c) the right of the data subject to withdraw his/her data. All these tasks are in general strictly regulated by the GDPR, but with significant exemptions for research, discussed in the following section. Under the assumption that the networks have been anonymized, then there is no problem because the GDPR no longer applies: the data are no longer personal. However, network anonymization is a complex issue, that we also discuss below.

3.4.1 Rights of the data subjects and the principle of storage limitation

When it comes to temporal storage limitation, the GDPR states that in general data can be "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", but more extended periods may apply in case of research as long as

³⁶<https://www.ucl.ac.uk/legal-services/research/data-protection-impact-assessment>

³⁷Art 35(3)(c) GDPR

appropriate safeguards are followed³⁸.

No exemption because of research is instead mentioned regarding the data subjects' right to check if there is data concerning them, and the right to obtain these data³⁹. This means that when requested the controller should provide the data, in a "commonly used and machine-readable format"⁴⁰ (even if there are possibly other national laws that may restrict this right of a data subject, such as for example secrecy acts⁴¹). Considering the average amount of data represented by a single node in a typical social network project, this should not be a problem. Nevertheless, as for other parts of this article, the size of the network may constitute a practical difference, and for large networks researchers should probably consider implementing an automated data filtering functionality.

Finally, the right to erasure, also known as right to be forgotten, grants to the data subject "the right to obtain from the controller the erasure of personal data concerning him or her without undue delay"⁴². However, also in this case the GDPR contains an exemption to this obligation if the erasure "is likely to render impossible or seriously impair the achievements of the objectives of that processing"⁴³. Many SNA measures are not so sensitive to a small amount of missing data [13] and the discipline has developed a set of techniques to handle missing data. Nevertheless, it should be acknowledged that a significant number of subjects requesting their data to be removed might seriously impair the research objectives, thus researchers would have the right to legally object to the data removal.

While these are the general guidelines emerging from the GDPR, according to art 89(2) Member States may further limit the data subjects right to access, rectification, restriction and to object in case of research if there are appropriate safeguards in place, and as long as the derogation is necessary for the fulfillment of the research.

3.4.2 Data anonymization

The GDPR asks for appropriate safeguards. The safeguards that are named in the GDPR are technical and organizational, e.g. data minimization, pseudonymization and anonymization. In addition, there can also be legal safeguards, such as contractual clauses between the controller and the processor, ethical vetting etc. [16]. Here we focus on anonymization, which should result in the data not being re-identifiable by the controller or any other person. In social network analysis, the typical approaches to anonymization are based on clustering, graph modification or network perturbation [26].

Data anonymization approaches in general are part of a considerable debate where some researchers argue that anonymization is impossible while others

³⁸Art 5(e) GDPR

³⁹Art 15 GDPR

⁴⁰Art 20 GDPR

⁴¹SOU p. 223

⁴²Art 17 GDPR

⁴³Art 17 GDPR

contend that it is in some cases [8, 1, 19]. Social network data is far more difficult to anonymize than other types of data and research on appropriate anonymization techniques is still in its relative infancy. Many of the simpler and more traditional approaches such as replacing node identifiers as well as more recent and complex approaches have been critiqued as insufficient [2, 1]. The knowledge of research being conducted in a particular location by a specific research group may be enough to reveal the identities of individuals encoded in the network to those who are familiar with these people more directly. In a small social network, such as for example a company division, it may be simple for the people in the network to recognize others based on just the revealed relational patterns [3]. Such an issue is not specific of social networks, but has been amply documented in qualitative and ethnographic research [22, 23]. As another case, if the data are public and indexed (e.g., by Web search engines), it can be very easy to find the original data using a part of it as a search key, such as finding the authors of a social media post based on the text of the post.

Whether anonymization or even just pseudonymisation are generally possible in a social network context is a difficult question. The GDPR states the necessity for privacy by design and by default but does not request specific privacy-preserving solutions: the controller should select and apply the appropriate measures for each case. In the GDPR, pseudonymisation requires the "additional information" to be "kept separately" and to be "subject to technical and organisational measures"⁴⁴, which is not really possible when the data source is public: if one removes the user identifier but keeps the text of the post (e.g., the tweet), a simple search on a search engine or on the social media platform can easily lead to the original, complete information. In this case, a possibility to be considered by the researchers (but not explicitly required by the GDPR) is to transform the text so that the analysis can still be performed but it becomes more complicated to fetch it from the Web, such as replacing it with a bag of words. The relevance of this discussion is that according to rec 26 pseudonymized data is identifiable, so the GDPR applies to that, while anonymized data is not, so the GDPR is no longer relevant. However, given the difficulty in fully anonymizing the data we should often assume that the GDPR is still the relevant regulation.

Even when we do not need identifiers to process social network data, because for example we are only interested in the structure of the network and its relationship with some indicators, we still need the identifiers if we want to extend the network, to know to which nodes the newly available information refers to. According to the GDPR we should at the very least pseudonymise the data "as soon as possible" (recital 78). However, it is not unusual in online network studies to keep collecting data for months or even years, which means that "as soon as possible" may be as late as the end of the study. One solution here is to develop or extend data collection systems with built-in network pseudonymisation functions, for example automatically removing identifiers and separately storing a mapping to user accounts in a location that requires special

⁴⁴Art 4(5) GDPR

access credentials. Such solutions may seem overly onerous given the current accepted practices, but the GDPR forces us to rethink our attitudes towards data collection and the impacts of our practices more broadly. In addition, the idea of designing ethically-related features in social network analysis software has already appeared in the literature [3].

As a final note, while in the previous paragraphs we have discussed the difficulty of network data anonymization, there are specific types of social network data where anonymization is indeed possible. In ego-network data collection different actors are asked about their own social ties and perhaps those of their neighbors. Ego-networks are then analysed without reconstructing a common network for all of the participants. In this case, there is typically no need to know who the individuals are, which means that we can design a data collection that is already anonymised at the source. As a result, these data are outside the scope of the GDPR given the definition of anonymous data as "data rendered anonymous in such a way that the data subject is not or no longer identifiable"⁴⁵.

4 Some more general issues and considerations

4.1 A code of conduct for social network research

The opportunity of writing a code of conduct for research in social network analysis has been under discussion for a long time. In the special issue of Social Networks on ethical dilemmas in social network research, there was a mention to "efforts now underway within INSNA, the professional association for social network researchers, to establish a code of ethics" [6]. Several years later, at a board meeting of the same association⁴⁶ this was still under discussion, and it was noted that many members of the association are also members of other associations for which codes of conduct already exist (e.g., by the American Anthropologist Association, the American Political Science Association, the American Sociological Association, the AOIR), questioning the need for an additional effort.

We believe that this article can contribute to this discussion. On the one hand, we note that many issues highlighted in the previous sections are common to other types of research not necessarily involving social networks, including for example social science research in general, Internet research and big data analysis, even though some specific aspects of social networks have also emerged and the combination of relevant issues is also unique. On the other end, the broad picture emerging from our analysis of the GDPR is a complex one, and a whole section of the regulation⁴⁷ indicates codes of conducts as a way to reduce this complexity. In fact, once a code of conduct proposed by an association has been approved, registered and published by a supervisory authority certifying

⁴⁵Rec 26 GDPR

⁴⁶2016 Report to INSNA Membership prepared by the INSNA Officer

⁴⁷Sec 5, art 40-43

that the code is compliant with the GDPR and "that it provides sufficient appropriate safeguards"⁴⁸, then showing compliance with the code exempts the data controller from a number of obligations. In summary, after the enforcement of the GDPR the benefits of codes of conducts have increased, but their establishment requires additional effort because they require an authority to verify their compliance.

4.2 Towards GDPR-compliant social network software

Through the analysis of the legal obligations emerging from the GDPR we have seen many cases where the law can be considered a bottom line for ethics, where individual researchers shall consider more restrictive actions. For example, as we have discussed above, the GDPR explicitly mentions "disproportionate effort" as a reason not to provide information to the data subjects. This, when framed within the context of online data or of secondary analysis of already collected large datasets, might easily be used as a solid reason to perform research without informing the data subjects. But if it is true that large online data could easily count millions of potential data subjects, one can expect that for online sources it can be possible to automatically send notifications or messages informing the data subjects. While this may result in a potentially significant overhead of communicating with confused data subjects, the effort may be a first step in acknowledging that people that produce data must be treated with dignity and respect regardless of research aims. Development of standards for notification in large-scale data collection endeavours is necessary and may need to be taken up at the level of professional codes of conduct.

In these cases, we should also consider developing tools that can take care of notification automatically, reducing the claim of disproportionate effort rather than leveraging it as a way to side-step responsibilities in research. For example, in the growing context of Twitter research, sending a short tweet mentioning those user accounts included in the social network data collected in a research project would be potentially interesting information for data subjects, contributing to the creation of an awareness about how much our public data is used. If done by a relevant share of researchers (which can theoretically be achieved if the main tool or tools for data collection are extended with this functionality) this increased awareness could result in a consequential generalized improvement in the way people manage their data online and an increased trust in science, showing how careful researchers are about this. However, while automatically sending the information to a list of users seems to require a limited effort, turning this into practice can be problematic, as described in the next section.

⁴⁸Art 40(5)

4.3 An experiment on automated data subject information on Twitter

To better understand the amount of effort needed to notify data subjects in the context of online social network research we have set up a protocol for a Twitter data collection process. This experiment, briefly reported below, highlighted the difficulty of performing even a task (appearing to be) as easy as sending some information to online users.

First, we had to consider a number of alternatives. First, when tweets are collected on Twitter the only contact information we have are the Twitter identifier and screen name of the accounts whose tweets were collected. This means that we can only inform Twitter users via Twitter, and as it is typically not possible to send direct (private) messages to generic Twitter users (for example, users not following us who have not explicitly allowed this in their privacy settings) we need to inform them in some way that is visible to others, such as using a public mention.

While this is not necessarily problematic, it is interesting to see how to inform data subjects so that we can protect their privacy we have to release additional information about them: our public message implies that those accounts have posted tweets with the hashtag we were monitoring, that by the principle of transparency we have to clearly indicate in the communication. Then, we must decide (1) how many users to mention in the same tweet and (2) whether we should check what their current screen name is. Both choices have an impact on the time needed to send the notifications: including more accounts in the same Tweet would reduce notification time, but would also again release more information as each notified user would see the other user names in the same tweet, knowing that they have also used the same hashtag. Checking the current screen name would require more accesses to the Twitter API but would avoid that we mention the wrong account because screen names can change in time. As an indication, the Twitter API currently allows us to send 2400 tweets per day, meaning that we would need around one year and two months to notify one million users (using a single notification account).

An alternative is to notify users through the hashtag, sending a tweet without mentions but with the monitored hashtag and specifying that we are collecting tweets containing the same hashtag. This is however also problematic, first because there is no guarantee that users will see that (they would have to search tweets containing that hashtag, at the right time), second because for some studies awareness of the data collection might result in a different behaviour.

Other decisions making the practical information process less trivial than one may think are whether we should also notify accounts mentioned in the collected tweets, even if they were not producing tweets themselves, or whether we should notify accounts retweeting other accounts' tweets.

After deciding on all these aspects, we started sending our tweets, including a link to the information about the project and the data processing and also information about the user rights. The procedure for the users to offer them (among other things) the possibility of retrieving the data about them we had

collected was also complicated, because to prove their identity the users were requested to follow our notification account - which is again revealing more information about the user and also requires some effort that might discourage potentially interested users. After sending notifications to 45 accounts we registered only one visit to the information page.

Finally, Twitter blocked our notification account. According to their rules, the account had been marked as having a spamming behaviour. In the process to reactivate the account we mentioned that despite the behaviour being compatible with their definition of spamming, the account was an attempt to enforce the rights of the users to know that their tweets had been collected and why, but this did not result in any exception, which led us to drop the experiment after considering that developing a "smarter" bot sending the notifications and trying to behave in a way not to be caught by Twitter's algorithms would have been ethically questionable.

4.4 Public vs. private actors

Another important point of discussion is the difference in classification of universities and commercial research and industry labs. The stark difference in legal basis for data processing and the impact on the consideration of whether consent is a legitimate legal basis is an important point to consider. What does consent constitute in the context of a commercial entity when it must be clearly uncoerced and freely given? What are the different obligations towards data subjects for researchers depending on the legal ground they employ for data collection? These questions have some answers in the GDPR but will be further evolving as time and litigation test the GDPR terms and definitions. Data availability for large scale computational social science and social network research is necessarily connected to commercial actors [14]. Collaborators across academic and commercial spheres have claimed the unalloyed public good that is possible from large-scale data collection, but what impact may the GDPR have, given the differentiation it makes between public and commercial research efforts? How much access will public university researchers continue to have to commercial data stores? How complex will these negotiations become? These questions are beyond the purview of this paper, but must be discussed and considered in the future.

5 Conclusions

Our main objective when we started writing this article was to provide a practical guide to GDPR-compliant social network data processing. Working on it, and also trying to apply our recommendations to our own research, it became evident that while some issues could be more easily translated into practical suggestions, other general indications and principles in the regulation are difficult to either interpret or apply in the context of social network research. The problems we have highlighted in the article include the difficulty of sending in-

formation to millions of users through a third-party API that does not allow it, the problems in pseudonymizing the data as soon as possible in a continuous network monitoring process performed with pre-GDPR software tools, the interpretation of concepts such as "manifestly made public data" and "publicly accessible areas", the problem of removing data by user request not knowing what impact this will have on network statistics, the practical impossibility of guaranteeing respondent anonymity, the inclusion of data about individuals not included in the study, as well as more general issues related to data protection emerging when social network analysis is applied to large-scale networks of social relations derived from social media data.

In summary, it is important that everyone involved in the processing of social network data invests some time to reflect about the implications of the GDPR on their research, seeking help from their institutions but not only relying on institutional support. While this may sound as an obvious statement, and legal and ethical problems related to social network analysis and Internet research for social behaviour have certainly received a lot of attention in the past as witnessed by the literature on the topic and by existing codes of conduct, the sudden explosion of online behavioural data has indeed affected the research landscape by both introducing new problems and involving new researchers from disciplines where some of these problems had not been traditionally accounted for.

6 Acknowledgments

This work was partially supported by the European Union's Horizon 2020 research and innovation programme under grant agreements No. 727040 (Virt-EU).

References

- [1] Narayanan A. and E.W. Felten. No silver bullet: De-identification still doesn't work. Technical report, 2014.
- [2] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography. In *16th international conference on World Wide Web*, pages 181–190, 2007.
- [3] S.P. Borgatti and J.-L. Molina. Towards ethical guidelines for ethical research in organizations. *Social Networks*, 27, 2005.
- [4] Cecile Bothorel, Juan David Cruz, Matteo Magnani, and Barbora Mícenkova. Clustering attributed graphs: models, measures and methods. *Network Science*, 3(3):408–444, 2015.

- [5] D. Boyd. *Social network sites as networked publics: Affordances, dynamics, and implications. A networked self*. Routledge, 2010.
- [6] R.L. Breiger. Introduction to special issue: ethical dilemmas in social network research. *Social Networks*, 27, 2005.
- [7] A. M. Capron. Where Did Informed Consent for Research Come From? *The Journal of Law, Medicine & Ethics*, 46(1):12–29, 2018.
- [8] A. Cavoukian and K. El Emam. Dispelling the myths surrounding de-identification: Anonymization remains a strong tool for protecting privacy. Technical report, Information and Privacy Commissioner of Ontario, Canada, 2011.
- [9] G. Chassang. The impact of the EU general data protection regulation on scientific research. *ecancer*, 11(709), 2017.
- [10] Michele Coscia, Fosca Giannotti, and Dino Pedreschi. A classification for community discovery methods in complex networks. *Statistical Analysis and Data Mining*, 4(5):512–546, oct 2011.
- [11] danah Boyd. *Taken Out of Context: American Teen Sociality in Networked Publics*. PhD thesis, University of California-Berkeley, School of Information, 2008.
- [12] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
- [13] Gueorgi Kossinets. Effects of missing data in social networks. *Social networks*, 28(3):247–268, 2006.
- [14] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, and T. Jebara. Life in the network: the coming age of computational social science. *Science*, 323(5915), 2009.
- [15] Matteo Magnani, Danilo Montesi, and Luca Rossi. Conversation retrieval for microblogging sites. *Information Retrieval*, 15(3-4):354–372, feb 2012.
- [16] C. Magnusson Sjöberg. Scientific Research and Academic e-Learning in Light of the EU’s Legal Framework for Data Protection. In *New Technology, Big Data and the Law. Perspectives in Law, Business and Innovation*. 2017.
- [17] A. Mantelero. The future of consumer data protection in the EU Rethinking the notice and consent paradigm in the new era of predictive analytics. *Computer Law & Security Review*, 30(6):643–660, 2014.
- [18] L. Marelli and G. Testa. Scrutinizing the EU General Data Protection Regulation. *Science*, 360(6388), 2018.
- [19] A. Narayanan, J. Huey, and E.W. Felten. A Precautionary Approach to Big Data Privacy. In *Data Protection on the Move*, pages 357–385. 2016.

- [20] S. Penasa, I. de M. Beriain, C. Barbosa, A. Bialek, T. Chortara, A.D. Pereira, P.N. Jiménez, T. Sroka, and M. Tomasi. The EU General Data Protection Regulation: How will it impact the regulation of research biobanks? Setting the legal frame in the Mediterranean and Eastern European area. *Medical Law International*, 2018.
- [21] K. Schaar. What is important for Data Protection in science in the future? General and specific changes in data protection for scientific use resulting from the EU General Data Protection Regulation. In *Working Paper Series of the German Council for Social and Economic Data*. German Council for Social and Economic Data (RatSWD), 2016.
- [22] I. Shklovski and J. Vertesi. Un-googling publications: the ethics and problems of anonymization. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 2169–2178, 2013.
- [23] W. C. van den Hoonaard. Is Anonymity an Artifact in Ethnographic Research? *Journal of Academic Ethics*, 1(2):141–151, 2003.
- [24] Davide Vega and Matteo Magnani. Foundations of Temporal Text Networks. *Applied Network Science*, 3(1):25:1—25:26, 2018.
- [25] Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications*, volume 8 of *Structural analysis in the social sciences*, 8. Cambridge University Press, 1994.
- [26] Jin Zhou, Xiaoke Xu, Jie Zhang, Junfeng Sun, Michael Small, and Jun-an Lu. Generating an Assortative Network With a Given Degree Distribution. *International Journal of Bifurcation and Chaos*, 18(11):3495–3502, 2008.