

Responsible Research on Social Networks: Dilemmas and Solutions

Jon Crowcroft, Hamed Haddadi, Tristan Henderson

May 30, 2018

The final version of this paper can be cited as Jon Crowcroft, Hamed Haddadi, and Tristan Henderson. Responsible research on social networks: Dilemmas and solutions. In Brooke Foucault-Welles and Sandra González-Bailón, editors, *Oxford Handbook of Networked Communication*. Oxford University Press, Oxford, UK, 2018. doi:10.1093/oxfordhb/9780190460518.013.28. This copy is reproduced by permission of Oxford University Press.

Abstract

Our digital footprint today spans over a large number of activities on the Web, smart-phone applications, wearable devices, Online Social Network (OSN) profiles, and the data collected about our online presence and behaviour by third-party analytics companies. As well as companies, scientists also often perceive OSNs as a goldmine for research into various aspects of social behaviour and inter-personal communication. For example, observing social interaction between individuals, their engagement in conversations, or performing sentiment analysis on these communications, are often carried out for research in a number of disciplines, e.g., mental health, sociology, or political predictions. Such studies introduce many challenges for conducting research in a responsible manner. Data may be repurposed or cross-correlated in ways that participants may not have anticipated or desired, private information may be measured or collected, or legal requirements may not be met. In this chapter, we explore some of the challenges and dilemmas met by industry, academia, regulators, privacy advocates, the data-driven society, and ultimately the individuals using these online services. We provide a number of arguments for and against the collection, analysis, and archiving of personal data specifically for digital research. We conclude by discussing a number of theoretical and practical approaches that target these dilemmas.

1 We are the data!

We have been gathering *Big Data* throughout history. Land surveys such as the Domesday Book or the Napoleonic Cadastre have long been conducted, while a population census of the Roman Empire is mentioned in the Bible. Of late, however, advances in technology have led to a sharp fall in the relative costs of gathering data about individuals and their traits, leading to increases in the so-called “three Vs” (volume, velocity and variety) of Big Data (Laney, 2001). The use of computers in the systems around us means that data are often generated as a side effect of some other task. Vehicle detectors in intelligent transport systems can be used for monitoring congestion or adjusting speed limits, but also for identifying cars or drivers. Smart electricity meters can be used for charging, but also for correlating usage with social deprivation demographic data. This extends to individuals themselves, as we increasingly use online social networks (OSNs), and they become integrated with smartphone applications (apps), and other wearable and *Internet of Things* technologies to create new *Social Computing* applications. These enable people to collect and generate their own data through the self-reporting of hobbies, interests, activities, and emotions, but at the same time enable accurate behavioural profiling by gathering browsing and shopping data from across the web through the use of social widgets (e.g., the Facebook

like button) and the rich interconnection between analytics and tracking services (Falahrastegar, Haddadi, Uhlig, & Mortier, 2014).

This increase in the amount and types of data, and the ease at which they can be collected, has led to a corresponding increase in research based around such data collection, with thousands of papers published using Facebook alone (Caers et al., 2013). Such research is essential for improving our understanding of this wide techno-social landscape. Indeed understanding whether human behaviour has changed psychologically or socially because of the Internet or technology use could be difficult without collecting such data. These new technologies have also enabled improvements in new research data collection methods, such as “citizen science” (Silvertown, 2009), while the resultant large quantities of data have enabled new methods of analysis, such as the so-called “discovery-driven” or “hypothesis-free” research (Aebersold, Hood, & Watts, 2000).

The data collected from people can also be used for societal and individual benefit. For example, location-based services can be used for predicting journeys, optimising traffic routes, or public transport services; but the data can also be used for delivering targeted advertising or providing detailed (or potentially intrusive) location-based services. Increasingly, the data *generated by us, or inferred about us*, creates a rich information ecosystem which has benefits and harms.

In this chapter, we focus on the challenges in conducting responsible and ethical research when using data collected from OSNs for research purposes, from active volunteers and passive participants. We discuss the potential benefits and also the potential pitfalls of such digital research, what it means to conduct responsible research, and the current technical limitations that might impede our ability to do so. We then describe some best practices, both in terms of what researchers can do technically, but also legally and socially.

Ethics itself is of course a rather overloaded term (in the computer science sense¹), with many definitions. For this reason some prefer to use the notion of “responsible research” (Owen, Macnaghten, & Stilgoe, 2012). But Shamoo and Resnik (2009) offer four definitions of ethics in relation to research, and for our purposes we focus on the first: “ethics as standards of conduct that distinguish between right and wrong, good and bad, and so on”. (Singer, 2011, p. 14) in his well-known book describes such ethical standards as being somehow decided by considering universal preferences, and making decisions that are ethical according to “preference utilitarianism”. But ethics and morality are closely linked (Farrimond, 2013), and when considering preferences, perhaps we should limit ourselves to moral actors.² Thus, when discussing the benefits and pitfalls of digital research, we need to consider the preferences of all actors involved in these systems.

2 The good bits

Many of the previously mentioned examples in this chapter and the book as a whole relate to the use of big data for understanding aggregate behaviours (e.g., overall road traffic, overall political patterns, or statistical distribution of mental or physical health conditions in individuals or the society). Such data are useful to governments, industry, and researchers aiming to optimise their services or improving scientific insight in relevant areas. They can be used to identify the spread of rumours and media in the society (Cha, Pérez, & Haddadi, 2009), or understanding bias in individuals’ views or those of the media (An, Cha, Gummadi, Crowcroft, & Quercia, 2012). Similarly, location and picture data from different OSNs such as Instagram and Twitter can be used to understand the health habits of individuals and communities with respect to issues such as physical activity and obesity (Widener & Li, 2014; Mejova, Haddadi, Noulas, & Weber, 2015), finding investors for crowdfunded projects (An, Quercia, & Crowcroft, 2014), or to aid relevant organisations in providing fast and efficient disaster relief in emergencies (Díaz, Aedo, & Herranz, 2014; Imran, Castillo, Lucas, Meier, & Vieweg, 2014). Data from OSNs can also be used for rec-

¹*Overloading* is a form of *polymorphism* where different functions or methods with the same name can be invoked depending on context. Informally the term can also be used to reference the use of the same term to invoke different meanings.

²“it is precisely the moral persons who are entitled to equal justice” (Rawls, 1999, p. 442).

ommending friends, new content, books, and entertainment to individuals using the social plugins provided by these services (Konstas, Stathopoulos, & Jose, 2009; Seth & Zhang, 2008).

The data collected from these sources can be used for inferring public health (Mejova et al., 2015; Paul & Dredze, 2011) and physical activity traits (Cavallo et al., 2012), as well as giving useful feedback and behaviour interventions to individual. The vast amount of data available from OSNs, complemented by the wealth of personal data obtainable from wearable devices and smartphones shared online, provide the research community with opportunities bounded only by consumer willingness, scientific creativity, data access regulations, and ethics (Mortier, Haddadi, Henderson, McAuley, & Crowcroft, 2014). These range of *Social Computing* applications often rely on collection of personally identifiable information (PII) from individuals. Though some services and organisations mandate ethical approval for dealign with such sensitive data, and some research fields are developing subject-specific guidelines (Bailey, Dittrich, Kenneally, & Maughan, 2012), there is as of yet no universally agreed framework for the collection, archiving, and use of personal information.

The examples above present a glimpse at opportunities and huge potentials in the use of social media. While we will not spend much longer on the benefits for the individual, or the public, we encourage the reader to consider the value of these services while reading the rest of this chapter where we present the challenges and limitations.

3 What could possibly go wrong?

Individuals on OSNs are often attracted by the *network effect*, with one of the main attractions being the ability to maintain their relationships with offline friends and family, but beyond these much useful information can be found through *weak ties* (Granovetter, 1973) and, from users self-organising to create well-connected communities (De Meo, Ferrara, Fiumara, & Provetti, 2014). In doing so, users may inherently divulge a large set of personal information about themselves (e.g., their location and activity), and their connections (social circles), through time. Researchers enjoy the wealth and depth of data, the interactions between individuals and organisations, and content sharing patterns on social media. However collecting this data has potential risks for the individual using these services as the data can be used to identify small groups or individuals, and thus potentially expose sensitive information.

Over the past few years several research studies have attracted attention in the media, perhaps unexpectedly, due to their use of sensitive data, jeopardising individuals' privacy, or affecting the participants' emotional wellbeing. In 2008, a set of university researchers employed students to help them crawl the Facebook "walls" of an entire undergraduate class (Selwyn, 2009). Similar research has been conducted by creating fake user profiles in order to collect individuals' profiles in certain geographic areas (Viswanath, Mislove, Cha, & Gummadi, 2009; Haddadi & Hui, 2010). These highlight the vulnerabilities of these OSNs to impersonation and honeypot attempts, though the intentions may have been scholarly studies.

One of the most well-known examples in this space was the 2014 "emotional contagion" Facebook/Cornell study (Kramer, Guillory, & Hancock, 2014), which faced much criticism over its legal and ethical aspects (Schroeder, 2014). This experiment involved tuning the individuals' newsfeed based on the sentiment of the posts from their network, and observing the effects of these posts on their consecutive posts. Participants did not consent to this manipulation, and the Facebook user agreement was only changed after the experiment to allow this use. The wide public reaction highlighted the broader impact and importance of scientific experiments, and the way their results are communicated to the public. In similar manner, an earlier Facebook experiment in 2008 received some criticism when the supposedly anonymous study was quickly deanonymised, revealing sensitive personal and private information (Zimmer, 2010).

Individuals on OSNs, or the Internet in general (e.g., search engines), are often unaware of the fact that they may be part of an ongoing experiment. The passive data collection from OSNs poses a level of threat on individuals' anonymity and privacy. Similarly, research on Internet openness or social network data mining can lead to ethically unacceptable practices, and may be

illegal in many jurisdictions and dangerous for individuals which reside in those locations (Wright, Souza, & Brown, 2011). Sometimes willingness to share OSN data is temporally sensitive (Bauer et al., 2013; Ayalon & Toch, 2013), though this does not reduce the severity of the risks involved. Research on mental health issues such as depression (De Choudhury, Counts, & Horvitz, 2013), sexual orientation, political views, personality, happiness and more personality traits (Kosinski, Stillwell, & Graepel, 2013) can have consequences on individuals’ employment, relations, and their personal lives. Future inferences and correlations (data collected today from an individual) may also be harmful due to dependancy on their future social, economical, or political status.

Anonymisation and de-identification by removing names, or mapping names to IDs, have long been utilised to release and use OSN and similar data for research purposes. However the de-anonymisation of Netflix Prize dataset and understanding the privacy risks of models such as K-anonymity (Sweeney, 2002) lead to widespread exposure of risks in traditional anonymisation techniques (Narayanan & Shmatikov, 2008; Machanavajjhala, Kifer, Gehrke, & Venkitasubramaniam, 2007). Despite further efforts in de-identification, this still remains a challenge for large scale OSN data usage (Narayanan & Felten, 2014). Location privacy studies have also identified the ease of predictability of mobility patterns and social ties using mobile phone data (De Domenico, Lima, & Musolesi, 2013), and linking OSN datasets with other mobility traces can make deidentification difficult (Ji, Li, Srivatsa, He, & Beyah, 2016). In Section 5 we will discuss some potential solutions in the literature.

Taking a step back and taking a broader view, we can categorise the risk levels and their probabilities of occurrence by looking at a number of example threats:

- Psychological harm: from embarrassment at falling for a deceptive experiment, through being exposed to offensive content, through being trolled or stalked.
- Economic harm: loss of money or property, identity theft, loss of job or reputation as an affect of content shared on social networks.
- Physical harm: threat to individuals’ lives due to their ideology, political affiliations, or religious views, based on knowledge gained from social media for personal purposes.

The above examples present just a few of the high risk threat categories to individuals using OSNs. But what are the challenges and obstacles which need to be overcome to protect individuals on these services? Are the outcomes and resulting inferences worth the potential risks? Can we protect against all potential risks, and indeed should we if protecting against a risk makes the resulting data less beneficial?

4 Technical limitations: can we overcome these?

In the last sections we discussed the pros and cons of using OSNs. One might ask: how can we use science to protect the individuals from privacy and ethics issues with large scale data analysis, yet enable us to perform research beneficial to society? In many studies, testing a certain hypothesis requires thorough understanding of the relationships in a network, or analysing the context and sentiment of tweets. Hence we need to consider a large number of varying privacy laws and ethics considerations in order to distinguish between cyber espionage and performing scientific research.

One of the most fundamental analysis techniques used by researchers is the use of graph analysis techniques on OSN users. Graph techniques can enable understanding of communication patterns and trends, or the role and influence of individuals in OSNs (Cha, Haddadi, Benevenuto, & Gummadi, 2010; Cha, Benevenuto, Haddadi, & Gummadi, 2012). Yet we are unable to study the topological characteristics of graphs fully de-identified in a way that some nodes cannot be de-anonymised and linked back to their neighbours (Narayanan & Felten, 2014; Narayanan & Shmatikov, 2008). Advances in applications of *differential privacy* have enabled progress in this space by achieving a tradeoff between maintaining structural similarity and privacy protection (Sala, Zhao, Wilson, Zheng, & Zhao, 2011). Other techniques such as *homomorphic encryption*, and subsequently *secure multi-party computation* (Lindell & Pinkas, 2009), have enabled

some computations to be done on sensitive data (Kocabag & Soyata, 2014; Torres, Bhattacharjee, & Srinivasan, 2014). However we are still unable to perform scalable contextual or sentiment analysis on fully encrypted content or large multi-dimensional datasets while fully preserving privacy. Full reliance on encryption remain an ideal goal in this space.

It is important to understand that technology does not, and indeed cannot, provide a complete panacea for many of the challenges and threats that we have outlined. For data to be useful, they need to be processed. If someone uploads a photo or a status update to an OSN, they expect that photo to be viewed, or that status to be read. If someone purchases a product based on a recommendation from an OSN friend, they expect that product to be delivered. It is the processing of these data that enable some fundamental attacks, as however legitimate data access may be, the persons to whom access has been granted may renege, and may:

1. *re-broadcast* data. This problem is well-known to the media industries, as no matter how you attempt to secure the storage and delivery of music or films, consumers have to be able to play or view these media at some point. Hence no DRM (Digital Rights Management) scheme, regardless of complexity or sophistication, has been able to block this “analog hole” (Sicker, Ohm, & Gunaji, 2007), i.e., been able to prevent people from recording their music coming out of a headphone jack or speakers. So copyright holders have pursued legal remedies rather than technological ones.
2. *re-identify* data. It has been demonstrated that the “power of 4”, e.g., four uses of a credit card online, or four check-ins on an OSN such as Foursquare or Facebook, is all that is needed to identify, with a high probability, a unique individual (de Montjoye, Hidalgo, Verleysen, & Blondel, 2013; de Montjoye, Radaelli, Singh, & Pentland, 2015).
3. *re-veal* data. Huge data breaches are a regular feature in today’s news.³ This is perhaps an unsurprising effect of the tendency to store lots of sensitive data in one system. Even though there may be good intentions for doing so (e.g., efficiencies and better healthcare through storing health data in a central system), the risks in terms of becoming an attractive target for attackers, or the total loss of privacy that results from such a data breach, must be considered.

5 Best practice from a technical perspective

What are the technical solutions to privacy-aware and ethical research on OSNs? How can we guarantee individuals’ privacy will be respected, their personal data will be protected, and data can not be used in future for malicious purposes? There have been a number of efforts in privacy-preserving, or private, information retrieval and analysis. In this section we will discuss a number of recent approaches to this challenge.

One of the most basic and fundamental responsibilities of researchers dealing with personal data is to ensure appropriate, role-based, access control measures to databases of personal data in place when collecting, analysing, and releasing personal data. This basic step would ensure basic protection against data being accidentally revealed to individuals who were not involved in the research, data collection or the analysis process. The natural follow-on procedure would be to perform strong encryption with careful management of decryption keys. These first steps would substantially raise the bar on effort level for an unintended individual trying to access the data. Although they may not be effective in protecting the individuals’ identity against the researchers, or a determined attacker.

More advanced techniques include:

- *Data fuzzing and noise addition*: The main aim in these techniques is aggregating data and building distributions in a way that an individual or the relationship between two subjects

³for instance, as we write this chapter, the US Office of Personal Management suffered a data breach that affected over 21 million people (Davis, 2015)

can not be identified. Methods such as differential privacy (Dwork, 2006) also provide guarantees about the accuracy of results of an analysis independent of an individuals' inclusion or exclusion in a database. They hence enable sharing data such as of OSN graphs (Sala et al., 2011), or addition of noise to survey results (Haddadi, Mortier, & Hand, 2012), without a major impact on the usefulness of final statistics.

- *Data retention/deletion policy:* Legislations such as the E.U. Data Protection Directive (*Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*, 1995) introduce rules on collection, use, retention and disposal of personal data intended for different purposes. In essence, introduction of more strict requirements such as the *Right To Be Forgotten* in the E.U. has led to introduction of mechanisms for individuals to request the removal of some of their digital footprint. Research in online reputation management and privacy has highlighted the importance of mechanisms for data deletion (Mayer-Schonberger, 2009).
- *Malleable encryption schemes:* These schemes provide an ability to perform certain computations on the ciphertext with results which match those which would have been performed on the raw data. The most important class of these schemes are partial and fully homomorphic encryption systems (Gentry, 2010). These schemes would allow the data from participants in research to be encrypted at the collection end by the user, hence limiting the exposure of sensitive data such as financial records or health data.
- *Privacy by Design:* Many research or surveying works can be done by limiting the data collection process in the first place, moving away from the traditional approach of collecting as much data as possible, in case one might find it useful at a later stage. Often, data collection and analysis can be performed at the user end, even for research involving resource-constrained devices such as smartphones or sensors (Haddadi, Hui, Henderson, & Brown, 2011; Lane & Georgiev, 2015). Researchers and data practitioners should ideally assess the cost-benefit tradeoffs of gathering data in the first place in order to prevent excessive data collection and processing of sensitive and unnecessary data.

These approaches present some of the ways in which researchers can reduce the unintended negative consequences of data collection process from OSNs or other sources of personal data. But are technical solutions ideal? Will they stand the test of time? Or should we expect responsible research subject to wider scrutiny and transparency?

6 Best practice from a legal and societal perspective

Technical systems can help us to ensure that data are stored or processed appropriately. But to understand what is appropriate, we need to understand how people, such as research participants, want their data to be used. Traditionally, researchers have been relying on ethics committee or Institutional Review Board (IRB) approvals, and *informed consent*. The notion of informed consent has long been the standard for determining whether participants' wishes have been met. "Informed" consent became enshrined in law in the 1957 Salgo case in the USA, where Martin Salgo became permanently paralysed in a translumbar aortography, and sued his physicians for negligence and failing to inform his of the risk of paralysis (Faden & Beauchamp, 1986). The court in this malpractice case suggested that doctors in such cases have a duty to disclose the risks and alternatives in treatments, while at the same time acknowledging that doctors should also use their discretion where appropriate. At the same time, similar ideas about informed consent in research ethics were first codified in the Nuremberg Code and the Declaration of Helsinki, and later in documents such as the CIOMS guidelines (Council for International Organizations of Medical Sciences, 2002), or the EU Data Protection Directive (European Parliament and the Council of the European Union, 1995). Probabilities of risks involved in de-anonymisation and

re-identification are not always easily comprehensible by researchers, let alone explainable in plain every day vocabulary to OSN users taking part in an experiment. On the other hand, it is not always clear what the role of researchers should be, in case they find legally questionable information from individuals' while performing these measurements.

Informed consent has a strong basis in moral theory, specifically the respect for autonomy. Although it has long been in use in medical ethics, Manson and O'Neill (2007) argue that it is insufficient. They find that current informed consent requirements are impractical, and that individual autonomy is only one part of the moral concerns in much research. One of their proposals is that universal standards for informed consent (e.g., standard forms that apply to all research situations), are unrealistic. Similarly, there is currently some debate as to whether informed consent is required for the collection of data from digital systems such as OSNs. Zimmer (2010) discusses the Facebook T3 study, concluding that people sharing data on an OSN does not give researchers the right to access those data for research. On the contrary, Solberg (2010) argues that Facebook research generally does not require ethics approval as it is of low risk. Buchanan, Aycock, Dexter, Dittrich, and Hvizdak (2011) looks at cybersecurity research, but proposes that we need a community process to determine standards: "we support a shared responsibility model, where researchers, REBs, and information technology experts work together to disclose, understand, and accommodate the unique ethical issues within CS research." This is akin to Manson and O'Neill (2007)'s concerns about standards, and the consideration of all actors within RRI.

Even if one does determine that informed consent is required or desired for a particular piece of research, it is not always clear whether obtaining consent is meaningful if the way in which people are informed is overly complex (Luger, Moran, & Rodden, 2013), or if, as in the Facebook example described in Section 3, terms of service may not be clear or even inaccurate (Check Hayden, 2012). Ioannidis (2013) analyses the claim that consent is meaningless for big data or discovery-driven research, and argues that we still need consent, in part to maintain trust between researchers and participants. Luger and Rodden (2013) further distinguish between *secured* consent, such as the traditional ticking of a checkbox at the start of an experiment or the acceptance of an EULA (End-User Licensing Agreement) when installing a piece of software, versus *sustained* consent, where participants might be probed in a sustained fashion to determine whether they truly consent to sharing data (a specific OSN example of this is provided by Sleeper et al. (2013), who studied self-censorship on Facebook by allowing participants to choose when and what to share with researchers over time). The need for the latter is borne out by Munteanu et al. (2015), who present a number of case studies, arguing that the consent process must be examined over the lifecycle of a study, while Neuhaus and Webmoor (2012) propose "agile ethics", akin to the agile software engineering process, where ethical considerations should be evaluated and revisited over time.

Sustained consent can be a burden for participants, who may be discouraged by continuous requests for consent in addition to requests for data. It may thus affect participation rates or even fatigue participants such that their consent decisions are incorrect. This has led to recent research into methods for reducing the burden of sustained consent while still providing more granular control than secured consent. Morrison, McMillan, and Chalmers (2014) show that allowing participants to periodically view representations of their data over the course of an experiment can improve engagement. Hutton and Henderson (2015) employ Nissenbaum (2004)'s framework of contextual integrity to predict when interactions are likely to violate privacy expectations and thus require new requests for consent, and show that this is a viable alternative to sustained consent for OSN studies. Gomer, Schraefel, and Gerding (2014) propose the use of artificial intelligence agents to achieve what they term *semi-autonomous* consent, where preferences can be elicited from participants and agents determine when consent would and would not be granted. Related approaches to this include the use of machine learning and recommender systems to set privacy preferences themselves, freeing users from the complex interfaces used for designing privacy policies in many OSNs (Toch, 2014; Zhao, Ye, & Henderson, 2014, 2016).

One particular concern is how data are used beyond the lifecycle of an experiment. This will become increasingly common as data sharing becomes mandated by funding bodies (EPSRC, n.d.; NSF, n.d.). Some participants may have concerns about using their data for other purposes, while others might welcome their use for some studies but not others. Kaye et al. (2014) recognise this

<i>Process</i>	<i>Techniques</i>
Data collection	Proportionate data collection, meaningful consent, pre-processing at source
Data storage	user-controlled data aggregation, source-based data anonymisation
Data analysis	Differential privacy, homomorphic encryption, Privacy by design
Data release	Distribution building, noise addition, differentially private anonymisation

Table 1: Example tools and techniques for best practices in dealing with sensitive data in research

tension through their notion of *dynamic* consent, whereby participants can be engaged in research over time, viewing how their data are used and allowing them to consider different contexts for their data. While proposed for biobank research, it could also be applicable to OSN data.

Beyond the notion of consent to participating in an experiment, we can consider how people would like to have more control of their data. Current systems enable very little control, with hardly any transparency about how or when data are collected (e.g., the various data protection controversies over Google’s Street View data collection using Wi-Fi-enabled cars), or transparency about how data are processed and transformed (e.g., controversies about how Facebook manipulates users’ news feeds to determine what information to present (Kramer et al., 2014)). We have proposed that understanding and enabling this control warrants a new research area of its own, so-called Human-Data Interaction (HDI) (Mortier et al., 2014). As the Internet of Things begins to be deployed, the amounts of data around us will be ever increasing, including information about health, well-being, food, home entertainment usage, and more. We need to understand what people will be willing to share with companies in return for services, or with researchers voluntarily, without assuming that they have to give up all of their privacy to allow such technology into their homes and lives.

Focusing back on legal aspects, another important challenge has been the recent introduction of the *Right To Be Forgotten* in Europe, which has enabled individuals to request removal of their data, or inaccurate articles referring to them, from search engines. Different social media and OSNs have also been increasingly enabling privacy tools and data removal requests. Some have gone to the extent of holding researchers responsible for removal of data items from their collections even long after the data collection has taken place, should an individual request the removal of a social media posting. Control also extends to legal requirements. For instance, people may wish to keep their cloud data in particular places, because they do not trust other governments, or because there are legal requirements that mandate the storing of data in particular jurisdictions (Hon & Millard, 2013).

The main barrier to treating responsible research and innovation as a purely legal issue is the often delayed response of the legal system to new advances in technology and the risks they can present. Hence researchers and scientists need to take into consideration not only legal aspects, but also ethical considerations, societal impact, and mutual trust when dealing with individuals data. Table 1 presents some potential tools and techniques for best practices in personal and social data collection, management, and analysis from OSNs. These are just examples of ways in which researchers can limit the usage of individuals’ private data beyond the intended scientific consequences.

7 What challenges and opportunities remain?

Today, OSNs occupy the largest part of screen time for hundreds of millions of Internet users across the world.⁴ With high levels of engagement modes, variety of content, and convenience of data collection mechanisms, these platforms provide the perfect environment for studies of human behaviour, benchmarking sociology and psychology theories, and opportunities for performing large scale A/B testing. However there are numerous ethical challenges which need to be considered

⁴In a recent earnings call, Facebook reported that on average people spend 50 minutes a day on their sites (D’Onfro, 2016).

when performing research using OSNs. As shown in this chapter, technical or legal solutions alone are not enough to prevent privacy disasters or trust misplacement. In essence one organisation or stakeholder is not enough for protecting the rights of individuals and the ecosystem needs co-creation of rules between the legal system, technology sector, analytics firms, consumer rights groups, researchers, and the individuals. As researchers and data practitioners, we also need develop new ethics framework to comply with new and evolving legal requirements and technical approaches (e.g., MSR (Bowser & Tsai, 2015)).

Some other challenges exist in the space of ethics-aware and responsible data collection for research on OSNs, including (not exclusively):

- Making secure systems easy to use (for all stakeholders);
- Communicating risks to ordinary users in simple terminology;
- Obtaining *meaningful* consent in a fashion that is accurate and yet not burdensome;
- Achieving nearly-perfect, yet feasible and scalable, anonymisation;
- Coping with legacy systems and technologies;
- Establishing the legal duties and liabilities of the researcher;
- Achieving near perfect reliability for Big and Small Data systems;
- Sustaining data and data protection over 100s of years.

Enabling more stringent requirements and tighter controls on data and processes also imply that people should own their own data. They should be allowed to decide what data are collected about them and to whom this should be shared, and they should be allowed to monetise these data, in similar ways to how supermarket loyalty cards work, with clearly identified parties, for transparent reasons, or through aggregators who turn large amounts of re-identifiable data into statistical information, using the statistical techniques described in the previous section. For instance, someone might be happy to share their financial data with service providers for helping with financial or budgeting assistance, but not to share their health data with these same services. Conversely, they might be willing to share food consumption data to other interested parties, e.g., supermarkets, in return for payment. Availability of a personal cloud, combined with techniques such as differential privacy, and appropriate use of access control and cryptographic techniques can all serve to make this work. There have been recently advances in systems and frameworks, such as OpenPDS (de Montjoye, Shmueli, Wang, & Pentland, 2014) and Databox (Chaudhry et al., 2015), which aim to enable a user-centric approach to personal data use.

In conclusion, there is an urgent need for the research community to consider personal, societal and ethical consequences of large-scale use of social media data in order to perform valuable scientific research without losing the respect and trust of the individuals involved in the ecosystem.

References

- Aebersold, R., Hood, L. E., & Watts, J. D. (2000, April). Equipping scientists for the new biology. *Nature Biotechnology*, 18(4), 359. doi: 10.1038/74325
- An, J., Cha, M., Gummadi, K. P., Crowcroft, J., & Quercia, D. (2012). Visualizing media bias through twitter. In *Proceedings of the 6th international AAAI conference on web and social media (ICWSM)*.
- An, J., Quercia, D., & Crowcroft, J. (2014). Recommending investors for crowdfunding projects. In *Proceedings of the 23rd international conference on world wide web* (pp. 261–270). doi: 10.1145/2566486.2568005
- Ayalon, O., & Toch, E. (2013). Retrospective privacy: Managing longitudinal privacy in online social networks. In *Proceedings of the ninth symposium on usable privacy and security*. New York, NY, USA: ACM. doi: 10.1145/2501604.2501608

- Bailey, M., Dittrich, D., Kenneally, E., & Maughan, D. (2012, March). The Menlo report. *IEEE Security & Privacy Magazine*, 10(2), 71–75. doi: 10.1109/msp.2012.52
- Bauer, L., Cranor, L. F., Komanduri, S., Mazurek, M. L., Reiter, M. K., Sleeper, M., & Ur, B. (2013). The post anachronism: The temporal dimension of Facebook privacy. In *Proceedings of the 12th ACM workshop on workshop on privacy in the electronic society* (pp. 1–12). New York, NY, USA: ACM. doi: 10.1145/2517840.2517859
- Bowser, A., & Tsai, J. Y. (2015). Supporting ethical web research: A new research ethics review. In *Proceedings of the 24th international conference on world wide web* (pp. 151–161). Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee. doi: 10.1145/2736277.2741654
- Buchanan, E., Aycok, J., Dexter, S., Dittrich, D., & Hvizdak, E. (2011, June). Computer science security research and human subjects: Emerging considerations for research ethics boards. *Journal of Empirical Research on Human Research Ethics*, 6(2), 71–83. doi: 10.1525/jer.2011.6.2.71
- Caers, R., De Feyter, T., De Couck, M., Stough, T., Vigna, C., & Du Bois, C. (2013, September). Facebook: A literature review. *New Media & Society*, 15(6), 982–1002. doi: 10.1177/1461444813488061
- Cavallo, D. N., Tate, D. F., Ries, A. V., Brown, J. D., DeVellis, R. F., & Ammerman, A. S. (2012). A social media-based physical activity intervention: a randomized controlled trial. *American journal of preventive medicine*, 43(5), 527–532.
- Cha, M., Benevenuto, F., Haddadi, H., & Gummadi, K. (2012). The world of connections and information flow in Twitter. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 42(4), 991–998. doi: 10.1109/tsmca.2012.2183359
- Cha, M., Haddadi, H., Benevenuto, F., & Gummadi, P. K. (2010). Measuring user influence in Twitter: The million follower fallacy. In *Proceedings of the 4th international AAAI conference on web and social media (ICWSM)*. Retrieved from <http://aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/viewPaper/1538>
- Cha, M., Pérez, J., & Haddadi, H. (2009). Flash floods and ripples: The spread of media content through the blogosphere. In *Proceedings of the 2009 ICWSM data challenge workshop*.
- Chaudhry, A., Crowcroft, J., Haddadi, H., Howard, H., Madhavapeddy, A., McAuley, D., & Mortier, R. (2015). Personal data: Thinking inside the box. *5th decennial Aarhus conferences*. doi: 10.7146/aahcc.v1i1.21312
- Check Hayden, E. (2012, June 20). Informed consent: A broken contract. *Nature*, 486(7403), 312–314. doi: 10.1038/486312a
- Council for International Organizations of Medical Sciences. (2002). *International ethical guidelines for biomedical research involving human subjects*. CIOMS. Retrieved from <http://www.cioms.ch/index.php/printable-publications?task=view&id=48&catid=57>
- Davis, J. H. (2015, July 9). Hacking exposed 21 million in U.S., Government says. *New York Times*. Retrieved from <http://www.nytimes.com/2015/07/10/us/office-of-personnel-management-hackers-got-data-of-millions.html>
- De Choudhury, M., Counts, S., & Horvitz, E. (2013). Social media as a measurement tool of depression in populations. In *Proceedings of the 5th annual ACM web science conference* (pp. 47–56). New York, NY, USA: ACM. doi: 10.1145/2464464.2464480
- De Domenico, M., Lima, A., & Musolesi, M. (2013, December). Interdependence and predictability of human mobility and social interactions. *Pervasive Mob. Comput.*, 9(6), 798–807. doi: 10.1016/j.pmcj.2013.07.008
- De Meo, P., Ferrara, E., Fiumara, G., & Provetti, A. (2014, October). On Facebook, most ties are weak. *Communications of the ACM*, 57(11), 78–84. doi: 10.1145/2629438
- de Montjoye, Y.-A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013, March 25). Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3(1376). doi: 10.1038/srep01376
- de Montjoye, Y.-A., Radaelli, L., Singh, V. K., & Pentland, A. . (2015, January 30). Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, 347(6221), 536–539. doi: 10.1126/science.1256297

- de Montjoye, Y.-A., Shmueli, E., Wang, S. S., & Pentland, A. S. (2014, July 9). openPDS: Protecting the privacy of metadata through safeanswers. *PLoS ONE*, *9*(7), e98790+. doi: 10.1371/journal.pone.0098790
- Díaz, P., Aedo, I., & Herranz, S. (2014). Citizen participation and social technologies: Exploring the perspective of emergency organizations. In C. Hanachi, F. Bénaben, & F. Charoy (Eds.), *Information systems for crisis response and management in Mediterranean countries* (Vol. 196, pp. 85–97). Springer International Publishing. doi: 10.1007/978-3-319-11818-5_8
- Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.* (1995, Nov). OJ L 281 pp.31-50.
- D’Onfro, J. (2016, April 27). *Here’s how much time people spend on Facebook, Instagram, and Messenger every day.* Retrieved from <http://businessinsider.com/how-much-time-do-people-spend-on-facebook-per-day-2016-4>
- Dwork, C. (2006). Differential Privacy. In M. Bugliesi, B. Preneel, V. Sassone, & I. Wegener (Eds.), *Automata, languages and programming* (Vol. 4052, pp. 1–12). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/11787006_1
- EPSRC policy framework on research data.* (n.d.). Retrieved from <http://www.epsrc.ac.uk/index.cfm/about/standards/researchdata/> (Accessed 14 October 2014)
- European Parliament and the Council of the European Union. (1995). Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Union*, *L 281*, 0031–0050. Retrieved from <http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:31995L0046&from=en>
- Faden, R. R., & Beauchamp, T. L. (1986). *A history and theory of informed consent.* Oxford, UK: Oxford University Press.
- Falahrastegar, M., Haddadi, H., Uhlig, S., & Mortier, R. (2014). Anatomy of the third-party web tracking ecosystem. *CoRR*, *abs/1409.1066*. Retrieved from <http://arxiv.org/abs/1409.1066>
- Farrimond, H. (2013). *Doing ethical research.* Basingstoke, UK: Palgrave Macmillan.
- Gentry, C. (2010, March). Computing arbitrary functions of encrypted data. *Communications of the ACM*, *53*(3), 97–105. doi: 10.1145/1666420.1666444
- Gomer, R., Schraefel, M. C., & Gerding, E. (2014). Consenting agents: Semi-autonomous interactions for ubiquitous consent. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing: Adjunct publication* (pp. 653–658). New York, NY, USA: ACM. doi: 10.1145/2638728.2641682
- Granovetter, M. S. (1973, May). The strength of weak ties. *American journal of sociology*, *78*(6), 1360–1380. doi: 10.2307/2776392
- Haddadi, H., & Hui, P. (2010). To add or not to add: privacy and social honeypots. In *Communications workshops (icc), 2010 ieee international conference on* (pp. 1–5).
- Haddadi, H., Hui, P., Henderson, T., & Brown, I. (2011, September). Targeted advertising on the handset: Privacy and security challenges. In J. Müller, F. Alt, & D. Michelis (Eds.), *Pervasive advertising* (pp. 119–137). London: Springer London. doi: 10.1007/978-0-85729-352-7_6
- Haddadi, H., Mortier, R., & Hand, S. (2012, April). Privacy analytics. *ACM SIGCOMM Computer Communication Review*, *42*(2), 94–98. doi: 10.1145/2185376.2185390
- Hon, W. K., & Millard, C. (2013). How do restrictions on international transfers of personal data work in clouds? In C. Millard (Ed.), *Cloud computing law.* Oxford, UK: Oxford University Press. doi: 10.1093/acprof:oso/9780199671670.003.0010
- Hutton, L., & Henderson, T. (2015, May). “I didn’t sign up for this!”: Informed consent in social network research. In *Proceedings of the 9th international AAAI conference on web and social media (ICWSM)* (pp. 178–187).
- Imran, M., Castillo, C., Lucas, J., Meier, P., & Vieweg, S. (2014). AIDR: Artificial intelligence for disaster response. In *Proceedings of the companion publication of the 23rd international conference on World Wide Web* (pp. 159–162). doi: 10.1145/2567948.2577034

- Ioannidis, J. P. A. (2013, March 20). Informed consent, big data, and the oxymoron of research that is not research. *The American Journal of Bioethics*, 13(4), 40–42. doi: 10.1080/15265161.2013.768864
- Ji, S., Li, W., Srivatsa, M., He, J. S., & Beyah, R. (2016, April). General graph data de-anonymization: From mobility traces to social networks. *ACM Trans. Inf. Syst. Secur.*, 18(4). doi: 10.1145/2894760
- Kaye, J., Whitley, E. A., Lund, D., Morrison, M., Teare, H., & Melham, K. (2014, May 07). Dynamic consent: a patient interface for twenty-first century research networks. *European Journal of Human Genetics*, 23(2), 141–146. doi: 10.1038/ejhg.2014.71
- Kocabaş, Ö., & Soyata, T. (2014). Medical data analytics in the cloud using homomorphic encryption. *Handbook of Research on Cloud Infrastructures for Big Data Analytics*, 471.
- Konstas, I., Stathopoulos, V., & Jose, J. M. (2009). On social networks and collaborative recommendation. In *Proceedings of the 32nd international ACM SIGIR conference on research and development in information retrieval* (pp. 195–202). New York, NY, USA: ACM. doi: 10.1145/1571941.1571977
- Kosinski, M., Stillwell, D., & Graepel, T. (2013, April 9). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15), 5802–5805. doi: 10.1073/pnas.1218772110
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014, June 17). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24), 8788–8790. doi: 10.1073/pnas.1320040111
- Lane, N. D., & Georgiev, P. (2015). Can deep learning revolutionize mobile sensing? In *Proceedings of the 16th international workshop on mobile computing systems and applications* (pp. 117–122). New York, NY, USA: ACM. doi: 10.1145/2699343.2699349
- Laney, D. (2001, February 6). *3-D data management: Controlling data volume, velocity and variety* (Tech. Rep. No. 949). META Group, Inc. Retrieved from <http://blogs.gartner.com/doug-laney/deja-vmvue-others-claiming-gartners-volume-velocity-variety-construct-for-big-data/>
- Lindell, Y., & Pinkas, B. (2009). Secure multiparty computation for privacy-preserving data mining. *Journal of Privacy and Confidentiality*, 1(1), 5.
- Luger, E., Moran, S., & Rodden, T. (2013). Consent for all: revealing the hidden complexity of terms and conditions. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 2687–2696). New York, NY, USA: ACM. doi: 10.1145/2470654.2481371
- Luger, E., & Rodden, T. (2013). An informed view on consent for UbiComp. In *Proceedings of the 2013 ACM international joint conference on pervasive and ubiquitous computing* (pp. 529–538). New York, NY, USA: ACM. doi: 10.1145/2493432.2493446
- Machanavajjhala, A., Kifer, D., Gehrke, J., & Venkitasubramaniam, M. (2007, March). L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1). doi: 10.1145/1217299.1217302
- Manson, N. C., & O’Neill, O. (2007). *Rethinking informed consent in bioethics*. Cambridge, UK: Cambridge University Press. doi: 10.1017/cbo9780511814600
- Mayer-Schonberger, V. (2009). *Delete: The virtue of forgetting in the digital age*. Princeton University Press.
- Mejova, Y., Haddadi, H., Noulas, A., & Weber, I. (2015, May). # foodporn: Obesity patterns in culinary interactions. *ACM conference on Digital Health*. doi: 10.1145/2750511.2750524
- Morrison, A., McMillan, D., & Chalmers, M. (2014, October). Improving consent in large scale mobile HCI through personalised representations of data. In *Proceedings of the 8th Nordic conference on human-computer interaction: Fun, fast, foundational* (pp. 471–480). New York, NY, USA: ACM. doi: 10.1145/2639189.2639239
- Mortier, R., Haddadi, H., Henderson, T., McAuley, D., & Crowcroft, J. (2014, October 1). Human-Data Interaction: The human face of the data-driven society. *SSRN*. doi: 10.2139/ssrn.2508051
- Munteanu, C., Molyneaux, H., Moncur, W., Romero, M., O’Donnell, S., & Vines, J. (2015,

- April). Situational ethics: Re-thinking approaches to formal ethics requirements for human-computer interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 105–114). doi: 10.1145/2702123.2702481
- Narayanan, A., & Felten, E. W. (2014). No silver bullet: De-identification still doesn't work. *White Paper*.
- Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *Proceedings of the IEEE symposium on security and privacy* (pp. 111–125). doi: 10.1109/SP.2008.33
- Neuhaus, F., & Webmoor, T. (2012, February). Agile ethics for massified research and visualization. *Information, Communication & Society*, 15(1), 43–65. doi: 10.1080/1369118x.2011.616519
- Nissenbaum, H. F. (2004, February). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–157.
- NSF policy on dissemination and sharing of research results*. (n.d.). Retrieved from <http://www.nsf.gov/pubs/policydocs/pappguide/nsf13001/aag-6.jsp#VID4> (Accessed 14 October 2014)
- Owen, R., Macnaghten, P., & Stilgoe, J. (2012, December). Responsible research and innovation: From science in society to science for society, with society. *Science and Public Policy*, 39(6), 751–760. doi: 10.1093/scipol/scs093
- Paul, M. J., & Dredze, M. (2011). You are what you tweet: Analyzing Twitter for public health. In *Proceedings of the 5th international AAAI conference on web and social media (ICWSM)*. Retrieved from <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2880>
- Rawls, J. (1999). *A theory of justice* (Revised ed.). Cambridge, MA, USA: Harvard University Press.
- Sala, A., Zhao, X., Wilson, C., Zheng, H., & Zhao, B. Y. (2011, November). Sharing graphs using differentially private graph models. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)* (pp. 81–98). doi: 10.1145/2068816.2068825
- Schroeder, R. (2014). Big Data and the brave new world of social media research. *Big Data & Society*, 1(2). doi: 10.1177/2053951714563194
- Selwyn, N. (2009). Faceworking: exploring students' education-related use of facebook. *Learning, Media and Technology*, 34(2), 157–174. doi: 10.1080/17439880902923622
- Seth, A., & Zhang, J. (2008). A Social Network Based Approach to Personalized Recommendation of Participatory Media Content. In *Proceedings of the 2nd international AAAI conference on web and social media (ICWSM)* (pp. 109–117). Seattle.
- Shamoo, A. E., & Resnik, D. B. (2009). *Responsible conduct of research*. Oxford, UK: Oxford University Press. doi: 10.1093/acprof:oso/9780195368246.001.0001
- Sicker, D. C., Ohm, P., & Gunaji, S. (2007, Spring). The analog hole and the price of music: An empirical study. *Journal on Telecommunications and High Technology Law*, 5(3), 573–588. Retrieved from http://www.jthtl.org/content/articles/V5I3/JTHTLv5i3_SickerOhmGunaji.PDF
- Silvertown, J. (2009, September). A new dawn for citizen science. *Trends in Ecology & Evolution*, 24(9), 467–471. doi: 10.1016/j.tree.2009.03.017
- Singer, P. (2011). *Practical ethics* (3rd ed.). Cambridge, UK: Cambridge University Press.
- Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., & Cranor, L. F. (2013). The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the 2013 conference on computer supported cooperative work* (pp. 793–802). New York, NY, USA: ACM. doi: 10.1145/2441776.2441865
- Solberg, L. (2010). Data mining on Facebook: A free space for researchers or an IRB nightmare? *University of Illinois Journal of Law, Technology & Policy*, 2010(2). Retrieved from <http://www.jltp.uiuc.edu/works/Solberg.htm>
- Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05), 557–570. doi: 10.1142/S0218488502001648

- Toch, E. (2014). Crowdsourcing privacy preferences in context-aware applications. *Personal and Ubiquitous Computing*, 18(1), 129–141. doi: 10.1007/s00779-012-0632-0
- Torres, W. A. A., Bhattacharjee, N., & Srinivasan, B. (2014). Effectiveness of fully homomorphic encryption to preserve the privacy of biometric data. In *Proceedings of the 16th international conference on information integration and web-based applications & services* (pp. 152–158). New York, NY, USA: ACM. doi: 10.1145/2684200.2684296
- Viswanath, B., Mislove, A., Cha, M., & Gummadi, K. P. (2009). On the evolution of user interaction in Facebook. In *Proceedings of the 2nd acm workshop on online social networks*. doi: 10.1145/1592665.1592675
- Widener, M. J., & Li, W. (2014). Using geolocated Twitter data to monitor the prevalence of healthy and unhealthy food references across the US. *Applied Geography*, 54, 189–197. doi: 10.1016/j.apgeog.2014.07.017
- Wright, J., Souza, T. D., & Brown, I. (2011). Fine-grained censorship mapping: Information sources, legality and ethics. In *Proceedings of the USENIX workshop on free and open communications on the internet*.
- Zhao, Y., Ye, J., & Henderson, T. (2014, December). Privacy-aware location privacy preference recommendations. In *Proceedings of the 11th international conference on mobile and ubiquitous systems: Computing, networking and services* (pp. 120–129). Brussels, Belgium: ICST. doi: 10.4108/icst.mobiquitous.2014.258017
- Zhao, Y., Ye, J., & Henderson, T. (2016, March). The effect of privacy concerns on privacy recommenders. In *Proceedings of the 21st international conference on intelligent user interfaces* (pp. 218–227). New York, NY, USA: ACM. doi: 10.1145/2856767.2856771
- Zimmer, M. (2010, December). “But the data is already public”: on the ethics of research in Facebook. *Ethics and Information Technology*, 12(4), 313–325. doi: 10.1007/s10676-010-9227-5