Beyond the EULA: Improving consent for data mining

Luke Hutton and Tristan Henderson

Abstract Companies and academic researchers may collect, process, and distribute large quantities of personal data without the explicit knowledge or consent of the individuals to whom the data pertains. Existing forms of consent often fail to be appropriately readable and ethical oversight of data mining may not be sufficient. This raises the question of whether existing consent instruments are sufficient, logistically feasible, or even necessary, for data mining. In this chapter, we review the data collection and mining landscape, including commercial and academic activities, and the relevant data protection concerns, to determine the types of consent instruments used. Using three case studies, we use the new paradigm of human-data interaction to examine whether these existing approaches are appropriate. We then introduce an approach to consent that has been empirically demonstrated to improve on the state of the art and deliver meaningful consent. Finally, we propose some best practices for data collectors to ensure their data mining activities do not violate the expectations of the people to whom the data relate.

1 Introduction

The ability of companies to collect, process, and distribute large quantities of personal data, and to further analyse, mine and generate new data based on inferences from these data, is often done without the explicit knowledge or consent of the individuals to whom the data pertains. Consent instruments such as privacy notices or End User License Agreements (EULAs) are widely deployed, often presenting

Luke Hutton
Centre for Research in Computing, The Open University, Milton Keynes, MK7 6AA, UK e-mail: luke.hutton@open.ac.uk

Tristan Henderson
School of Computer Science, University of St Andrews, St Andrews KY16 9SX, UK e-mail: tnhh@st-andrews.ac.uk

individuals with thousands of words of legal jargon that they may not read nor comprehend, before soliciting agreement in order to make use of a service. Indeed, even if an individual does have a reasonable understanding of the terms to which they have agreed, such terms are often carefully designed to extend as much flexibility to the data collector as possible to obtain even more data, distribute them to more stakeholders, and make inferences by linking data from multiple sources, despite no obvious agreement to these new practices.

The lack of transparency behind data collection and mining practices threatens the agency and privacy of data subjects, with no practical way to control these invisible data flows, nor correct misinformation or inaccurate and inappropriate inferences derived from linked data. Existing data protection regimes are often insufficient as they are predicated on the assumption that an individual is able to detect when a data protection violation has occurred in order to demand recourse, which is rarely the case when data are opaquely mined at scale.

These challenges are not unique to commercial activities, however. Academic researchers often make use of datasets containing personal information, such as those collected from social network sites or devices such as mobile phones or fitness trackers. Most researchers are bound by an obligation to seek ethical approval from an institutional review board (IRB) before conducting their research. The ethical protocols used, however, are inherited from post-war concerns regarding biomedical experiments, and may not be appropriate for Internet-mediated research, where millions of data points can be collected without any personal interventions. This raises the question of whether existing consent instruments are sufficient, logistically feasible, or even necessary, for research of this nature.

In this chapter we first review the data collection and mining landscape, including commercial and academic activities, and the relevant data protection laws, to determine the types of consent instruments used. Employing the newly-proposed paradigm of Human-Data Interaction, we examine three case studies to determine whether these mechanisms are sufficient to uphold the expectations of individuals, to provide them with sufficient agency, legibility and negotiability, and whether privacy norms are violated by secondary uses of data which are not explicitly sanctioned by individuals. We then discuss various new dynamic and contextual approaches to consent, which have been empirically demonstrated to improve on the state of the art and deliver meaningful consent. Finally, we propose some best practices that data collectors can adopt to ensure their data mining activities do not violate the expectations of the people to whom the data relate.

2 Background

Data mining is the statistical analysis of large-scale datasets to extract additional patterns and trends [16]. This has allowed commercial, state, and academic actors to answer questions which have not previously been possible, due to insufficient data, analytical techniques, or computational power. Data mining is often characterised by the use of aggregate data to identify traits and trends which allow the identification and characterisation of clusters of people rather than individuals, associations between events, and forecasting of future events. As such, it has been used in a number of real-world scenarios such as optimising the layout of retail stores, attempts to identify disease trends, and mass surveillance. Many classical data mining and knowledge discovery applications involve businesses or marketing [12], such as clustering consumers into groups and attempting to predict their behaviour. This may allow a business to understand their customers and target promotions appropriately. Such profiling can, however, be used to characterise individuals for the purpose of denying service when extending credit, leasing a property, or acquiring insurance. In such cases, the collection and processing of sensitive data can be invasive, with significant implications for the individual, particularly where decisions are made on the basis of inferences that may not be accurate, and to which the individual is given no right of reply. This has become more important of late, as more recent data mining applications involve the analysis of personal data, much of which is collected by individuals and contributed to marketers in what has been termed “self-surveillance” [25]. Such personal data have been demonstrated to be highly valuable [48], and have even been described as the new “oil” in terms of the value of their resource [55]. Value aside, such data introduce new challenges for consent as they can often be combined to create new inferences and profiles where previously data would have been absent [17].

Table 1: Some relevant EU legislation that may apply to data mining activities

Legislation	Some relevant sections
Data Protection Directive 95/46/EC	Data processing (Art 1), fair processing (Art 6(1)), purpose limitation (Art 6(2)), proportionality (Art 6(3)), consent (Art 7(1)), sensitive data (Art 8)
E-Privacy Directive 2002/58/EC as amended by 2009/136/EC	cookies (Art 5), traffic data (Art 6), location data (Art 9), unsolicited communication (Art 13)
General Data Protection Regulation 2016/679	Consent (Art 7), right to be forgotten (Art 17), right to explanation (Art 22), privacy by design (Art 25)

Data mining activities are legitimised through a combination of legal and self-regulatory behaviours. In the European Union, the Data Protection Directive [10], and the forthcoming General Data Protection Regulation (GDPR) that will succeed it in 2018 [11] govern how data mining can be conducted legitimately. The e-Privacy Directive also further regulates some specific aspects of data mining such as cookies (Table 1). In the United States, a self-regulatory approach is generally preferred, with the Federal Trade Commission offering guidance regarding privacy protections [42], consisting of six core principles, but lacking the coverage or legal backing of the EU’s approach.

Under the GDPR, the processing of personal data for any purpose, including data mining, is subject to explicit opt-in consent from an individual, prior to which the data controller must explicitly state what data are collected, the purpose of processing them, and the identity of any other recipients of the data. Although there are a number of exceptions, consent must generally be sought for individual processing activities, and cannot be broadly acquired *a priori* for undefined future uses, and there are particular issues with data mining, transparency and accountability [7]. Solove [47] acknowledges these regulatory challenges, arguing that paternalistic approaches are not appropriate, as these deny people the freedom to consent to particular beneficial uses of their data. The timing of consent requests and the focus of these requests need to be managed carefully; such thinking has also become apparent in the GDPR.[1] The call for dynamic consent is consistent with Nissenbaum's model of contextual integrity [39], which posits that all information exchanges are subject to context-specific norms, which governs to whom and for what purpose information sharing can be considered appropriate. When the context is disrupted, perhaps by changing with whom data are shared, or for what purpose, privacy violations can occur when this is not consistent with the norms of the existing context. Therefore, consent can help to uphold contextual integrity by ensuring that if the context is perturbed, consent is renegotiated, rather than assumed.

Reasoning about how personal data are used has resulted in a new paradigm, *human-data interaction*, which places humans at the centre of data flows and provides a framework for studying personal data collection and use according to three themes [35]:

- *Legibility*: Often, data owners are not aware that data mining is even taking place. Even if they are, they may not know what is being collected or analysed, the purpose of the analysis, or the insights derived from it.
- *Agency*: The opaque nature of data mining often denies data owners agency. Without any engagement in the practice, people have no ability to provide meaningful consent, if they are asked to give consent at all, nor correct flawed data or review inferences made based on their data.
- *Negotiability*: The context in which data are collected and processed can often change, whether through an evolving legislative landscape, data being traded between organisations, or through companies unilaterally changing their privacy policies or practices. Analysis can be based on the linking of datasets derived from different stakeholders, allowing insights that no single provider could make. This is routinely the case in profiling activities such as credit scoring. Even where individuals attempt to obfuscate their data to subvert this practice, it is often possible to re-identify them from such linked data [38]. Data owners should have the ability to review how their data are used as circumstances change in order to uphold contextual integrity.

Early data protection regulation in the 1980s addressed the increase in electronic data storage and strengthened protections against unsolicited direct marketing [49]. Mail order companies were able to develop large databases of customer details to

[1] e.g., Article 7(3) which allows consent to be withdrawn, and Article 17 on the "right to be forgotten" which allows inferences and data to be erased.

enable direct marketing, or the trading of such information between companies. When acquiring consent for the processing of such information became mandatory, such as under the 1984 Data Protection Act in the UK, this generally took the form of a checkbox on paper forms, where a potential customer could indicate their willingness for secondary processing of their data. As technology has evolved away from mail-in forms being the primary means of acquiring personal information, and the scope and intent of data protection moves from regulating direct marketing to a vast range of data-processing activities, there has been little regulatory attention paid to how consent is acquired. As such, consent is often acquired by asking a user to tick a checkbox to opt-in or out of secondary use of their data. This practice is well-entrenched, where people are routinely asked to agree to an End-User Licence Agreement (EULA) before accessing software, and multiple terms of service and privacy policies before accessing online services, generally consisting of a long legal agreement and an “I Agree” button.

A significant body of research concludes that such approaches to acquiring consent are flawed. Luger et al. find that the terms and conditions provided by major energy companies are not sufficiently readable, excluding many from being able to make informed decisions about whether they agree to such terms [29]. Indeed, Obar and Oeldorf-Hirsch find that the vast majority of people do not even read such documents [40], with all participants in a user study accepting terms including handing over their first-born child to use a social network site. McDonald and Cranor measure the economic cost of reading lengthy policies [31], noting the inequity of expecting people to spend an average of ten minutes of their time reading and comprehending a complex document in order to use a service. Freidman et al. caution that simply including more information and more frequent consent interventions can be counter-productive, by frustrating people and leading them to making more complacent consent decisions [13].

Academic data mining is subject to a different regulatory regime, with fewer constraints over the secondary use of data from a data protection perspective. This is balanced by an ethical review regime, rooted in post-war concern over a lack of ethical rigour in biomedical research. In the US, ethical review for human subjects research via an institutional review board (IRB) is necessary to receive federal funding, and the situation is similar in many other countries. One of the central tenets of ethical human research is to acquire informed consent before a study begins [5]. As such, institutions have developed largely standardised consent instruments [1] which researchers can use to meet these requirements. While in traditional lab-based studies, these consent procedures can be accompanied by an explanation of the study from a researcher, or the opportunity for a participant to ask any questions, this affordance is generally not available in online contexts, effectively regressing to the flawed EULAs discussed earlier.

Some of these weaknesses have been examined in the literature. Hamnes et al. find that consent documents in rheumatological studies are not sufficiently readable for the majority of the population [15], a finding which is supported by Vučemilo and Borovečki who also find that medical consent forms often exclude important information [53]. Donovan-Kicken et al. examine the sources of confusion when reviewing such documents [8], which include insufficient discussion of risk and lengthy or overly complex language. Munteanu et al. examine the ethics approval process in a number of HCI research case studies, finding that participants often agreed to consent instruments they have not read or understood, and the rigidity of such processes can often be at odds with such studies where a "situational interpretation" of an agreed protocol is needed [36]. There also lacks agreement among researchers about how to conduct such research in an ethical manner, with Vitak et al. finding particular variability regarding whether data should be collected at large scale without consent, or if acquiring consent in such cases is even possible [52].

Existing means of acquiring consent are inherited from a time when the scope of data collection and processing was perhaps constrained and could be well understood. Now, even when the terms of data collection and processing are understood as written, whether registering for an online service, or participating in academic research, it is not clear that the form of gaining the consent was meaningful, or sufficient. Someone may provide consent to secondary use of their data, without knowing what data this constitutes, who will be able to acquire it, for what purpose, or when. This is already a concern when considering the redistribution and processing of self-disclosed personally identifiable information, but becomes increasingly complex when extended to historical location data, shopping behaviours, or social network data, much of which are not directly provided by the individual, and are nebulous in scale and content. Moreover concerns may change over time (the so-called "privacy paradox" [3] that has been demonstrated empirically [2, 4]), which may require changes to previously-granted consent.

Returning to our three themes of *legibility*, *agency*, and *negotiability*, we can see that:

- Existing EULAs and consent forms may not meet a basic standard of *legibility*, alienating significant areas of the population from understanding what they are being asked to agree to. Furthermore, the specific secondary uses of their data are often not explained.
- EULAs and consent forms are often only used to secure permission once, then often never again, denying people *agency* to revoke their consent when a material change in how their data are used arises.
- Individuals have no power to meaningfully *negotiate* how their data are used, nor to intelligently adopt privacy-preserving behaviours, as they generally do not know which data attributed to them is potentially risky.

3 Case studies

In this section we examine a number of real-world case studies to identify instances where insufficient consent mechanisms were employed, failing to provide people with legibility, agency, and negotiability.

3.1 Taste, Ties, and Time

In 2006, researchers at Harvard University collected a dataset of Facebook profiles from a cohort of undergraduate students, named “Tastes, Ties and Time” (T3) [28]. At the time, Facebook considered individual universities to comprise networks where members of the institution could access the full content of each other’s profiles, despite not having an explicit friendship with each other on the service. This design was exploited by having research assistants at the same institution manually extract the profiles of each member of the cohort.

Subsequently, an anonymised version of the dataset was made publicly available, with student names and identifiers removed, and terms and conditions for downloading the dataset made it clear that deanonymisation was not permitted. Unfortunately, this proved insufficient, with aggregate statistics such as the size of the cohort making it possible to infer the college the dataset was derived from, and as some demographic attributes were only represented by a single student, it was likely that individuals could be identified [56].

Individuals were not aware that the data collection took place, and did not consent to its collection, processing, nor subsequent release. As such, this case falls short in our themes for acceptable data-handling practices:

- **Legibility**: Individuals were not aware their data was collected or subsequently released. With a tangible risk of individuals being identified without their knowledge, the individual is not in a position to explore any legal remedies to hold Facebook or the researchers responsible for any resulting harms. In addition, even if consent were sought, it can be difficult for individuals to conceptualise exactly which of their data would be included, considering the large numbers of photos, location traces, status updates, and biographical information a typical user might accrue over years, without an accessible means of visualising or selectively disclosing these data.
- **Agency**: Without notification, individual users had no way to opt-out of the data collection, nor prevent the release of their data. As a side-effect of Facebook’s university-only network structure at the time, the only way for somebody to avoid their data being used in such a manner was to leave these institution networks, losing much of the utility of the service in the process. This parallels Facebook’s approach to releasing other products, such as the introduction of News Feed in 2006. By broadcasting profile updates to one’s network, the effective visibility of their data was substantially increased, with no way to opt-out of the feature without leaving the service entirely. This illusory loss of control was widely criticised at the time [20].
- **Negotiability**: In this respect, the user’s relationship with Facebook itself is significant. In addition to IRB approval, the study was conducted with Facebook’s permission [28], but Facebook’s privacy policy at the time did not allow for Facebook to share their data with the researchers.[2] Therefore, the existing context

[2] Facebook Privacy Policy, February 2006: `http://web.archive.org/web/20060406105119/http://www.facebook.com/policy.php`

for sharing information on Facebook was disrupted by this study. This includes the normative expectation that data are shared with Facebook for the purpose of sharing information with one's social network, and not myriad third parties. In addition, no controls were extended to the people involved to prevent it from happening, or to make a positive decision to permit this new information-sharing context.

3.2 Facebook emotional contagion experiment

In 2012, researchers at Facebook and Cornell University conducted a large-scale experiment on 689,003 Facebook users. The study manipulated the presentation of stories in Facebook's News Feed product, which aggregates recent content published by a user's social network, to determine whether biasing the emotional content of the news feed affected the emotions that people expressed in their own disclosures [27].

While the T3 study highlighted privacy risks of nonconsensual data sharing, the emotional contagion experiment raises different personal risks from inappropriate data mining activities. For example, for a person suffering from depression, being subjected to a news feed of predominantly depressive-indicative content could have catastrophic consequences, particularly considering the hypothesis of the experiment that depressive behaviour would increase under these circumstances. Considering the scale at which the experiment was conducted, there was no mechanism for excluding such vulnerable people, nor measuring the impact on individuals to mitigate such harms. Furthermore, as the study was not age-restricted, children may have unwittingly been subjected to the study [19]. Rucuber notes that the harms to any one individual in such experiments can be masked by the scale of the experiment [43].

Beyond the research context, this case highlights the broader implications of the visibility of media, whether socially-derived or from mainstream media, being algorithmically controlled. Napoli argues that this experiment highlights Facebook's ability to shape public discourse by altering the news feed's algorithm to introduce political bias [37], without any governance to ensure that such new media are acting in the public interest. The majority of Facebook users do not know that such filtering happens at all, and the selective presentation of content from one's social network can cause social repercussions where the perception is that individuals are withholding posts from someone, rather than an algorithmic intervention by Facebook [9].

This case shows one of the greater risks of opaque data mining. Where people are unaware such activities are taking place, they lose all power to act autonomously to minimise the risk to themselves, even putting aside the responsibility of the researchers in this instance. We now consider how this case meets our three core themes:

- **Legibility**: Individuals were unaware that they were participants in the research. They would have no knowledge or understanding of the algorithms which choose which content is presented on the news feed, and how they were altered for this experiment, nor that the news feed is anything other than a chronological collection of content provided by their social network. Without this insight, the cause of a perceptible change in the emotional bias in the news feed can not be reasoned. Even if one is aware that the news feed is algorithmically controlled, without knowing which data are collected or used in order to determine the relevance of individual stories, it is difficult to reason why certain stories are displayed.

- **Agency**: As in the T3 case, without awareness of the experiment being conducted, individuals were unable to provide consent, nor opt-out of the study. Without an understanding of the algorithms which drive the news feed, nor how they were adjusted for the purposes of this experiment, individuals are unable to take actions, such as choosing which information to disclose or hide from Facebook in an effort to control the inferences Facebook makes, nor to correct any inaccurate inferences. At the most innocuous level, this might be where Facebook has falsely inferred a hobby or interest, and shows more content relating to that. Of greater concern is when Facebook, or the researchers in this study, are unable to detect when showing more depressive-indicative content could present a risk.
- **Negotiability**: In conducting this study, Facebook unilaterally changed the relationship its users have with the service, exploiting those who are unable to control how their information is used [45]. At the time of the study, the Terms of Service to which users agree when joining Facebook did not indicate that data could be used for research purposes [19], a clause which was added after the data were collected. As a commercial operator collaborating with academic researchers, the nature of the study was ambiguous, with Facebook having an internal product improvement motivation, and Cornell researchers aiming to contribute to generalisable knowledge. Cornell's IRB deemed that they did not need to review the study because Facebook provided the data, [3], but the ethical impact on the unwitting participants is not dependent on who collected the data. As Facebook has no legal requirement to conduct an ethical review of their own research, and without oversight from the academic collaborators, these issues did not surface earlier. Facebook has since adopted an internal ethics review process [24], however it makes little reference to mitigating the impact on participants, and mostly aims to maximise benefit to Facebook. Ultimately, these actions by researchers and institutions with which individuals have no prior relationship serve to disrupt the existing contextual norms concerning people's relationship with Facebook, without extending any ability to renegotiate this relationship.

3.3 NHS sharing data with Google

In February 2016, the Google subsidiary DeepMind announced a collaboration with the National Health Service's Royal Free London Trust to build a mobile application

[3] Cornell statement: `https://perma.cc/JQ2L-TEXQ`

titled Streams to support the detection of acute kidney injury (AKI) using machine learning techniques. The information sharing agreement permitting this collaboration gives DeepMind ongoing and historical access to all identifiable patient data collected by the trust's three hospitals [22].

While the project is targeted at supporting those at risk of AKI, data relating to all patients are shared with DeepMind, whether they are at risk or not. There is no attempt to gain the consent of those patients, or to provide an obvious opt-out mechanism. The trust's privacy policy only allows data to be shared without explicit consent for "your direct care". [4] Considering that Streams is only relevant to those being tested for kidney disease, it follows that for most people, their data are collected and processed without any direct care benefit [21], in violation of this policy. Given the diagnostic purpose of the app, such an application could constitute a medical device, however no regulatory approval was sought by DeepMind or the trust [21].

Permitting private companies to conduct data mining within the medical domain disrupts existing norms, by occupying a space that lies between direct patient care and academic research. Existing ethical approval and data-sharing regulatory mechanisms have not been employed, or are unsuitable for properly evaluating the potential impacts of such work. By not limiting the scope of the data collection nor acquiring informed consent, there is no opportunity for individuals to protect their data. In addition, without greater awareness of the collaboration, broader public debate about the acceptability of the practice is avoided, which is of importance considering the sensitivity of the data involved. Furthermore, this fairly limited collaboration can normalise a broader sharing of data in the future, an eventuality which is more likely given an ongoing strategic partnership forged between DeepMind and the trust [21]. We now consider this case from the perspective of our three themes:

- **Legibility**: Neither patients who could directly benefit from improved detection of AKI, nor all other hospital patients, were aware that their data were being shared with DeepMind. Indeed, this practice in many data mining activities – identifying patterns to produce insight from myriad data – risks violating a fundamental principle of data protection regulation; that of proportionality [11].
- **Agency**: The NHS collects data from its functions in a number of databases, such as the Secondary Uses Service which provides a historical record of treatments in the UK, and can be made available to researchers. Awareness of this database and its research purpose is mostly constrained to leaflets and posters situated in GP practices. If patients wish to opt-out of their data being used they must insist on it, and are likely to be reminded of the public health benefit and discouraged from opting-out [6]. Without being able to assume awareness of the SUS, nor individual consent being acquired, it is difficult for individuals to act with agency. Even assuming knowledge of this collaboration, it would require particular understanding of the functions of the NHS to know that opting-out of the SUS would limit historical treatment data made available to DeepMind. Even where someone is willing to share their data to support their direct care, they may wish

[4] Royal Free London Trust Privacy Statement: `https://perma.cc/33YE-LYPF`

to redact information relating to particularly sensitive diagnoses or treatments, but have no mechanism to do so.
- **Negotiability**: The relationship between patients and their clinicians is embodied in complex normative expectations of confidentiality which are highly context-dependent [44]. Public understanding of individual studies is already low [6], and the introduction of sophisticated data mining techniques into the diagnostic process which existing regulatory mechanisms are not prepared for disrupts existing norms around confidentiality and data sharing. The principle of negotiability holds that patients should be able to review their willingness to share data as their context changes, or the context in which the data are used. Existing institutions are unable to uphold this, and the solution may lie in increased public awareness and debate, and review of policy and regulatory oversight to reason a more appropriate set of norms.

How each of these case studies meets the principles of legibility, agency, and negotiability is summarised in Table 2.

Case study	**HDI principle**	**How was principle violated?**
Taste, Ties and Time	Legibility	Individuals unaware of data collection
	Agency	No way to opt-out of data collection
	Negotiability	Data collection violated normative expectation with Facebook
Facebook emotional contagion experiment	Legibility	Individuals unaware they were participants
	Agency	No way to opt-out of participation
	Negotiability	Research not permitted by Facebook's terms
NHS sharing data with Google	Legibility	Patients unaware of data sharing
	Agency	Poor awareness of secondary uses of data and difficulty of opting out
	Negotiability	Data mining can violate normative expectations of medical confidentiality

Table 2: Summary of how the three case studies we examine violate the principles of human-data interaction

4 Alternative consent models

In Section 2 we discussed some of the shortcomings with existing means of acquiring consent for academic research and commercial services including data mining, and discussed three case studies in Section 3. Many of the concerns in these case studies revolved around an inability to provide or enable consent on the part of participants. We now discuss the state-of-the-art in providing meaningful consent for today's data-mining activities.

The acquisition of informed consent can broadly be considered to be *secured* or *sustained* in nature [30]. Secured consent encompasses the forms we discussed in Section 2, where consent is gated by a single EULA or consent form at the beginning of the data collection process and not revisited. Conversely, sustained consent involves ongoing reacquisition of consent over the period that the data are collected or used. This might mean revisiting consent when the purpose of the data collection or processing has changed, such as if data are to be shared with different third parties, or if the data subject's context has changed. Each interaction can be viewed as an individual consent *transaction* [32]. In research, this can also mean extending more granular control to participants over which of their data are collected, such as in Sleeper et al.'s study into self-censorship behaviours on Facebook, where participants could choose which status updates they were willing to share with researchers [46]. This approach has a number of advantages. Gaining consent after the individual has had experience with a particular service or research study may allow subjects to make better-informed decisions than a sweeping form of secured consent. Furthermore, sustained consent can allow participants to make more granular decisions about what they would be willing to share, with a better understanding of the context, rather than a single consent form or EULA being considered *carte blanche* for unconstrained data collection.

The distinction between secured and sustained consent reveals a tension between two variables: *burden* – the time spent and cognitive load required to negotiate the consent process, and *accuracy* – the extent to which the effect of a consent decision corresponds with a person's expectations. While a secured instrument such as a consent form minimises the burden on the individual, with only a single form to read and comprehend, the accuracy is impossible to discern, with no process for validating that consent decision in context, nor to assess the individual's comprehension of what they have agreed to. Conversely, while a sustained approach – such as asking someone whether they are willing for each item of personal data to be used for data mining activities – may improve accuracy, the added burden is significant and can be frustrating, contributing to attrition, which is particularly problematic in longitudinal studies [18].

In some domains other than data mining, this distinction has already been applied. In biomedical research, the consent to the use of samples is commonly distinguished as being broad or dynamic. Broad consent allows samples to be used for a range of experiments within an agreed framework without consent being explicitly required [50], whereas dynamic consent involves ongoing engagement with participants, allowing them to see how their samples are used, and permitting renegotiation of consent if the samples are to be used for different studies, or if the participant's wishes change [26]. Despite the differences from the data mining domain, the same consent challenges resonate.

Various researchers have proposed ways of minimising the burden of consent, while simultaneously collecting meaningful and accurate information from people.

Williams et al. look at sharing medical data, enhancing agency with a dynamic consent model that enables control of data electronically, and improved legibility by providing patients with information about how their data are used [54]. Gomer et al. propose the use of agents who make consent decisions on behalf of individuals to reduce the burden placed on them, based on preferences they have expressed which are periodically reviewed [14]. Moran et al. suggest that consent can be negotiated in multi-agent environments by identifying interaction patterns to determine appropriate times to acquire consent [33].

We have discussed legibility as an important aspect of HDI, and Morrison et al. study how to visualise collected data to research participants [34]. Personalised visualisations led participants in an empirical study to exit the study earlier, which might mean that secured consent was leading participants to continue beyond the appropriate level of data collection. In much earlier work, Patrick looked at presenting user agreements contextually (rather than at the beginning of a transaction as in secured consent) and developed a web-based widget to do so [41].

Such dynamic approaches to consent are not universally supported. Steinsbekk et al. suggest that where data are re-used for multiple studies, there is no need to acquire consent for each one where there are no significant ethical departures because of the extra burden, arguing that it puts greater responsibility on individuals to discern whether a study is ethically appropriate than existing governance structures [50].

In previous work, we have developed a method for acquiring consent which aims to maximise accuracy and minimise burden, satisfying both requirements, bringing some of the principles of dynamic consent to the data mining domain [23], while aiming to maintain contextual integrity by respecting prevailing social norms. While many of the consent approaches discussed in this chapter may satisfy a legal requirement, it is not clear that this satisfies the expectations of individuals or society at large, and thus may violate contextual integrity.

In a user study, we examine whether prevailing norms representing willingness to share specific types of Facebook data with researchers, along with limited data about an individual's consent preferences, can be used to minimise burden and maximise accuracy. The performance of these measures were compared to two controls: one for accuracy and for burden. In the first instance, a sustained consent process which gains permission for the use of each individual item of Facebook data would maximise accuracy while pushing great burden on to the individual. Secondly, a single consent checkbox minimises burden, while also potentially minimising the accuracy of the method. The contextual integrity method works similarly to this approach, by asking a series of consent questions until the individual's conformity to the social norm can be inferred, at which point no more questions are asked.

For 27.7% of participants, this method is able to achieve high accuracy of 96.5% while reducing their burden by an average of 41.1%. This is highlighted in Figure 1, showing a cluster of norm-conformant participants achieving high accuracy and low burden. This indicates that for this segment of the population, the contextual integrity approach both improves accuracy and reduces burden compared to their respective controls. While this does indicate the approach is not suitable for all people,

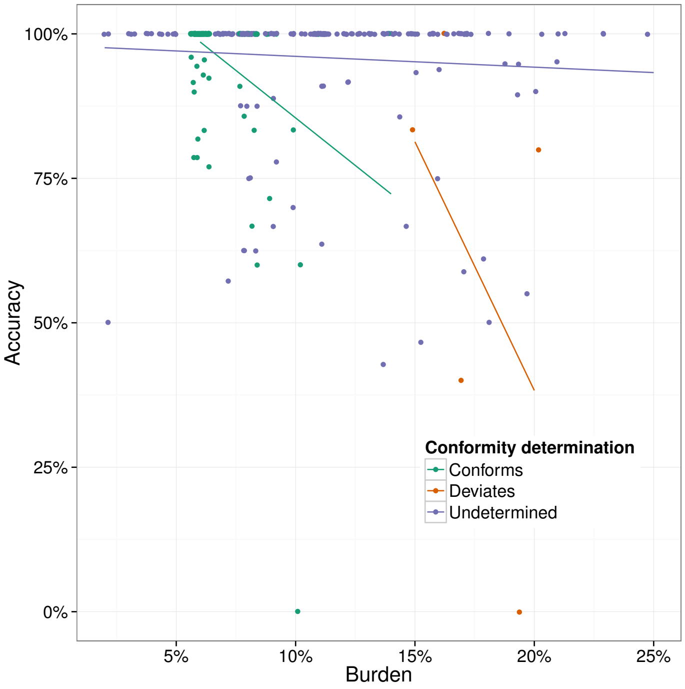

Fig. 1: Scatterplot showing the relationship between norm conformity and accuracy. As indicated by the cluster after 5 questions (5% burden), high accuracy can be preserved when norm conformity is detected quickly, although the technique is not useful for people who are highly norm deviant. Note that the points have been jittered to improve readability. [23]

norm conformity is able to be quickly determined within six questions. Where one does not conform to the norm, the sustained approach can be automatically used as a fallback, which maintains accuracy at the cost of a greater time burden on the individual. Even in less optimal cases, the technique can reduce the burden by an average of 21.9%.

While the technique assessed in this user study is prototypical in its nature, it highlights the potential value of examining alternative means of acquiring consent, which has seen little innovation in both academic and commercial domains. Moreover, while this technique is not universally applicable, this only highlights that the diversity of perspectives, willingness to engage, and ability to comprehend consent language requires a plurality of approaches.

5 Discussion

In this chapter we have illustrated how data mining activities, in both academic and commercial contexts, are often opaque by design. Insufficient consent mechanisms can prevent people from understanding what they are agreeing to, particularly where the scope of the data collected or with whom it is shared is changed without consent being renegotiated. Indeed, as in our three case studies, consent is often not sought at all.

We have considered the impacts of opaque data mining in terms of legibility, agency, and negotiability. We now propose some best practices for conducting data mining which aim to satisfy these three themes.

5.1 Legibility

In order to make data mining more acceptable, it is not sufficient to simply make processes more transparent. Revealing the algorithms, signals, and inferences may satisfy a particularly technically competent audience, but for most people does not help them understand what happens to their data, in order to make an informed decision over whether to consent, or how they can act with any agency.

The incoming General Data Protection Regulation (GDPR) in the European Union requires consent language to be concise, transparent, intelligible and easily accessible [11], which as indicated in the literature, is currently not a universal practice. As highlighted in our three case studies, the absence of any meaningful consent enabling data to be used beyond its original context, such as a hospital or social network site, is unacceptable. Even without adopting more sophisticated approaches to consent as discussed in Section 4, techniques to notify and reacquire consent such that people are aware and engaged with ongoing data mining practices can be deployed. As discussed earlier, a practical first step is to ensure all consent documents can be understood by a broad spectrum of the population.

5.2 Agency

Assuming that legibility has been satisfied, and people are able to understand how their data are being used, the next challenge is to ensure people are able to act autonomously and control how their data are used beyond a single consent decision. Some ways of enabling this include ensuring people can subsequently revoke their consent for their data to be used at any time, without necessarily being precipitated by a change in how the data are used. In the GDPR, this is enshrined through the right to be forgotten [11] that includes the cascading revocation of data between data controllers.

Legibility can also enable agency by allowing people to act in a certain way in order to selectively allow particular inferences to be made. By being able to choose what they are willing to share with a data collector in order to satisfy their own utility, some of the power balance can be restored, which has been previously tipped towards the data collector who is able to conduct analyses at a scale beyond any individual subject's capabilities.

5.3 Negotiability

As discussed in Section 4, Nissenbaum's contextual integrity [39] can be used to detect privacy violations when the terms of data-handling have changed in such a way that existing norms are breached. The principle of negotiability is key to preventing this, by allowing people to make ongoing decisions about how their data are used as contexts evolve, whether their own, environmentally, or that of the data collector.

Dynamic consent in the biobanks context [26] could be adapted to allow data subjects to be notified and review how their data are being used, whether for new purposes or shared with new actors, allowing consent to be renegotiated. Our consent method informed by contextual integrity [23] is one such approach which aims to tackle this problem, by allowing people to make granular consent decisions without being overwhelmed. Adopting the principles of the GDPR, which emphasises dynamic consent, can support negotiability, with guidance made available for organisations wishing to apply these principles [51].

6 Conclusion

Data mining is an increasingly pervasive part of daily life, with the large-scale collection, processing, and distribution of personal data being used for myriad purposes. In this chapter, we have outlined how this often happens without consent, or the consent instruments used are overly complex or inappropriate. Data mining is outgrowing existing regulatory and ethical governance structures, and risks violating entrenched norms about the acceptable use of personal data, as illustrated in case studies spanning the commercial and academic spheres. We argue that organisations involved in data mining should provide legible consent information such that people can understand what they are agreeing to, support people's agency by allowing them to selectively consent to different processing activities, and to support negotiability by allowing people to review or revoke their consent as the context of the data mining changes. We have discussed recent work which dynamically negotiates consent, including a technique which leverages social norms to acquire granular consent without overburdening people. We call for greater public debate to negotiate these new social norms collectively, rather than allowing organisations to unilaterally impose new practices without oversight.

7 Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/L021285/1].

References

1. Akkad, A., Jackson, C., Kenyon, S., Dixon-Woods, M., Taub, N., Habiba, M.: Patients' perceptions of written consent: questionnaire study. BMJ **333**(7567), 528+ (2006). DOI 10.1136/bmj.38922.516204.55
2. Ayalon, O., Toch, E.: Retrospective privacy: Managing longitudinal privacy in online social networks. In: Proceedings of the Ninth Symposium on Usable Privacy and Security. ACM, New York, NY, USA (2013). DOI 10.1145/2501604.2501608
3. Barnes, S.B.: A privacy paradox: Social networking in the United States. First Monday **11**(9) (2006). DOI 10.5210/fm.v11i9.1394
4. Bauer, L., Cranor, L.F., Komanduri, S., Mazurek, M.L., Reiter, M.K., Sleeper, M., Ur, B.: 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. ACM, New York, NY, USA (2013). DOI 10.1145/2517840.2517859
5. Berg, J.W., Appelbaum, P.S.: Informed consent legal theory and clinical practice. Oxford University Press (2001)
6. Brown, I., Brown, L., Korff, D.: Using NHS patient data for research without consent. Law, Innovation and Technology **2**(2), 219–258 (2010). DOI 10.5235/175799610794046186
7. Carmichael, L., Stalla-Bourdillon, S., Staab, S.: Data mining and automated discrimination: A mixed legal/technical perspective. IEEE Intelligent Systems **31**(6), 51–55 (2016). DOI 10.1109/mis.2016.96
8. Donovan-Kicken, E., Mackert, M., Guinn, T.D., Tollison, A.C., Breckinridge, B.: Sources of patient uncertainty when reviewing medical disclosure and consent documentation. Patient Education and Counseling **90**(2), 254–260 (2013). DOI 10.1016/j.pec.2012.10.007
9. Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., Sandvig, C.: "I always assumed that I wasn't really that close to [her]": Reasoning about invisible algorithms in news feeds. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15, pp. 153–162. ACM, New York, NY, USA (2015). DOI 10.1145/2702123.2702556
10. European Parliament and the Council of the European Union: 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 (1995)
11. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union **L119/59** (2016)
12. Fayyad, U., Piatetsky-Shapiro, G., Smyth, P.: From data mining to knowledge discovery in databases. AI Magazine **17**(3) (1996). DOI 10.1609/aimag.v17i3.1230

13. Friedman, B., Lin, P., Miller, J.K.: Informed consent by design. In: L.F. Cranor, S. Garfinkel (eds.) Security and Usability, chap. 24, pp. 495–521. O'Reilly Media (2005)
14. Gomer, R., Schraefel, M.C., Gerding, E.: Consenting agents: Semi-autonomous interactions for ubiquitous consent. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, UbiComp '14 Adjunct, pp. 653–658. ACM, New York, NY, USA (2014). DOI 10.1145/2638728.2641682
15. Hamnes, B., van Eijk-Hustings, Y., Primdahl, J.: Readability of patient information and consent documents in rheumatological studies. BMC Medical Ethics **17**(1) (2016). DOI 10.1186/s12910-016-0126-0
16. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning: Data Mining, Inference, and Prediction, corrected edn. Springer (2003)
17. Heimbach, I., Gottschlich, J., Hinz, O.: The value of user's Facebook profile data for product recommendation generation. Electronic Markets **25**(2), 125–138 (2015). DOI 10.1007/s12525-015-0187-9
18. Hektner, J.M., Schmidt, J.A., Csikszentmihalyi, M.: Experience sampling method: measuring the quality of everyday life. SAGE Publications, Thousand Oaks, CA, USA (2007)
19. Hill, K.: Facebook Added 'Research' To User Agreement 4 Months After Emotion Manipulation Study (2014). URL `http://onforb.es/15DKfGt`. Accessed 30 November 2016
20. Hoadley, C.M., Xu, H., Lee, J.J., Rosson, M.B.: Privacy as information access and illusory control: The case of the Facebook News Feed privacy outcry. Electronic Commerce Research and Applications **9**(1), 50–60 (2010). DOI 10.1016/j.elerap.2009.05.001
21. Hodson, H.: Did Google's NHS patient data deal need ethical approval? (2016). URL `https://www.newscientist.com/article/2088056-did-googles-nhs-patient-data-deal-need-ethical-approval/`. Accessed 30 November 2016
22. Hodson, H.: Google knows your ills. New Scientist **230**(3072), 22–23 (2016). DOI 10.1016/s0262-4079(16)30809-0
23. Hutton, L., Henderson, T.: "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 (2015). URL `http://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10493`
24. Jackman, M., Kanerva, L.: Evolving the IRB: Building robust review for industry research. Washington and Lee Law Review Online **72**(3), 442–457 (2016). URL `http://scholarlycommons.law.wlu.edu/wlulr-online/vol72/iss3/8/`
25. Kang, J., Shilton, K., Estrin, D., Burke, J., Hansen, M.: Self-surveillance privacy. Iowa Law Review **97**(3), 809–848 (2012). DOI 10.2139/ssrn.1729332
26. Kaye, J., Whitley, E.A., Lund, D., Morrison, M., Teare, H., Melham, K.: Dynamic consent: a patient interface for twenty-first century research networks. European Journal of Human Genetics **23**(2), 141–146 (2014). DOI 10.1038/ejhg.2014.71
27. Kramer, A.D.I., Guillory, J.E., Hancock, J.T.: Experimental evidence of massive-scale emotional contagion through social networks. Proceedings of the National Academy of Sciences **111**(24), 8788–8790 (2014). DOI 10.1073/pnas.1320040111
28. Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., Christakis, N.: Tastes, ties, and time: A new social network dataset using Facebook.com. Social Networks **30**(4), 330–342 (2008). DOI 10.1016/j.socnet.2008.07.002
29. Luger, E., Moran, S., Rodden, T.: Consent for all: revealing the hidden complexity of terms and conditions. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13, pp. 2687–2696. ACM, New York, NY, USA (2013). DOI 10.1145/2470654.2481371
30. Luger, E., Rodden, T.: An Informed View on Consent for UbiComp. In: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, pp. 529–538. ACM, New York, NY, USA (2013). DOI 10.1145/2493432.2493446
31. McDonald, A.M., Cranor, L.F.: The cost of reading privacy policies. I/S: A Journal of Law and Policy for the Information Society **4**(3), 540–565 (2008). URL `http://www.is-journal.org/files/2012/02/Cranor_Formatted_Final.pdf`

32. Miller, F.G., Wertheimer, A.: Preface to a theory of consent transactions: Beyond valid consent. In: F. Miller, A. Wertheimer (eds.) The Ethics of Consent, chap. 4, pp. 79–105. Oxford University Press, Oxford, UK (2009). DOI 10.1093/acprof:oso/9780195335149.003.0004
33. Moran, S., Luger, E., Rodden, T.: Exploring Patterns as a Framework for Embedding Consent Mechanisms in Human-Agent Collectives. In: D. Ślęzak, G. Schaefer, S. Vuong, Y.S. Kim (eds.) Active Media Technology, *Lecture Notes in Computer Science*, vol. 8610, pp. 475–486. Springer International Publishing (2014). DOI 10.1007/978-3-319-09912-5_40
34. Morrison, A., McMillan, D., Chalmers, M.: 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, NordiCHI '14, pp. 471–480. ACM, New York, NY, USA (2014). DOI 10.1145/2639189.2639239
35. Mortier, R., Haddadi, H., Henderson, T., McAuley, D., Crowcroft, J., Crabtree, A.: Human-data interaction. In: M. Soegaard, R.F. Dam (eds.) Encyclopedia of Human-Computer Interaction, chap. 41. Interaction Design Foundation, Aarhus, Denmark (2016). URL `https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/human-data-interaction`
36. Munteanu, C., Molyneaux, H., Moncur, W., Romero, M., O'Donnell, S., Vines, J.: Situational ethics: Re-thinking approaches to formal ethics requirements for human-computer interaction. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15, pp. 105–114. ACM, New York, NY, USA (2015). DOI 10.1145/2702123.2702481
37. Napoli, P.M.: Social media and the public interest: Governance of news platforms in the realm of individual and algorithmic gatekeepers. Telecommunications Policy **39**(9), 751–760 (2015). DOI 10.1016/j.telpol.2014.12.003
38. Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: Proceedings of the IEEE Symposium on Security and Privacy, pp. 173–187. IEEE, Los Alamitos, CA, USA (2009). DOI 10.1109/sp.2009.22
39. Nissenbaum, H.: Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford Law Books, Stanford, CA, USA (2009)
40. Obar, J.A., Oeldorf-Hirsch, A.: The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services. Social Science Research Network Working Paper Series (2016). DOI 10.2139/ssrn.2757465
41. Patrick, A.: Just-in-time click-through agreements: Interface widgets for confirming informed, unambiguous consent. Journal of Internet Law **9**(3), 17–19 (2005). URL `http://nparc.cisti-icist.nrc-cnrc.gc.ca/npsi/ctrl?action=rtdoc&an=8914195&lang=en`
42. Pitofsky, R., Anthony, S.F., Thompson, M.W., Swindle, O., Leary, T.B.: Privacy online: Fair information practices in the electronic marketplace: A report to congress. Security (2000). URL `http://www.ftc.gov/reports/privacy2000/privacy2000.pdf`
43. Recuber, T.: From obedience to contagion: Discourses of power in Milgram, Zimbardo, and the Facebook experiment. Research Ethics **12**(1), 44–54 (2016). DOI 10.1177/1747016115579533
44. Sankar, P., Mora, S., Merz, J.F., Jones, N.L.: Patient Perspectives of Medical Confidentiality. Journal of General Internal Medicine **18**(8), 659–669 (2003). DOI 10.1046/j.1525-1497.2003.20823.x
45. Selinger, E., Hartzog, W.: Facebook's emotional contagion study and the ethical problem of co-opted identity in mediated environments where users lack control. Research Ethics **12**(1), 35–43 (2016). DOI 10.1177/1747016115579531
46. Sleeper, M., Balebako, R., Das, S., McConahy, A.L., Wiese, J., Cranor, L.F.: The Post That Wasn't: Exploring Self-censorship on Facebook. In: Proceedings of the 2013 Conference on Computer Supported Cooperative Work, CSCW 2013, pp. 793–802. ACM, New York, NY, USA (2013). DOI 10.1145/2441776.2441865

47. Solove, D.J.: Privacy self-management and the consent dilemma. Harvard Law Review **126**(7), 1880–1903 (2013). URL http://heinonline.org/HOL/Page?handle=hein.journals/hlr126&id=&page=&collection=journals&id=1910
48. Staiano, J., Oliver, N., Lepri, B., de Oliveira, R., Caraviello, M., Sebe, N.: Money walks: A human-centric study on the economics of personal mobile data. In: Proceedings of Ubicomp 2014 (2014). DOI 10.1145/2632048.2632074
49. Steinke, G.: Data privacy approaches from US and EU perspectives. Telematics and Informatics **19**(2), 193–200 (2002). DOI 10.1016/s0736-5853(01)00013-2
50. Steinsbekk, K.S., Kare Myskja, B., Solberg, B.: Broad consent versus dynamic consent in biobank research: Is passive participation an ethical problem?;. European Journal of Human Genetics **21**(9), 897–902 (2013). DOI 10.1038/ejhg.2012.282
51. Tankard, C.: What the GDPR means for businesses. Network Security **2016**(6), 5–8 (2016). DOI 10.1016/s1353-4858(16)30056-3
52. Vitak, J., Shilton, K., Ashktorab, Z.: Beyond the Belmont Principles: Ethical challenges, practices, and beliefs in the online data research community. In: Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, CSCW '16, pp. 941–953. ACM, New York, NY, USA (2016). DOI 10.1145/2818048.2820078
53. Vučemilo, L., Borovečki, A.: Readability and content assessment of informed consent forms for medical procedures in Croatia. PLoS ONE **10**(9), e0138,017+ (2015). DOI 10.1371/journal.pone.0138017
54. Williams, H., Spencer, K., Sanders, C., Lund, D., Whitley, E.A., Kaye, J., Dixon, W.G.: Dynamic consent: A possible solution to improve patient confidence and trust in how electronic patient records are used in medical research. JMIR Medical Informatics **3**(1), e3+ (2015). DOI 10.2196/medinform.3525
55. World Economic Forum: Personal data: The emergence of a new asset class (2011). URL http://www.weforum.org/reports/personal-data-emergence-new-asset-class
56. Zimmer, M.: "But the data is already public": on the ethics of research in Facebook. Ethics and Information Technology **12**(4), 313–325 (2010). DOI 10.1007/s10676-010-9227-5