A Geoprivacy Manifesto

Carsten Keßler
Department of Development and Planning
Aalborg University Copenhagen, Denmark

Grant McKenzie
Department of Geographical Sciences
University of Maryland, USA

August 7, 2017

Abstract

As location-enabled technologies are becoming ubiquitous, our location is being shared with an ever-growing number of external services. Issues revolving around location privacy – or *geoprivacy* – therefore concern the vast majority of the population, largely without knowing how the underlying technologies work and what can be inferred from an individual's location, especially if recorded over longer periods of time. Research, on the other hand, has largely treated this topic from isolated standpoints, most prominently from the technological and ethical point of view. This article therefore reflects upon the current state of geoprivacy from a broader perspective. It integrates technological, ethical, legal, and educational aspects and clarifies how they interact and shape how we deal with the corresponding technology, both individually and as a society. It does so in the form of a manifesto, consisting of 21 theses that summarise the main arguments made in the article. These theses argue that location information is different from other kinds of personal information and, in combination, show why geoprivacy (and privacy in general) needs to be protected and should not become a mere illusion. The fictional couple of Jane and Tom is used as a running example to illustrate how common it has become to share our location information, and how it can be used – both for good and for worse.

1 Introduction

After Jane wakes up to the chime of her smartphone's alarm, she brings up the weather app to check how to dress for the day. While she skims her inbox, her phone brings up an alert, telling her that her commute might take a little longer today because of construction work on her subway line. She quickly gets ready and leaves the house to make sure she will not be late for work and swipes her

monthly pass to enter the subway station. Since the construction work only caused a few minutes delay in her commute, she still has time to stop at her favorite coffee shop, using her credit card to pay for a cappuccino. When she enters the office building she works in, her phone brings up the reminder she had set the day before to make sure she downloads the client presentation she had been working on last night from her cloud storage to her office computer.

Jane's husband Tom left the house early this morning for a two-day meeting out of town. He did not really mind the two hour drive, since this was his first opportunity for an extended trip in his brand new car. When he purchased it the week before, he had signed up for the roadside assistance plan after his old station wagon had left him stranded several times. Following the GPS instructions, he takes the toll bridge to get out of town and onto the highway. Before he arrives at the meeting, he decides to find a place for breakfast, checking for on-line ratings and recommendations first. Later that day, Tom goes out to have dinner with his colleagues, *checking in* at the restaurant with his favorite social network to let his friends know about their fantastic selection of red wines. After paying with his company credit card, he uses the limousine app on his phone to find a nearby driver to take him back to the hotel.

In these two short, yet very common examples, Jane and Tom have shared their location with a dozen parties: the weather app provider, the operator of the digital assistant, the subway operator, two credit card companies, the reminder app, the cloud storage provider, the roadside assistance provider, the toll station operator, the restaurant ratings service, the social network, and the limousine app service. While some of these services may be offered by the same provider – such as the operating system provider running the weather service and the digital assistant – this demonstrates how we share our location information with a large number of entities on a daily basis, together with other personal identifiable information (PII) such as credit card numbers, user names, license plates, or customer numbers. Such location information does not always come as readily mappable pairs of geographic coordinates, but rather as the ID of a subway turnstile, a toll gate, or the name of a restaurant. However, such qualitative location information can still easily be georeferenced (Vasardani et al., 2013, for example), and in combination provide a detailed picture of an identifiable individual's whereabouts.

While these samples of their location history are always triggered by a specific action such as a credit card payment or the swipe of a subway pass, cell phones act as permanent (coarse) positioning devices through the cellular towers they are connected to. With the vast majority of adults in the world – including many developing countries – carrying cell phones today (World Bank, 2016), network providers have a continuous record of their users' locations that goes far beyond the samples in our introductory example (Ahas et al., 2015). Moreover, having these records for a large number of users and long periods of time bears the potential for analyses at the social network level (Eagle et al., 2009), especially if linked to the users' communication through phone calls and text messages. Likewise, operators of WiFi hotspots can keep track of devices that pass by frequently, even if they do not connect to the hotspot (Miller,

2013).

Whether location information comes as point samples or as a continuous track stored by a mobile phone operator (Zeit Online, 2011; Sascha Venohr, 2012), Jane and Tom do not know, or have no control of, what happens to their location information. This holds even if they *did* read the terms of service for the applications they have been using and the provider contracts they signed. They most often do not know if the location will be stored and for how long, and whether the information will be shared with other parties, or whether it may at some point be accessed illegally by malicious hackers. Even if said terms of service contain information about the storage times, there is no way for users to check whether service providers stick to the promised time limits, let alone enforcing them. They do not know at what resolution the location is stored, and whether the recorded location has been correct in the first place. Morerover – and maybe most importantly – they may not even be aware that their location information is being recorded at all.

The goal of this article is to reflect upon the current state of individual location privacy – or *geoprivacy* – and attempt to set the research agenda in this area for the coming years. While it does give a coarse overview of the current state of the art in this field, its primary goal is not to serve as a review article; Beresford and Stajano (2003), Duckham and Kulik (2006), and Krumm (2009) already provide excellent and more comprehensive reviews. Wegener and Masser (1996) have outlined four different scenarios concerning the development of geoprivacy in the mid-1990s, and, 20 years later, came to the conclusion that these scenarios cannot be viewed in isolation. Moreover, they "underestimated the huge privacy problems connected with these technologies" (Masser and Wegener, 2016, p. 1158). The main argument put forward here is therefore that the community has been approaching geoprivacy independently either only from a technological standpoint, or only from an ethical standpoint so far. We argue that those technological and ethical perspectives need to be combined and integrated with educational and legal aspects of this problem space to address the fact that geoprivacy concerns almost every individual in the developed world, and increasingly also in developing countries.

Summarizing these observations in a *manifesto* as a series of theses seems to be an appropriate format for this purpose, which clearly states and reflects on the current state of the art, points out goals, and formulates challenges. One does not have to go back to Luther or Marx and Engels for examples – famous examples that are closer to the topic of this article include the naïve physics manifesto (Hayes et al., 1978), which has been transferred to our domain as naïve geography (Egenhofer and Mark, 1995); the manifesto for agile software development (Beck et al., 2001); the manifesto of transdisciplinarity (Nicolescu, 2002); the object-oriented database system manifesto (Atkinson et al., 1989); and the cluetrain manifesto (Levine, 2000). As in all of these examples, the theses put forward in the geoprivacy manifesto will necessarily reflect the authors' personal views to a certain extent; however, this is intentional, since one of the reasons for a manifesto is to trigger a discussion in the community.

2 Yet again, spatial is special

Geoprivacy[1] has been defined "as a special type of information privacy which concerns the claim of individuals to determine for themselves when, how, and to what extent location information about them is communicated to others. In short, control of location information is the central issue in location privacy" (Duckham and Kulik, 2006, p. 36). One might argue that geoprivacy is just a special case of general information privacy. Both have recently declined, as more and more PII is being collected, stored and shared between different governmental and private actors. Examples include credit ratings, purchasing behavior, and health records (including the increasing use of fitness apps, wearables and the quantified self movement) (Libert, 2015; Lupton, 2016), to name but a few. As in many other cases though (Anselin, 1989; Egenhofer, 1993; Gould et al., 1996; Hart and Dolbear, 2009), *spatial is special* for the sharing (and collection) of PII for the reasons discussed in the remainder of this section.

Thesis 1 Information about an individuals' location is substantially different from other kinds of personally identifiable information.

2.1 Access to location information

The insight that geographic information has gone mainstream and is now being used and produced by a much bigger audience than just GIS professionals is by no means new. Dunn (2007, p. 618) already commented on this in 2007, even before smartphones with embedded GPS chips were available.[2] The initial spark for this development was the discontinuation of selective availability for the Global Positioning System (GPS) in May 2000,[3] which accelerated the development of ever smaller and cheaper receiver units. GPS receivers are now embedded in the main chips for most mobile devices. While there seem to be no reliable numbers on the total number of GPS enabled smartphones, the 2015 median for adults who report owning a smartphone was 68% for advanced economies, and 54% for emerging/developing economies (Pew Research Center, 2014). Taking wearable devices such as sports watches and fitness trackers into account, it is safe to assume that billions of devices worldwide can now be positioned via GPS, most of which are more or less permanently with their owners. Other positioning techniques based on WiFi or Bluetooth (Zandbergen, 2009) can be used to augment or even replace GPS-based positioning, making the trackable population even larger.

In order to make it easy for app developers to access the location information gathered by these hardware sensors, application programming interfaces (APIs)

[1] This article uses the terms *geoprivacy* and *location privacy* synonymously, as does the literature in this field.

[2] In fact, the first iPhone was presented the same year (which only got built-in GPS with the 3G model in 2008), and the first version of Google Maps was introduced only two years before.

[3] See `http://www.gps.gov/systems/gps/modernization/sa/`.

have been added to all major mobile operating systems.[4] For web apps based on HTML and JavaScript, the World Wide Web Consortium (W3C) has developed a standardized location API as part of the HTML5 effort. It has been implemented in all major mobile web browsers. Many browsers running on desktop operating systems also offer location-based services through this API, albeit at a coarser resolution, since the location is usually an estimate based on IP address.

The developments both on hardware and software levels outlined above have made a user's location extremely easy to capture, often with only a few lines of code, and the required hardware sensors are built into even relatively low-end devices. Therefore, location data is much easier to obtain than other kinds of data about users.

Thesis 2 Ubiquitous positioning devices and easy-to use APIs make information about an individuals' location much easier to capture than other kinds personally identifiable information.

2.2 Utility of location information

On top of the technical ease of sharing and collecting location information, users have a high incentive to share their whereabouts with information services. In other areas such as finance, market research, or health, it is often not in the interest of the user to share any of their information. It is in the interest of a bank, online shop, or health insurance company to know as much as possible about their (potential) customers, but sharing their information hardly ever improves the services significantly for the users (despite online advertising firms' claims that tracking users allows them to show more *relevant* advertising, for example). On the contrary, users can be denied a mortgage (or granted only at high interest rates), or insurance companies can turn down applications based on the applicant's financial or health profile.

Sharing a user's location, however, often does improve a service significantly. When Jane is searching the web for a pediatrician or evening classes to learn a new language, getting results independent of her location is quite useless, as she won't take her child to a pediatrician across the country, and she won't attend Danish classes every Thursday night in Chicago if she lives in Seattle. As a result, users either add at least their current city to the query (Rose and Levinson, 2004), or they allow the search engine to use their current location directly through the APIs discussed in Section 2.1. Admittedly, the information that the pediatrician with the best online ratings is across town, but the third best is just around the corner would have been available without sharing the user's location *in principle*. However, *automatically* getting results that have been adapted to their current location is a significant improvement: it saves a number of clicks and visiting websites that are not relevant, and therefore,

[4] See `https://developers.google.com/android/reference/com/google/android/gms/location/package-summary` and `https://developer.apple.com/library/ios/documentation/UserExperience/Conceptual/LocationAwarenessPG/CoreLocation/CoreLocation.html`.

time. While convenience seems to be the main factor for such services, others do not work at all without knowing the user's current (and/or future) location. Examples for such services include directions to get from A to B using different modes of transportation, *check-ins* in social networks, finding nearby friends, or location-based notifications (Fechner et al., 2016) or reminders such as "remind me to drop off that letter next time I'm near a post office".

Thesis 3 Users of information services have a substantial incentive to share their location with service providers, as location information can significantly improve the quality of a service and make it more useful.

All of this is not to say that users always share their location knowingly and willingly. Numerous examples have been documented where mobile apps track their users' locations, while the apps have clearly no requirement to know where the user is. Examples of such *coerced geographic information* (McKenzie and Janowicz, 2014) include flashlight apps and games without any functional requirement to know their users' locations (Hong, 2012), thus violating their geoprivacy to learn as much as possible about them. A recent study has also unveiled a new type of coerced geographic information, using ultrasound signals for location tracking (Arp et al., 2017). These signals, inaudible for the human ear, are emitted by speakers placed in stores and picked up by an app through a smartphone's microphone. These apps use one of the APIs developed by several companies in this business. This type of information provides the participating stores with information about frequent customers and the API provider with detailed personal-level location profiles. These get more detailed as the number of participating stores increases. It is safe to assume that the vast majority of users are not aware that they are being tracked this way.

Thesis 4 Users often share their current location unknowingly.

2.3 Location-based inferences

A user's location as such, even over longer periods of time, is only part of the story. Location information is also substantially different from other kinds of PII from a service providers perspective, as it allows for inferences about many other kinds of information about its users. Knowing where Tom goes shopping for groceries, what restaurants he prefers, and what kind of coffee shops he frequents can accumulate to detailed consumer profiles over time, even without knowing what he purchased. Knowing home and work locations alone can already put Tom in a narrow socio-economic bracket, with commuting patterns and information about frequent visits to locations such as schools or hospitals making that profile even more detailed. Therefore, while a user's location as such can be classified as *observed* PII, it acts as an enabler for *inferred* PII (OECD, 2013).

Thesis 5 Having access to a user's location history allows for a broad range of *location-based inferences*, such as information about their health, consumer behavior, or social status.

Location-based inferences are not always made with the user's consent, or even awareness. They can easily reconstruct information about a user that they never agreed to or intended to share with a service provider. Moreover, the inferences made are not necessarily correct. While a service provider making location-based inferences does have an interest in making correct inferences, these inferences are predictions, which can go wrong. Service providers strive for a high correlation between those predictions and reality, however, the underlying algorithms will always fail in some cases, especially if they do not match the typical profiles: A user may be in perfect health despite never going near a gym and visiting pubs twice a week, another one may have significant disposable income every month despite living in a working class neighborhood, to name just two examples. Such individual cases are not very problematic for the service providers, as they aim for a high correlation *overall*. They may have significant adverse effects for the affected individuals, however, as they lower their consumer profile rating, making them less attractive customers and potentially driving up costs for services, loans, or insurance. Moreover, there is little the user can do about such incorrect inferences, as they are usually unaware of them in the first place.

Thesis 6 Location-based inferences can reveal information that the user never intended or agreed to share with a service.

In extreme cases, this can extend to individuals who do not even use any location-based services themselves. The default locations recorded for large administrative areas such as countries or states recorded in IP-based geocoding services, for example, have been shown to cause significant problems for indivuals living close to those locations. They have been accused of different crimes based on the geographic location provided for an IP address that the service could not localize accurately, and hence returned the default location for the region – which happened to be at or near their home address (Hill, 2016). In this case, a rather obvious mistake has been made inferring their identity for the location provided by an IP-based geocoder.

Thesis 7 Incorrect location-based inferences can have severe adverse effects for the affected individuals, with little to no opportunity to rectify those errors.

3 Economic value of location information

This section discusses the market value of location information from a business as well as a customer or user perspective. It concludes with observations about an emerging market for location information that connects these two prespectives.

3.1 The business perspective

The success of companies such as Google and Facebook shows that PII clearly has a significant economic value, as the data such companies collect are core assets for them. For some companies, e.g., Foursquare, their primary business model centers around selling their users' location information (Finley, 2016). However, assessing the concrete economic value of PII – i.e., putting a number on it – is not straightforward. Feijóo et al. (2014) have explored different approaches and conclude that revenue per data record is a useful measure. The Organization for Economic Co-operation and Development (OECD) notes, however, that this is only an indirect measure, and proposes the consideration of market prices at which personal data records are offered as a direct measure (OECD, 2013). In the same report, the costs for a data breach are also discussed as a measure, following the assumption that a data record is more valuable the more confidential it is.

While these papers consider location information as part of a data record, they do not discuss the value of information about an individual's location as such. A report from McKinsey estimates the "potential annual consumer surplus from using personal location data globally" at $600 billion (Manyika et al., 2011, p. VII). While the report acknowledges that the value of an individual's location information varies from country to country (and certainly even between much smaller spatial units), it gives no indication as to how the value of an individual record could be assessed. In an attempt to address this problem, Baccelli and Bolot (2011) develop a complex economic model based on potential revenue from a customer in the vicinity of a business. While the model targets a very specific use case – which is just one among many existing and potential for individual-level location information – it does take uncertainty in the location information into account.

Thesis 8 Knowing a customer's location is an economic asset for a business.

3.2 The user perspective

Other researchers have looked at the same problem from the user side, investigating how much money (or other kind of compensation) users would ask for sharing their location. Table 1 summarizes the results of the different studies, which show that users seem to base their asking price both on the type of location information they are supposed to share and the use of that information.

Some of those studies used auctions to simulate a market situation with supply and demand (Danezis et al., 2005; Cvrcek et al., 2006), while others only asked for the price participants were asking for their information (Barak et al., 2013; Krumm, 2009). Nonetheless, in combination, this body of research points to the fact that the situation is not quite as bad as stated by Krumm, who concluded that "people do not care about location privacy" (Krumm, 2009, p. 392). Users seem to be aware – at least partially – that their location information is an economic asset, and that it can potentially be used against them. This insight is supported by the fact that most of the studies cited in this section

Table 1: Prices individuals asked for different aspects of their location information in different studies.

	N	Location type	Price asked
Danezis et al. (2005)	74	Constant tracking for one month	£10 (academic use) £20 (commercial use)
Cvrcek et al. (2006)	1200	Constant tracking for one month	€40–50 (academic use) €80–90 (commercial use)
Krumm (2009)	250	2 weeks of car GPS tracks	1% chance to win $200 MP3 player
Barak et al. (2013)	25	Work and home locations	€8.00 (home) €5.40 (work)

also report participants who refused to share their location information, either altogether or for specific (commercial) use.

Thesis 9 Users value their own location information based on level of detail and use case.

3.3 An emerging market

The recent trend of offering discounts to customers who share their location information might provide insights that combine the business perspective and the user perspective. Users have been sharing their location in exchange for a free service for years now, including navigation or lookup of points of interest. More recently, however, businesses have started offering monetary compensation to customers who agree to share their location with them. Auto insurance companies have started offering *pay while you drive* plans that require customers to install an OBD-II adapter with GPS and cellular connection in their vehicles. The incentive for the customers is that they only pay for the time and distance they actually drive, thus saving money compared to the classic flat-rate model. For example, *Progressive Snapshot* customers saved an average of $150 on their car insurance in 2015 (Progressive News Release, 2015), so Tom decided to sign up for this plan with his new car. In return, the insurance company gets to learn not only when and where he goes by car, but also learns about Tom's driving style, gas consumption, and other engine parameters that can be read through the standardized OBD-II interface. Likewise, first health insurance companies are offering discounts on their premiums or rewards such as gift cards for customers who reach certain goals measured on their fitness trackers (Eastwood, 2016).

Thesis 10 A new market is currently emerging in which businesses and users directly trade personal level location information.

Such business models are still relatively new to the market and only a small percentage of customers seem to have signed up for such plans. Therefore, it

remains to be seen whether this model will gain more popularity in the future. If it does become the norm, however, we can expect that customers concerned about sharing their PII will increasingly incur financial penalties, since they are forced to stay in the more expensive flat-rate plans if they refuse to share their information.

Thesis 11 Discounts for customers who agree to share their location with a business are effectively penalizing customers who refuse to do so.

4 Safeguarding geoprivacy

The preservation of geoprivacy is the focus of a substantial body of work which can generally be divided into approaches that address the privacy-preserving querying of large, already accumulated datasets, and approaches that address the sharing and collection of individual-level location information. More recent research has also attempted to assess or classify the degree of (location) privacy an individual can expect. We will argue that such measures are unrealistic, because they only consider datasets in isolation.

4.1 Privacy-preserving data collection and querying

An issue that often plagues both private industry and research communities is the ability to extract meaning from large, location-identifiable datasets, yet preserve the privacy and anonymity of individuals or groups within the data. Significant headway is being made in this arena and a number of methods are currently in use. K-anonymity (Wang et al., 2014) and differential privacy (Dwork, 2008; Mir et al., 2013) are two of the more common approaches that have been adopted in the spatial sciences with the aim of maximizing the value of a dataset containing location information, while minimizing the chances of identifying individuals or groups in the data. The act of obfuscating location information to allow some groups access while restricting others (Duckham and Kulik, 2005; Seidl et al., 2015) has gained traction as well. Examples include Yahoo's *Fire Eagle* (Kiss, 2008), a third-party location-broking application, and Apple introducing rotating MAC addresses (Zebra Technologies, 2015) with the purpose of protecting individuals from identification through the MAC address of their device.

While the majority of work in privacy-preservation has focused on the geographic coordinates, it is important to consider the non-spatial information we share. Other sources of published data such as the time and language used in social media posts can be used to identify place types and knowing the type of place, e.g., Mexican Restaurant, that Jane visits, can significantly increase the ability to identify her spatial location (McKenzie et al., 2016). There is an increasing trend in people choosing to obfuscate their "traditional" location information, e.g., geographic coordinates, but not realize that their location can still be inferred through other information such as their stated interests, so-

cial activities and temporal behavior, sometimes referred to as *digital exhaust* (Adams and Janowicz, 2012).

Thesis 12 Preserving geoprivacy involves more than obfuscating geographic coordinates. Location can be inferred from non-explicitly geospatial information such as interests, activities and socio-demographics.

4.2 Assessing geoprivacy

Privacy is a graded concept, and so is geoprivacy. While there may be occasions where we want to keep our location completely private, there are often cases where we would agree to share our location up to a certain precision in order for a better quality of service. On the dating website where Tom met Jane, for example, he had listed the city he lived in at the time, but not the exact address. This allowed him to filter potential matches to the same city.

Figure 1 shows this graded concept as the level of privacy intrusion, depicted as an abstract graph. Other researchers have already presented attempts at developing more concrete measures of (location) privacy. Lin et al. (2012) propose to measure privacy as degrees of user comfort, which are based on the users' expectations towards an application's use of a resource, such as location services. Their research has been used in the development of `http://privacygrade.org`, which rates mobile apps on a scale from A to D based on how they respect their users' privacy.[5] Research from the same group has later proposed the use of qualitative privacy profiles that describe different user types such as "Privacy Conservatives" or "The Unconcerned" (Lin et al., 2014). Ultimately these assessments classify the degree of private information a user is willing to share with a service, and the quality of service she can expect in return. In general, there is a consensus that by increasing the privacy of the location information that one chooses to share, the usability of the data decreases (Ardagna et al., 2007; Bhumiratana and Bishop, 2009).

Thesis 13 Any location-based service offered to a user is limited by the amount of private information a user is willing to share.

Mobile operating system settings, letting the user control how and when her phone shares her location with an app, are testament to the difficulty in generating an intuitive measure of geoprivacy. Apple's iOS, for example, currently lets a user select on a per-app basis whether location services should be disabled entirely, available while she is using the app, or available anytime, including when the app is not open. While this is clearly an improvement over a simple *on/off* switch, it does not take the context of each respective app into account, namely why the app needs to know the user's location. Moreover, it is not possible to obfuscate the location information to a certain degree before it is forwarded to the app.

[5] The website is offline at the time of writing of this paper, but can still be retrieved through the wayback machine at `https://web.archive.org/web/20161220030726/http://www.privacygrade.org/`.

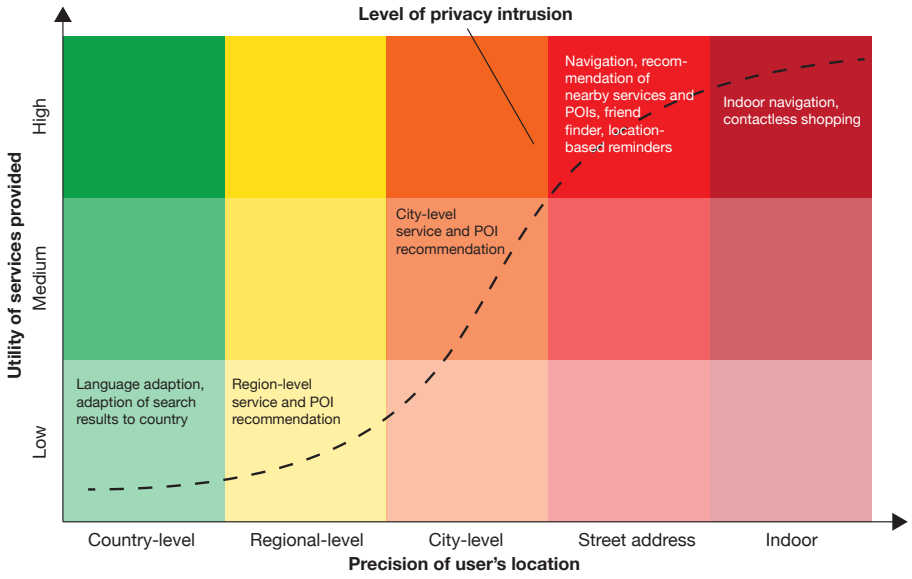

Figure 1: The most useful location-based services depend on knowing the user's precise location. The more useful a service is, the more severe the intrusion on the user's geoprivacy.

Thesis 14 Mobile operating systems lack fine-grained control mechanisms for location services, thus severely limiting the degree of control users have over their location information.

While these classification approaches are all based on the assumption that a user is knowingly sharing her location, a different strand of research has looked at geoprivacy from a security point of view. Yang et al. (2015) introduce a measure based on a how *revealing* a place category can be based on its uniqueness in the vicinity. Fechner and Kray (2012) discuss different human strategies to re-identify previoulsy anonymized individuals. In an experiment involving a location-based game, they have shown that such strategies can be extremely successful. In other words, it is often not very hard to re-identify individuals with some common sense.

Auxilliary data has also been shown to play a vital role in re-identifying individuals within datasets that appeared to be perfectly anonymized when inspected on their own (El Emam et al., 2011). The discussions at and proceedings of two workshops on Geoprivacy organized in 2015 and 2016 outline this dilemma (Keßler et al., 2014; McKenzie et al., 2015), which renders it virtually impossible to reliably assess how well an individual's privacy is protected, given that it is impossible to know what other information a potential attacker could obtain to aid de-anonymization. Hartzog and Rubinstein (2017) have recently argued that it would be better to approach this problem from a standpoint of

risk of de-anonymization, comparing it to the field of IT security. The case for geoprivacy, however, is much more complex than the relatively straighforward calculation of how long it will take a computer to guess a well-chosen password of a certain length. The complexity here arises both from the vast range of sources and potential uses of location information, as well as the range of strategies for de-anonymization.

Thesis 15 An individual's level of geoprivacy cannot be reliably assessed because it is impossible to know what auxialliary information a third party may have access to.

5 Legal and ethical aspects

In many ways, location privacy is a battle still being fought in court rooms around the world. The old adage that technology moves faster than policy or the law, very much holds true here. Law suits either involving or centered around location-aware technology have been, and will continue to be debated as technology, privacy and law collide. For example in 2003, a customer of the Payless car rental company was charged over $3000 for traveling out of state (Office of the Attorney General, 2004), information that was obtained, unbeknownst to the customer, from a GPS unit integrated into the rental vehicle. Similarly, a privacy campaign group in the UK raised legal concerns over devices embedded in recycling bins around London that logged the MAC addresses of passer-bys' wi-fi enabled mobile devices (Miller, 2013). It also appears that a legal gray area has arisen around location tracking of company-issued property (United States District Court, E.D. California, 2015), a discussion that will continue as devices, such as mobile phones, are used in both private and professional situations.

Laws already exist restricting the placement of GPS devices on suspect vehicles by law enforcement agencies (United States Supreme Court, 2015). At the time of writing of this paper, the United States Surpreme Court has just decided to hear argument in a petition dealing with cell-tower based location data (United States Supreme Court, 2017). The petitioner has been convicted of a series of armed robberies, in part due to his presence in the vicinity of the crimes. He is now claiming that the information which cell tower his phone was connected to – acting as a proxy measure for his location in this case – falls under the Fourth Amendmend. The case will be heard later this year and can be expected to provide a legal definition of geoprivacy, at least for the United States.

Legal issues aside, the ethics surrounding location privacy are complex. Social media companies such as Facebook have already begun to blur the lines on what is ethical and socially responsible (Kramer et al., 2014). It is not only privately held companies or government groups that should be focused on these ethical concerns. As researchers we must also be aware of our ethical responsibility to preserving the privacy of individuals (Zimmer, 2010). While research into new methods of de-anonymization, for example, is important, we must not

lose sight of the fact that research scientists have an ethical responsibility to bring concerns over geoprivacy to light. In an evocative paper title *Geoslavery*, Dobson and Fisher (2003) expressed concern about the impact of GIS technology and location-based services, outlining a number of ways that this technology could be used to violate human rights. While many of the ideas appeared farfetched at the time, a number of them are in-use today including real-time location tracking (Goldstein, 2014) and citizen tracking (Nebeker et al., 2016). Some groups have gone so far as to patent *employee tracking systems* (Stoller and Silverstein, 2005). Furthermore, unforeseen ethical dilemmas have arisen from location-enabled technology such as location-based cyberbullying (Black et al., 2016) and celebrity stalking via publicly available transit data (Neustar, 2014).

Thesis 16 The ethical ramifications of advances in location-enabled technology are often viewed as an afterthought and legal concerns over privacy aspects lag behind technological advances.

6 Geoprivacy as a tension field

The previous paragraphs show that geoprivacy is affected and influenced by a number of different factors. Substantial research has been conducted in many of these areas, particularly concerning the technological aspects. A more holistic approach, however, is yet to be developed. In the following, we will attempt to shed some light on the relationships between the different aspects that build this tension field of geoprivacy.

Thesis 17 Geoprivacy as a research topic is situated in a tension field between technological, ethical, economical, legal, and educational aspects that have only been addressed separately so far.

Figure 2 gives an overview of those different aspects that affect a user's geoprivacy. At the center of this field are the *user* and the *tools* the user interacts with, such as mobile apps or devices such as fitness trackers. They are connected through the *utility* offered by the tool, which the user is looking for. In order to leverage this utility, the user needs to provide their *location* at a certain *precision*, where the utility often depends on this precision (see also Figure 1).

This functional core is affected by additional external factors. The user's *eduction* and level of *information* have a major impact on judging whether a tool is safe to use and whether the utility justifies providing some potentially very personal information. The tools, on the other end, depend on *technological developments* that define what can be built. *Service providers* such as software companies, banks, or insurance companies leverage these technologies to build the tools and provide them to the user. Evidently, these processes are most often driven by *economic interests*, which, together with *ethical considerations*, shape the *legal framework* the tools need to act within.

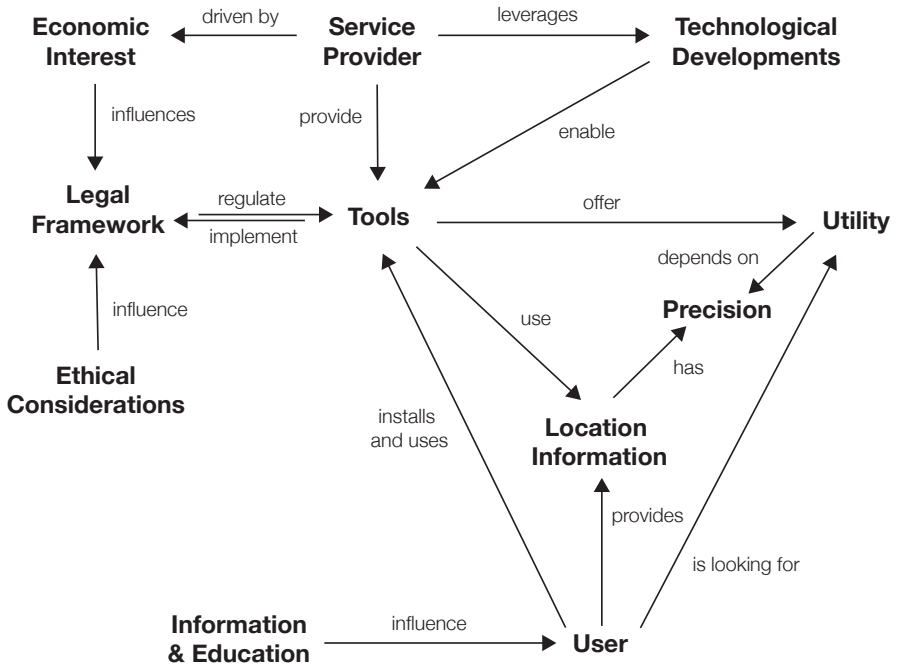

Figure 2: Geoprivacy is influenced by a number of different aspects, creating a tension field that makes it difficult to tackle as a whole.

The direct link between the legal framework and the tools reveal a dilemma of geoprivacy as there is no intermediary control mechanism. In many cases, users simply have to trust that a tool adheres to the legal framework (and to the developer's description and privacy policy), but unless it is an open source tool, there is no way to validate that – and even in the case of open source tools, checking requires an advanced level of education in the corresponding technologies.

Thesis 18 Users often have no way of checking whether the location-aware services and devices they use act within the legal framework and adhere to the provided description and privacy policy.

While it is certainly unrealistic to expect that everybody acquires the programming skills required to do so, certification by independent experts has also been proposed for this purpose, and is already considered best practice in other fields dealing with PII (Hasselbalch and Tranberg, 2016). In cases where no such certification is available, better educated users can be expected to be able to make more informed judgements which tools seem trustworthy, who they share their location information with, and whether this is a reasonable tradeoff for the utility they are getting. Anecdotal evidence from one of the authors' 101-level

college classes supports this claim. In a quiz, almost half of the students believed that the GPS receiver in their phone transmits their location back to the GPS satellite element. Most of them were also surprised to learn that mobile network operators can track them in real time through the cell tower their phone is connected to, at a relatively high precision in urban areas with a dense network of towers. A 2014 survey confirms this, finding that "physical location data is seen as more sensitive among the college educated" (Pew Research Center, 2014, p. 34). Given the almost ubiquituous use of location-based services, education of everyday users with little or no background in technology should therefore be a high priority, ideally as part of general highschool education.

Thesis 19 A higher level of user education in the area of position tracking and location-based services is required to allow them to make more informed decisions about the tools and services they are using.

This lack of technical understanding also seems to hamper initiatives that push for more restrictive legislation in the collection and use of PII, including location information. The fact that many users do not seem to understand what information is being collected about them, and how it is being used, result in a low priority of these topics on the political agenda. Code becomes law (Lessig, 2006), and whatever is technically possible, will be done.

Thesis 20 A better educated user base can push for more restrictive legislation and force service providers to be more transparent about their data collection and use policies.

7 Conclusions

The tension between technical developments, commercial interests, legislation, and an often uninformed user base that consists of a large (and increasing) part of the population make geoprivacy a pressing topic that needs to be addressed beyond the many technical approaches that can be found in the literature. Ethical aspects of recording, processing, and storing a user's location information need a broader discussion that should lead to a clear legislation and more transparency for the users, who in turn need to be better informed about how location-based services work and what can be done with corresponding information. Highschool seems an appropriate level to teach the corresponding topics, which could be covered together with other media- and technology-related contents. Evidently, such initiatives are against many service providers' interests, who want to learn as much as possible about their users' whereabouts. Examples include insurance companies, banks, advertising companies, and customer loyalty programs, among others. In many cases, the information about what happens with collected location information and other PII is hidden in lengthy terms of service and privacy policies that the vast majority of users accept without reading them (Obar and Oeldorf-Hirsch, 2016) – including the authors of this paper.

While the nearly constant surveillance of one's location may be largely abstract to most of us, it becomes very concrete if we know who exactly is interested in where we spend time and who we meet with. This goes for employees who know that their employer can always locate them – e.g., through a company-provided smartphone –, but also for children. An increasing number of parents make use of tracking apps to know where their children are at any given time.[6] While there seems to be no research yet about the long-term effects of constant surveillance by a big brother (or mother), it stands to reason whether it helps raising independent individuals.

The potential for location information to be used as a tool of oppression cannot be overstated, particularly in countries where authoritarian regimes are in control. From participation in demonstrations to meetings with dissidents, an individual's location history bears the potential for imprisonment, or worse. Even in Western democracies, being in the wrong place at the wrong time can have serious adverse effects. In 2011, the phone numbers of thousands of german citizens' were registered based on nearby cell towers after a group of left-wing activists had blocked a neonazi march in Dresden, putting them in the focus of an investigation without any wrongdoing (Biermann, 2017). Such news will arguably lead to preemptive obedience for citizens who fear potential adverse consequences.

Thesis 21 Constant surveillance of citizens' locations can be used as a tool for oppression and to limit freedom of speech, even in democracies.

While these are by no means new observations, there is still no broad discussion, let alone consensus in society as to what uses of location information are acceptable, and where the *red line* is that should not be crossed. Areas such as health or finance are already seeing stricter regulation concerning the use of PII, and, even more important, raised everyday users' awareness. Topics that used to be of interest only to security specialists, such as two-factor authentication or end-to-end encryption, are now being discussed in mainstream media. We hope that this article contributes to this debate, as well as to the manifestation of the right to (location) privacy as an achievement of modern civilization, and not just a mere "blip in human history" (Weigend, 2017, p. 47).

References

Adams, B. and K. Janowicz (2012). On the Geo-Indicativeness of Non-Georeferenced Text. In *ICWSM*, pp. 375–378.

Ahas, R., A. Aasa, Y. Yuan, M. Raubal, Z. Smoreda, Y. Liu, C. Ziemlicki, M. Tiru, and M. Zook (2015). Everyday space–time geographies: using mobile phone-based sensor data to monitor urban activity in Harbin, Paris, and

[6] For example, at the time writing of this paper, the app *Find My Kids – GPS Tracker* for Android is listed with 1–5 million installs, and a 4 out of 5 rating: `https://play.google.com/store/apps/details?id=com.fsp.android.g`.

Tallinn. *International Journal of Geographical Information Science 29*(11), 2017–2039.

Anselin, L. (1989). What is Special About Spatial Data? Alternative Perspectives on Spatial Data Analysis. In *Spring 1989 Symposium on Spatial Statistics, Past, Present, and Future, Department of Geography, Syracuse.*

Ardagna, C. A., M. Cremonini, E. Damiani, S. D. C. Di Vimercati, and P. Samarati (2007). Location privacy protection through obfuscation-based techniques. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pp. 47–60. Springer.

Arp, D., E. Quiring, C. Wressnegger, and K. Rieck (2017). Privacy Threats through Ultrasonic Side Channels on Mobile Devices. In *Proc. of 2nd IEEE European Symposium on Security and Privacy (EuroS&P).*

Atkinson, M. P., F. Bancilhon, D. J. DeWitt, K. R. Dittrich, D. Maier, and S. B. Zdonik (1989). The Object-Oriented Database System Manifesto. In *DOOD*, Volume 89, pp. 40–57.

Baccelli, F. and J. Bolot (2011, April). Modeling the economic value of the location data of mobile users. In *INFOCOM, 2011 Proceedings IEEE*, pp. 1467–1475.

Barak, O., G. Cohen, A. Gazit, and E. Toch (2013). The price is right?: Economic value of location sharing. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, UbiComp '13 Adjunct, New York, NY, USA, pp. 891–900. ACM.

Beck, K., M. Beedle, A. Van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, et al. (2001). Manifesto for Agile Software Development.

Beresford, A. R. and F. Stajano (2003). Location privacy in pervasive computing. *IEEE Pervasive computing 2*(1), 46–55.

Bhumiratana, B. and M. Bishop (2009). Privacy aware data sharing: balancing the usability and privacy of datasets. In *Proceedings of the 2nd International Conference on PErvasive Technologies Related to Assistive Environments*, pp. 73. ACM.

Biermann, K. (2017). Dresdner Polizei fischte mit Datenschleppnetzen. Available from `http://www.zeit.de/digital/datenschutz/2011-06/polizei-dresden-vorratsdaten`.

Black, E. W., K. Mezzina, and L. A. Thompson (2016). Anonymous social media–Understanding the content and context of Yik Yak. *Computers in Human Behavior 57*, 17–22.

Cvrcek, D., M. Kumpost, V. Matyas, and G. Danezis (2006). A study on the value of location privacy. In *Proceedings of the 5th ACM workshop on Privacy in electronic society (WPES '06)*, pp. 109–118.

Danezis, G., S. Lewis, and R. J. Anderson (2005). How much is location privacy worth? *WEIS 5*.

Dobson, J. E. and P. F. Fisher (2003). Geoslavery. *IEEE Technology and Society Magazine 22*(1), 47–52.

Duckham, M. and L. Kulik (2005). A formal model of obfuscation and negotiation for location privacy. In *Pervasive computing*, pp. 152–170. Springer.

Duckham, M. and L. Kulik (2006). Location privacy and location-aware computing. *Dynamic & mobile GIS: Investigating change in space and time*, 34–51.

Dunn, C. E. (2007). Participatory GIS—a people's GIS? *Progress in human geography 31*(5), 616–637.

Dwork, C. (2008). Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pp. 1–19. Springer.

Eagle, N., A. S. Pentland, and D. Lazer (2009). Inferring friendship network structure by using mobile phone data. *Proceedings of the national academy of sciences 106*(36), 15274–15278.

Eastwood, B. (Jul 28, 2016). How Wearing a Fitness Tracker Can Lower Your Insurance. Available online from `http://www.tomsguide.com/us/fitness-trackers-insurance,news-23053.html`. *Tom's Guide*.

Egenhofer, M. and D. Mark (1995). Naïve Geography. In A. Frank and W. Kuhn (Eds.), *Proceedings of COSIT'95*, September 1995, Semmering, Austria, pp. 1–15.

Egenhofer, M. J. (1993). What's Special about Spatial?: Database Requirements for Vehicle Navigation in Geographic Space. In *ACM SIGMOD 1993, Washington, DC, USA*, Volume 22, pp. 398–402. ACM.

El Emam, K., E. Jonker, L. Arbuckle, and B. Malin (2011). A systematic review of re-identification attacks on health data. *PloS one 6*(12), e28071.

Fechner, T. and C. Kray (2012). Attacking Location Privacy: Exploring Human Strategies. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, New York, NY, USA, pp. 95–98. ACM.

Fechner, T., D. Schlarmann, and C. Kray (2016). Facilitating citizen engagement in situ: assessing the impact of pro-active geofenced notifications. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services–Mobile HCI '16*, New York, New York, USA. ACM Press.

Feijóo, C., J. L. Gómez-Barroso, and P. Voigt (2014). Exploring the economic value of personal information from firms' financial statements. *International Journal of Information Management 34*(2), 248–256.

Finley, K. (2016). Foursquare's Plan to Use Your Data to Make Money – Even if You Aren't a User. *Wired*.

Goldstein, J. (2014). To increase productivity, UPS monitors driver's every move. Available from `http://www.npr.org/sections/money/2014/04/17/303770907/to-increase-productivity-ups-monitors-drivers-every-move`.

Gould, M., M. Brand, M. Craglia, S. Fotheringham, and A. Frank (1996). What's so special about spatial? *GIS Europe 5*(10), 22,24–26.

Hart, G. and C. Dolbear (2009). What's So Special about Spatial? In *The Geospatial Web*, pp. 39–44. Springer.

Hartzog, W. and I. Rubinstein (2017). The Anonymization Debate Should Be About Risk, Not Perfection. *Communications of the ACM 60*, 22–24.

Hasselbalch, G. and P. Tranberg (2016). *Data Ethics – The New Competitive Advantage*. PubliShare.

Hayes, P. J. et al. (1978). *The Naïve Physics Manifesto*. Université de Genève, Institut pour les études sémantiques e cognitives.

Hill, K. (2016). How an internet mapping glitch turned a random Kansas farm into a digital hell. *Fusion*.

Hong, J. (2012). Analysis of Most Unexpected Permissions for Android Apps. Available from `http://confabulator.blogspot.com/2012/11/analysis-of-top-10-most-unexpected.html`.

Keßler, C., G. McKenzie, and L. Kulik (Eds.) (2014). *GeoPrivacy '14: Proceedings of the 1st ACM SIGSPATIAL International Workshop on Privacy in Geographic Information Collection and Analysis*. November 4, 2014, Dallas/Fort Worth, Texas, USA: ACM.

Kiss, J. (2008). Yahoo launches fire eagle location too. *The Guardian*.

Kramer, A. D., J. E. Guillory, and J. T. Hancock (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences 111*(24), 8788–8790.

Krumm, J. (2009). A Survey of Computational Location Privacy. *Personal and Ubiquitous Computing 13*(6), 391–399.

Lessig, L. (2006). *Code: And Other Laws of Cyberspace, Version 2.0*. New York: Basic Books.

Levine, R. (2000). *The Cluetrain Manifesto: The End of Business as Usual*. Cambridge, Mass: Perseus Books.

Libert, T. (2015). Privacy implications of health information seeking on the web. *Communications of the ACM 58*(3), 68–77.

Lin, J., S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang (2012). Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pp. 501–510. ACM.

Lin, J., B. Liu, N. Sadeh, and J. I. Hong (2014). Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Symposium On Usable Privacy and Security (SOUPS 2014)*, pp. 199–212.

Lupton, D. (2016). *The quantified self*. John Wiley & Sons.

Manyika, J., M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers (2011). Big data: The next frontier for innovation, competition, and productivity.

Masser, I. and M. Wegener (2016). Brave New GIS Worlds Revisited. *Environment and Planning B: Planning and Design 43*(6), 1155–1161.

McKenzie, G. and K. Janowicz (2014). Coerced geographic information: The not-so-voluntary side of user-generated geo-content. In *Eighth International Conference on Geographic Information Science*.

McKenzie, G., K. Janowicz, and G. Kossinets (Eds.) (2015). *GeoPrivacy '15: Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Privacy in Geographic Information Collection and Analysis*. November 3, 2015, Seattle, Washington, USA: ACM.

McKenzie, G., K. Janowicz, and D. Seidl (2016). Geo-privacy beyond coordinates. In *Geospatial Data in a Changing World: Proceedings of the 2016 AGILE Conference*, pp. 157–175. Springer.

Miller, J. (2013). City of London calls halt to smartphone tracking bins. *BBC News*.

Mir, D. J., S. Isaacman, R. Cáceres, M. Martonosi, and R. N. Wright (2013). Dp-where: Differentially private modeling of human mobility. In *Big Data, 2013 IEEE International Conference on*, pp. 580–588. IEEE.

Nebeker, C., T. Lagare, M. Takemoto, B. Lewars, K. Crist, C. S. Bloss, and J. Kerr (2016). Engaging research participants to inform the ethical conduct of mobile imaging, pervasive sensing, and location tracking research. *Translational behavioral medicine 6*(4), 577–586.

Neustar, R. (2014). Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset. Available from https://research.neustar.biz/2014/09/15/riding-with-the-stars-passenger-privacy-in-the-nyc-taxicab-dataset/.

Nicolescu, B. (2002). *Manifesto of Transdisciplinarity*. SUNY Press.

Obar, J. A. and A. Oeldorf-Hirsch (2016). *The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services*.

OECD (2013). Exploring the Economics of Personal Data: A Survey of Methodologies for Measuring Monetary Value. *OECD Digital Economy Papers* (220).

Office of the Attorney General (2004). Attorney General Lockyer Announces Consumer Protection Settlement with Bay Area Rental Car Firm. *State of California Department of Justice*.

Pew Research Center (November 2014). Public Perceptions of Privacy and Security in the Post-Snowden Era. http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/.

Progressive News Release (2015). Lead Foot Report from Progressive® Insurance busts industry braking standards. Available from https://www.progressive.com/newsroom/article/2015/may/lead-foot-report-from-progressive/.

Rose, D. E. and D. Levinson (2004). Understanding user goals in web search. In *Proceedings of the 13th international conference on World Wide Web (WWW2004)*, pp. 13–19. ACM.

Sascha Venohr (2012). *The Tell-All Telephone*, pp. 93–95. O'Reilly Media.

Seidl, D. E., G. Paulus, P. Jankowski, and M. Regenfelder (2015). Spatial obfuscation methods for privacy protection of household-level data. *Applied Geography 63*, 253 – 263.

Stoller, G. and S. Silverstein (2005). Employee tracking system with verification. US Patent App. 11/152,279.

United States District Court, E.D. California (2015). Arias v. Intermex Wire Transfer, LLC. Case No. 1:15-cv-01101 JLT.

United States Supreme Court (2015). United States v. Jones. 615 F. 3d 544, affirmed.

United States Supreme Court (2017). Carpenter v. United States. Case Nos. 14–1572.

Vasardani, M., S. Winter, and K.-F. Richter (2013). Locating place names from place descriptions. *International Journal of Geographical Information Science 27*(12), 2509–2532.

Wang, Y., L. Xie, B. Zheng, and K. Lee (2014). High Utility K-anonymization for Social Network Publishing. *Knowledge and Information Systems 41*(3), 697–725.

Wegener, M. and I. Masser (1996). *Brave new GIS worlds*, pp. 9–22. Taylor and Francis, London.

Weigend, A. (2017). *Data for the People: How to Make Our Post-Privacy Economy Work for You*. Basic Books.

World Bank (2016). Mobile cellular subscriptions (per 100 people). Available from `http://data.worldbank.org/indicator/IT.CEL.SETS.P2`.

Yang, J., Z. Zhu, J. Seiter, and G. Tröster (2015). Informative Yet Unrevealing: Semantic Obfuscation for Location Based Services. In *Proceedings of the 2Nd Workshop on Privacy in Geographic Information Collection and Analysis*, GeoPrivacy'15, New York, NY, USA, pp. 4:1–4:8. ACM.

Zandbergen, P. A. (2009). Accuracy of iphone locations: A comparison of assisted gps, wifi and cellular positioning. *Transactions in GIS 13*(s1), 5–25.

Zebra Technologies (2015). Analysis of the iOS 8 Mac randomization on locationing. Technical report.

Zeit Online (2011). Tell-all telephone. Available from `http://www.zeit.de/datenschutz/malte-spitz-data-retention`.

Zimmer, M. (2010). "But the data is already public": on the ethics of research in Facebook. *Ethics and information technology 12*(4), 313–325.