university of groningen

university college

InferID

The Possibilities And Implications Of Covert User Identification Via Behavioural Biometrics

Bachelor Thesis

Greg Charitonos

3367649

University College Groningen

10.06.20

Supervisor: Dr. Tjeerd Andringa
Co-assessor: dr. S.M. (Simon) Friederich

Table Of Contents

1. Introduction 4

2. Existing Behavioural Biometric Systems **5**

3. System Design And Applications **8**

a. System Design 8

b. Applications 9

c. Limitations 11

4. Mitigation Techniques **12**

5. Ethical And Societal Implications **14**

6. Public Opinion Survey On InferID **16**

Participants 16

Design 17

Quantitative Analysis 18

Qualitative Analysis 19

Reflection 21

7. Conclusion **24**

8. References **25**

Appendix **29**

a. Survey Questions 29

b. Paragraphs And Themes 31

c. ANOVA Repeated Measures | Traits v. Paragraph 32

d. Boxplots Of Traits 33

e. Spearman Correlations Traits v. [interest] 35

f. Word Cloud 36

g. Samples Of [short_open] By [interest] 36

1. Introduction

"Could a behavioural biometric system be used to identify individuals without their knowledge or consent? What moral and societal implications would arise from its adoption?"

The rise of Machine Learning and Big Data has paved the way for powerful algorithms capable of extrapolating information based on our behaviour (Yu et al., 2015). In the past decade it has been shown that our digital selves can betray sensitive aspects of our nature without our intention (Kosinski et al., 2013), which can be aggregated without our consent. Everything from our aesthetic preferences (Sieu & Gavrilova, 2020) to the way we walk (Connor & Ross, 2018) has been demonstrated as powerful vectors for accurately identifying individuals. Now, given the uptake of mobile computing, the public are armed with small, self-reporting devices containing a number of inputs, including accelerometers, gyroscopes and touchscreens, neatly packaged into an internet-connected smartphone. A number of papers describe the potential for non-traditional forms of authentication methods that function on currently in-use hardware (Dargan & Kumar, 2020), including gesture-based authentication methods, as well as more broadly defined "behavioural profiling" techniques. In some cases, the methods described may be covert, allowing for seamless, behind-the-scenes authentication without further interaction from the user than what is required. At present, such systems are being developed as alternative means of authentication.

In this paper, the notion that similar methods could be used to identify, rather than merely authenticate individuals, and the potential effects of such technologies on society will be examined. Existing biometric authentication systems will be dissected to better understand currently available techniques and the capacity for identification via behavioural biometrics to exist (Section 2). Following this, the design for an example InferID system is proposed, as well as its applications in society (Section 3). Mitigation techniques will be discussed (Section 4), in which several methods for reducing the efficacy of the system will be outlined.

The implications of such a system are vast and complex, and detailed in a section of their own (Section 5). Finally, to better examine the ethical and societal implications of such a system, a survey was conducted, featuring both quantitative and qualitative data (Section 6).

Current Methods for Tracking Online Users

At present two methods for tracking users online are largely in use: "cookies", and "browser fingerprinting".

Cookies

A cookie is "a small piece of information a remote website stores on your computer" (Horn, 2016, p. 183). It can be used to authenticate a user, and / or track their behaviour across a website, by assigning them a unique ID. Typically cookies are temporary, and are intended to be deletable, giving users the option to "opt-out" of online tracking. Several browsers include a “private” or “incognito” mode that automatically deletes cookies after closing the browser. In recent years more persistent forms of tracking have been put to use, despite demonstrations by privacy advocates on the effectiveness of such techniques, and the threat they pose to user privacy. One such technique is Samy Kamkar’s “Evercookie” (Goth, 2011), which demonstrates persistent cookies, capable of being stored in multiple parts of the computer, and recreating itself when the user attempts to delete it (Kamkar, n.d.).

Browser Fingerprinting

Browser fingerprinting is an approach to user tracking that is not wholly reliant on cookies. Instead, a user is identified through various properties of their browser that can be retrieved by a website (Al-Fannah & Mitchell, 2020). As such, even if a cookie is deleted, a web authority would still be capable of identifying a user if they use the same browser, and recreate the cookie. Thus, browser fingerprinting presents itself as a persistent way to track users. It is also largely out of the control of the user, whether or not websites attempt browser fingerprinting, and in some cases can be virtually undetectable (see below). Given this, browser fingerprinting has come under scrutiny from privacy advocates, whose concern is that its effectiveness, undetectability, and the difficulty in developing effective countermeasures make it a threat to user privacy.

Two kinds of browser fingerprinting exist: “passive” and “active” (Al-Fannah & Mitchell, 2020). The former relies solely on identifiable information that is sent in any web request. This includes the brand and version of the browser, the IP address of the client, operating system, and other client-specific information. As this data is present in all web requests, passive fingerprinting is virtually undetectable. The latter form is more invasive, and makes use of client-side scripts (typically in the form of JavaScript files) to gauge and report on other variables, including but not limited to, the width and height of the screen, as well as permission-specific information such as GPS coordinates.

Browser fingerprinting is browser-specific. As such, using a different browser, or a different device, would be an effective countermeasure against user tracking.

As technologies evolve, it can be expected that new methods of tracking users will arise; the field of behavioural biometrics shows promise in this domain. Behavioural biometrics and its capacities in the fields of user tracking, identification and surveillance form the focus of this paper.

InferID: Covert, Permissionless, Behavioural Biometric Identification

The technique to be discussed in this paper would be considered a form of “user fingerprinting” in that a collection of user-specific characteristics are collected and used for identification and tracking. This is achieved through behavioural biometrics, and poses as an even more persistent and insidious threat to user privacy. Given that it is designed to be user-specific, the traits collected and analysed would persist across devices and browsers.

Behavioural biometrics, also known as “soft” biometrics, refers to the analysis of behavioural tendencies. It is a concept that has existed since at least the Second World War, when

Morse code operators noticed that they could identify the sender of a message by subtle differences in the timing of letters (Jenkins et al., 2011). In recent years the field has picked up traction as advancements in technology have led to more robust analytical methods, and an abundance of behavioural data.

Behavioural biometrics could allow cross-domain, cross-platform, cross-device tracking and data mining, potentially allowing an authority to extrapolate offline traits and activities based on online behaviour, without the subject's knowledge or consent. Furthermore, such techniques could void a subject's attempt at remaining anonymous. This would be of interest to a number of authorities, including governmental organisations and law enforcement, business corporations, marketing agencies etc.
The capacity to identify users through behavioural biometrics would be an innovative approach to cybersecurity and related fields, and would grant authorities the ability to unmask the individual behind multiple online identities, devices and platforms.

Throughout this paper the name "InferID" will be used to reference a hypothetical system that infers the identity of an individual, and meets the following criteria:

1. It uses behavioural biometrics to infer identities
2. It does so in a permissionless, unprompted manner. (The individual does not need to consent)
3. It functions cross-platform and cross-device (non-specific environment)
4. Acts seamlessly and in the background (covert).

More information on an example InferID-like system is detailed in Section 3 "System Design And Application".

The following section explores existing behavioural biometric systems, to better understand what InferID may be capable of

2. Existing Behavioural Biometric Systems

At present, a number of papers describe behavioural biometric authentication methods that yield promising results, as listed below. However, all current methods are limited to authentication rather than identification, due to the lack of multiclass classification algorithms in use. By using either binary or unary classification methods, the current systems described are limited to outputting a "true" or "false" label on a dataset, describing whether or not it belongs to an individual of a particular ID. This is a significant difference between existing systems and InferID, and begs the crucial question "can a multiclass classifier be trained to identify a user?" as opposed to merely authenticating a subject (Figure 1).

Figure 1

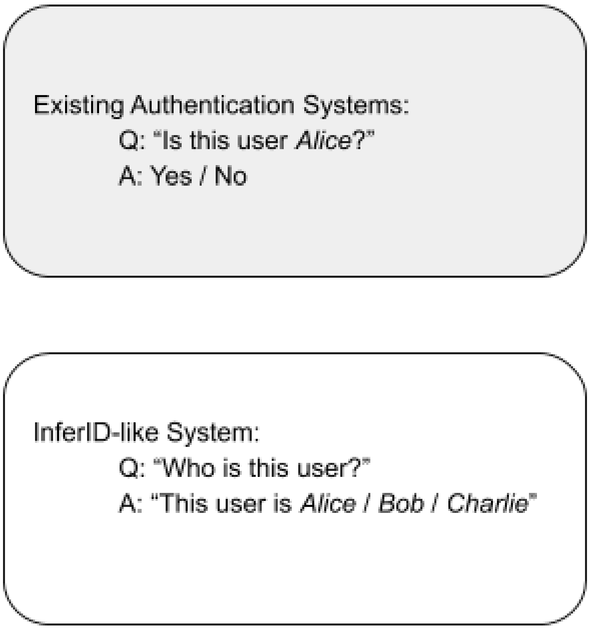

Existing Authentication Systems:
Q: "Is this user *Alice?*"
A: Yes / No

InferID-like System:
Q: "Who is this user?"
A: "This user is *Alice / Bob / Charlie*"

Using constrained horizontal swipe patterns and experimenting with both binary and unary classification algorithms, (Antal & Szabó, 2016) demonstrated accurate and significant authentication via a set of features that include touchscreen gestures and device orientation. None of the features utilised in their experiment required the user to give explicit permission, and all the data collected could be done on any modern iOS or Android device. Using their dataset of 11 features, they achieved an Equal Error Rate (EER)[1] as low as 0.002 when up to five consecutive swipes were used in the authentication process.

Though the promise of such a system makes it tempting to pursue further research into it, (Antal & Szabó, 2016) falls short at its use of a *constrained* swiping motion, ie. specific (not general) swiping motion. The system required users to be prompted (via visual cues) to swipe in a particular direction. For its intended purpose (authentication) it serves well, though not so for InferID's continuous, background identification system. In order to identify users in a covert manner, InferID would accept non-constrained swiping motion (if it were to use swipes to populate its featureset[2]) as the user cannot be prompted to swipe in any particular manner.

Another behavioural biometric authentication system dubbed "Touchalytics" (Frank et al., 2013) demonstrates such potential, and has inspired a number of characteristics in the example InferID outlined in Section 3. Touchalytics is designed to run in the background (i.e covert), uses non-constrained sets of gestural swipes, and obtains them without prompting the user. It can be placed in-app or within websites, and passively collects data without

[1] Equal Error Rate: the value of the false-rejection rate (FRR) and false-acceptance rate (FAR) when both rates are equal.

[2] Featureset: the set of behavioural attributes used as input for an InferID-like system. Eg. the rotational tilt of the device.

direct user interaction or interference. It is a system designed for continuous authentication, running in the background to continually verify the authenticity of the current user, ergo it is still limited to authentication.

In “Touchalytics”, the researchers performed a set of experiments on the system, including experiments a week apart in order to measure the robustness of the system, achieving an EER of 0.04 using a Support Vector Machine as the classification algorithm.

Similar to Touchalytics is “SilentSense”, a behavioural biometric technique that moves beyond the realm of “authentication” and aims to identify individuals in a covert manner (Bo et al., 2013) using touchscreens, accelerometers and gyroscopes, all found on most modern smartphones. The researchers had 100 participants interact with an android phone “freely” with “SilentSense” functioning in the background. In their results they obtained an FRR and FAR of 0.2, with 10 observations, where 1 observation refers to 1 event, eg. a gestural swipe. With 12 observations this number dropped to nearly 0, implying that all identities were inferred accurately. “SilentSense” may be the most promising example of an InferID-like software as of yet.

Taking a different Machine Learning approach, Artificial Neural Networks are the suggested machine learning algorithm for InferID classification, due to their flexibility and robustness. Neural networks have been demonstrated as an effective algorithm for analysing behavioural biometrics, as shown using keystroke analysis (Brown & Rogers, 1993). In an InferID system, that supports multiple devices, a flexible featureset may be a requirement. Although Support Vector Machines have been developed to deal with multiclass classification in the past (Crammer & Singer, 2002), they are designed to suit a specific featureset, making neural networks a more ideal candidate.

More insight into potential design choices for an InferID-like system are detailed in the following section.

Although behavioural biometrics are primarily being explored with authentication in mind, the capacities demonstrated by some of the aforementioned systems suggest that the technology is capable of identification. “SilentSense” in particular has demonstrated an impressive capacity to identify individuals, and is a working proof of concept that identification through behavioural biometrics is possible.

3. System Design And Applications

In this section the design of an example InferID system will be outlined, to better understand some of its functions and capabilities. A discussion on the applications of such a system will follow, bringing to light the various technologies and fields that InferID could disrupt. Finally, there will be a discussion on the various known limitations of the technology, many of which plague current behavioural biometric techniques and similar practices.

a. System Design

InferID could take on a number of different forms, and would likely evolve with any new trends in technology. The system proposed is merely one such example, designed with mobile computing in mind. Smartphones, unlike desktops, are often taken with us wherever we go, and are accessed throughout the day, making them ideal candidates for tracking human behaviour. This system is focused on collecting data using touchscreens and accelerometers, both of which are hardware that is largely available in modern day mobile devices (Pires et al., 2018).

The example is designed to comply to the following:

1. Input
    a. Touch-based (swipe) data: in the form of a set of *n (x,y)* points, where *n* is the number of recorded points in the swipe. (start [*x,y*], *[x,y]...*end)
    b. Device Orientation: *(α, β, ɣ)* rotation recorded once per swipe.
2. Output
    a. *ID*, a number representing the identity of the individual, chosen from the set of all individuals (classes)

In order to achieve this, the example uses a multiclass, deep artificial neural network (ANN). It takes a sequence of *N* inputs. As this example is swipe-based, it takes each touch point *p* and separates its *x,y* values as individual parts of the sequence.

each *x,y* pair for each point *p* is provided, followed by each rotation *α,β,ɣ*. The final input set is:

n(*x,y)...α, β, ɣ* where n(*x,y*) refers to the x,y for each *n* points in the swipe.

The length of the sequence must be constant. This linear set of inputs is capable of being fed into the input layer of ANN of size *N*, without the need for further formatting.

To ensure the input length is of size *N*, the number of touch points *p* (*x,y*) is limited to *Pn* = (*N*- (3 + *p**2)). Excess points *p* in the swipe are removed, preserving the first and last touch points, and a set of artificially included points are added as padding if the number of points is less than *Pn*. The artificially added points would have the mean *x* and mean *y* of all actual recorded points, so as not to alter the trajectory of the gesture. (See figure 2 for illustration).

Figure 2

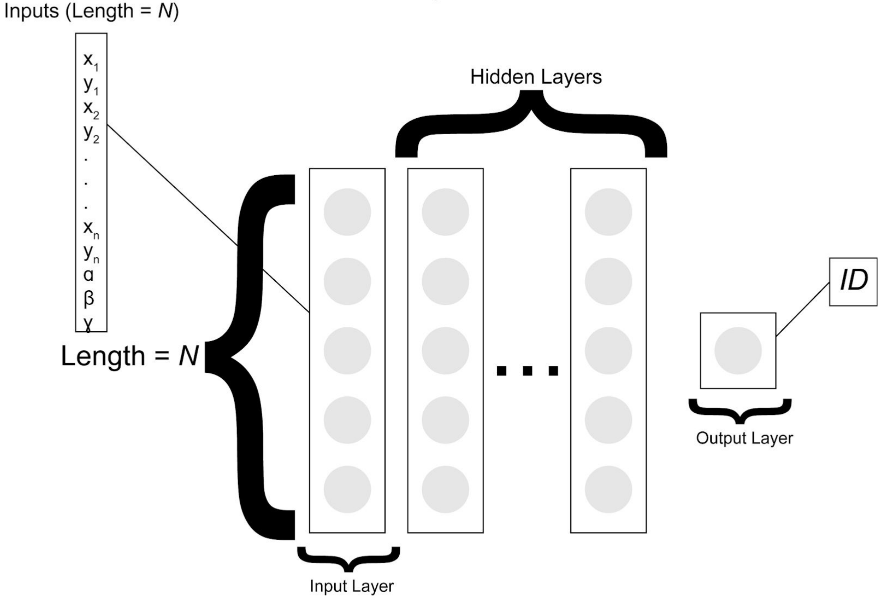

InferID Artificial Neural Network

There are three distinct phases for any InferID system: a collecting phase, training phase, and identification phase, similar to “SilentSense” (Bo et al., 2013). The collecting phase happens continuously on the client-side (user’s device), collecting data at set intervals, or at certain triggered events (eg. at the beginning of any touch), and sends the data to the server, to be stored and processed. A non-identified user will initially be given a unique random ID until the data can successfully be used to classify the user. During the training phase, the data, now server-side, is used in tandem with the user’s ID, initialising a new ID, and fortifying pre-existing IDs with new data, that may be more relevant to the user. The identification phase parses the data through a trained neural-network, whose return value is the most suitable ID. Alternatively, InferID may return multiple IDs, in tandem with a percentage confidence, and select the most appropriate ID, or, if none of the percentages are high enough, may determine that this is a new user, or that it requires more data to return an accurate identity. Similarly to SilentSense, the more data that is available, the more accurate the results for identification.

Existing users would be listed in a database, which holds their ID, as well as their data, or a sample of the most recently uploaded (and hence most relevant) data. This data can be used to retrain the neural network at set intervals (eg. daily, weekly) in order to keep the classifier up-to-date (see “Temporal Instability”, Section 3c, Limitations).

b. Applications

InferID can be used for a variety of applications. Covert behavioural biometrics have a number of boons that highlight their potential to replace passwords. They are difficult to forge, require no memory, and thus cannot be “forgotten” (Dargan & Kumar, 2020). Their seamless nature could usher in an error of passwordless identification.
If a hypothetical InferID system was shared across several sites (perhaps as a new generation of Open Authentication (Leiba, 2012)), it would boost the robustness of such a

system, by collecting different behaviours across different sites presenting different forms of media, and maintaining a single ID for each individual user. Combine this ID with other observed behaviours and a more robust version of current day targeted advertising (Cornière & Nijs, 2016) could be achieved. Online Behavioural Advertising (Boerman et al., 2017) has existed for several years and has been proven to be effective at providing targeted ads towards consumers. The data could be used to better understand and improve user experience, and aid organisations in improving their standards and tailoring experiences to the individual.

As the system can follow users across a number of devices, offline behaviours and traits can be inferred in a manner not yet made possible with existing user-tracking systems. In and of themselves, behavioural biometrics have shown the promise of revealing characteristics of an individual. Researchers have demonstrated the capacity to infer one's gender based on their typing habits (Li et al., 2019) with significant results. Behavioural biometrics have also been used to infer an individual's emotional state. Once again using keystroke dynamics and text patterns researchers accurately inferred user emotions based solely on online behaviour (Nahin et al., 2014). On its own, InferID may be capable of more than just user identification, but may also be able to infer other characteristics of a user. Gauging one's emotions would allow advertisers and user experience designers to tailor ads and media not only to individuals but to individuals *in the moment*, leading to more contextually appropriate content being served.

Such a system could also prove worthwhile for governments across the globe. A robust identification system that spans across a number of devices, applications or sites, and assigns accurate and precise IDs to users could be beneficial in a number of situations. For example, the COVID-19 crisis of 2020, and the capacity to track such outbreaks. Google has demonstrated the ability to track flu seasons based on search patterns in the past (Watts, 2008). With a similar technique, governments could track and trace the spread of contagions through the monitoring of individuals via InferID, and other statistics gathered across-sites; eg. search history, combined with the searches and public posts of individuals in their social circle. In this way governmental organisations could rapidly curb the spread of bio-contagions and spare their citizens from an outbreak. InferID would prove a powerful tool in monitoring cyber-crime, as it would provide a single ID that ties information on an individual across a number of platforms.

If InferID were to run in the background of all legally purchased devices, it would reduce the capacity for any citizen to engage in illicit online activity anonymously. As such, drug and illegal arms trade, piracy, illegal gambling and other cyber-crimes could rapidly decline (Lusthaus, 2012), as the threat of being unmasked becomes ever more prevalent. If law enforcement gained access to a darknet market server, surreptitiously appending an InferID script to the site would prove trivial and may provide insight into the frequent users of the site, both buyers and sellers, and unmask several individuals in the process.

Overall this section outlined a number of innovative solutions to existing issues that could be resolved through the help of InferID; from law enforcement to developing better user experiences. All systems have some limit to their capacities, this will be explored in the following subsection.

c. Limitations

InferID, as it is currently conceived, has a number of limitations, which will be explored in this section. Some limitations may be exploitable during an adversarial attack, which will be explored further in Section 4, Mitigation Techniques. Others merely limit the efficacy of the system. Many of the limitations in InferID are already found in existing behavioural biometric systems.

Context-Specific Behaviour

One glaring limitation lies in context-specific behaviour — different applications / mediums invite different kinds of behaviours. This limits the robustness and precision of the system over multiple contexts, as the classifier would have to assign labels to a wide array of behaviours. As stated in “Touchalytics”, taking the context into question (by adding it as an element in the featureset, for example) both allows the classifier to condition certain behaviours within context, and uses the context as “soft evidence” for the user itself, owing to the fact that some users may use certain apps more than others (Frank et al., 2013). With enough data and a powerful enough machine learning algorithm, InferID may be capable of inferring identities across a wide array of contexts, without necessarily requiring the context to be explicitly included in the featureset.

Diversity In Device Configuration

Another limitation is the wide assortment of hardware and software configurations to cater towards. As a prime example, take the fact that Google Chrome and Mozilla Firefox treat the angular rotation differently, even on the same device (*Detecting Device Orientation*, n.d.). Other limitations involve device-specific behaviour (eg. engaging with a larger device with two hands instead of one). Once again adding the device and / or application (browser) as a label in the featureset may be sufficient, and act as soft evidence. In any of these scenarios the added features can also be manipulated by the software to yield misleading results. Web browsers, for example, send their software name (called a user-agent) to servers by default, though this can be faked, and the tools to do so are a part of most modern web browsers (*Chrome DevTools | Tools for Web Developers*, n.d.), (*Firefox Developer Tools*, n.d.).

Temporal Instability

A wholly separate argument is that over time the individual may change the way in which they engage with the device / application, referred to as “temporal instability” (Frank et al., 2013). The further one ventures from the initial training phase of InferID, the more likely it becomes that the user has altered their behaviour, especially if the engagement context was novel when training first occurred. To fortify against this particular issue, frequent retraining sessions would be wise, to maintain an up-to-date classifier. InferID could train the classifier on a time-specific basis (eg. daily) and on the condition that the classifier already has a high certainty of the identity of the current individual. This would only work if temporal instability results in micro-changes in behaviour, small enough that the classifier can still identify the individual from training session to training session, but large enough that over longer periods of time the results become more inconclusive. In the follow-up experiment conducted using Touchalytics one week after the first set of experiments, the classifier was capable of authenticating users with an EER of 0% for vertical scrolling and 4% for horizontal scrolling (Frank et al., 2013), which is remarkably low considering the classifier had no new information to train on. There is evidence behind this to suggest that user behaviour does not (usually) involve drastic changes in short spaces of time. There are exceptions, of course, that cannot be accounted for eg. sudden injury, that may alter the way in which an individual holds their smartphone.

The discussed limitations bring into question the efficacy of such a system, which would need to be explored with experimentation. These are limitations that might result in unintentional attacks against the system. Other limitations, in particular those reliant on corrupting data streams, are discussed in Section 4.
Although a number of difficult limitations plague the field of behavioural biometrics as a whole, the efficacy of pre-existing systems, even those dominantly used in authentication, suggest that an InferID-like system may not remain hypothetical for very long

4. Mitigation Techniques

The robustness of a hypothetical InferID may be undermined via a number of creative methods. In the process of designing such a system, potential mitigation techniques should be explored, in the interest of improving security and reliability, but also as an adversarial exercise. Supposing that such a system would ever come into use, it would prove vital to the ever-raging battle between privacy and security to have mitigation techniques (MT) to undermine the efficacy of InferID.

At its core, InferID is a data-collection and classification system. A device running InferID would:

1. Collect behavioural biometric data
2. Send data to the server
3. Get response from server (containing ID)
4. Repeat

As the system is reliant on the sending and receiving of data, it is vulnerable to a number of attack vectors affecting data transmission. For example:

1. Blocking data transmission

   In this scenario, data transmission and reception is directly blocked by the MT, in the same manner as traditional ad blockers use. The MT would match the domain or IP Address of any incoming or outgoing requests against an existing blacklist, and if a match occurs, blocks the data before transmission can occur. This technique works well in ad blockers for blocking the majority of ads, due to the fact that the ads are generally loaded from external sources — not from the servers providing the other content on the site / platform. This

attack would not work if the server-side aspect of InferID was being hosted on the same platform as the content.
Other forms of blacklisting could be used, including matching different paths on the url to differentiate between InferID and other resources, and block InferID at that particular location.

A workaround for this technique would be to make data transmission to InferID a requirement in order to access other site / platform resources. By locking the site, or aspects of the site, until a specific response from InferID is provided, platform developers can incentivise disabling a blocker in order to maintain functionality. This is a technique used against ad blockers that has proved effective in the past (Iqbal et al., 2017).

2. Forging misleading packet data

   Another technique would be to transmit forged packets in the form of legitimate data (to pass through filters), in order to reduce InferID's capacity to identify the user with any sort of precision. If InferID relies on classifying behavioural biometric information to individual IDs, then adding noise in the form of forged packets could result in imprecise classification, and reduce the system's efficacy (Co et al., 2019).

   A limitation to this technique is that in order to pass through filters, the data would have to strongly resemble legitimate data, and be sent frequently enough that it would be difficult for InferID to differentiate between legitimate and forged data. Forging realistic data may prove difficult without a complex ruleset to follow.

3. Manipulating packet data

   This technique involves capturing and manipulating data in unpredictable ways, to yield seemingly legitimate packets. This could be achieved by capturing packets before they are sent, replacing any of the data points with a random variable, or by applying randomised transformations to the data. Thus, noise is added to the packet itself, as opposed to transmitting wholly forged packets. Done appropriately, the semblance of the manipulated packet would be of a similarity to a legitimate packet, making it more likely that it would escape any filters.

4. Using stolen packets

   Finally, it may be possible to not only mitigate identification, but to steal an identity, if one gains access to existing packets from another user. An altered version of the client-side InferID code could be used to retrieve data packets from another user (that has obtained the code through a malware infection, for example) and these packets could be sent in place of legitimate packets, or could be used to construct a ruleset for forging new packets, stealing the identity of another user.

The packets of dead users may be put on sale, their identities linked to death certificates yet their online selves living on, as the mask of a cyber-criminal, or revolutionist.

5. Identity forging via behavioural mimicry

   Although improbable, it may not be impossible for an individual to “forge” behaviours in order to mask their true identity. Much as an individual might practice several sets of handwriting, one might be capable of practicing unique behaviours in an attempt to assume multiple identities. On one hand, this could prove difficult, as it would require practicing a number of different behaviours that are unique and unconscious. On the other hand, it may be as simple as swapping hands. Using your left hand to swipe, as opposed to your right hand, or vice versa. This attack would require the ability to repeat the behaviours whilst deliberately avoiding gestures and behaviours that would belong to the true identity of the individual.

Mitigation techniques and adversarial attacks against a system are an important part in gauging its robustness. Understanding these techniques progresses the development of such systems, and enlightens the pitfalls in existing technologies. It also provides a stepping stone for those that wish to undermine the efficacy of an InferID-like system, that might be used by totalitarian regimes. Concerns about unethical or otherwise improper use of such a technology will be explored in the following section.

5. Ethical And Societal Implications

Elaborating on the points discussed in 3b, a vital discussion of the hypothetical InferID would be the ethical and societal implications of such a system.

Unmasking Dangerous Individuals

A covert, permissionless behavioural biometric system would be a powerful tool for unmasking individuals and connecting their numerous online identities back to the biological being from which they spawned. For businesses seeking to avoid illegal content being distributed via their platforms, such a system could be used to verify that an individual does not already have a history of such activities on other sites, or have a criminal record. Employers would better vet their potential employees without the risk of fake or multiple identities being used to hide dark histories and past issues. Much the same techniques could be used by those seeking to make business arrangements, social connections, or dates. Individuals would be unable to hide behind false identities in order to deceive those that they engage with.

In the field of Law Enforcement, such a technique could prove essential, as it would provide virtually indisputable evidence that online illicit activities were perpetrated by a specific individual. By better connecting a user’s behaviour across-site, and perhaps inferring offline behaviour through such activity, law enforcement might better piece together the history of a

crime, or prevent one from ever occurring. Stolen devices could rapidly be retrieved as criminals attempting to use the device are quickly identified and apprehended.
A psychoanalysis of the inner-workings of criminals could be brought to light through the collection of cross-platform engagements, as well as any behaviour that can be directly inferred through behavioural analysis (Kosinski et al., 2014), rapidly improving our understanding of the brain.

A Price To Pay

All such boons come at the high price of the individual's right to privacy. Many privacy advocates fear that currently available technologies threaten our privacy (Al-Fannah & Mitchell, 2020), let alone new technologies. By connecting all online activity to a single identity, one's digital self is augmented and stripped of their masks. A government mandated requirement to have InferID installed on all legally purchased devices would have the consequence that any and all online activities could be traced to a specific individual, threatening to eliminate the right to a private life. Governments have already acknowledged a privacy-threatening desire to instantiate backdoors into devices (Lear, 2017), to imagine a mandate requiring government controlled InferID daemons[3] to be active on all devices is not difficult. Such a system could pave the way towards a centralised database of information on various aspects of an individual's life, and be used to influence and outright control citizens. The "Social Credit System" that is being developed in China has already demonstrated the potential for connecting a number of digital identities, and the effects this has in influencing citizens (Liang et al., 2018). Combine this powerful technique with InferID, and the capability one has to act outside of the omniscient gaze of the government will be significantly undermined. A system such as this, could pave the way for a totalitarian regime to rise, one where every digital, internet-connected device is an extension of the watchful eyes of the government. In much the same way as a government might use the system to curb the spread of diseases, it could attempt to curb the spread of ideas, annihilating their sources and reducing the capacity for citizens to revolt.

Nefarious governments aside, the promise of more robust and precise targeted advertising would prove desirable for organisations selling their product, to the detriment of user privacy. Privacy-enforcing policies have decreased the effect of advertising on consumers (Goldfarb & Tucker, 2011) indicating that it is not in the interest of companies to seek privacy-first ad policies. In 2017 documents surfaced from Facebook demonstrating to advertisers their ability to target teenagers in emotionally vulnerable states (Tiku, 2017), and they have explored this domain in the past (Kramer et al., 2014). Targeting vulnerable populations has been proven to work effectively (Stanton & Guion, 2013) and such controversial targeting methods have come under scrutiny by ethicists, and have been forewarned (Nairn & Berthon, 2003), though they persist to this day.

InferID encroaches on one's right to live a private life, and by extension threatens one's rights to freedom of thought and freedom of expression. It is a dangerous concept in that it directly reduces one's capacity to act anonymously, to produce or digest digital information

[3] Daemon: A computer program that runs in the background, as a service.

without risk of slander or discrimination. In an ideal world such a system could be used justly, and with mercy; to protect the innocent, without prejudice and without intolerance. This is not an ideal world, and the abuse of such an omniscient system has already been forewarned, in fiction and in our encounters with already existing technologies.
More on the potential effects of InferID on society are explored in the following section

6. Public Opinion Survey On InferID

As elaborated above, a behavioural biometric identification system may impact society on a number of levels. Understanding public opinion towards the integration of such a system into society would provide useful insight into its possible effects, as well as public knowledge and opinion on pre-existing surveillance and user-tracking techniques.

This section aims to answer the following question: *"Would participants be averse to InferID because they believe it is morally wrong?"*
In order to examine this, an online survey was conducted. This survey aimed to capture the societal and ethical concerns of the public, with a focus on the ethical implications of InferID. In order to examine this, the survey asked participants for their opinion on three other traits, aside from morality, that may contribute to their aversion. The four traits are:

- Importance `[important]`
- Accessibility `[accessible]`
- Morality `[moral]`
- Safety `[safe]`

Overall desirability of InferID was inferred from the 5-point Likert scale question *"I am interested in using this technology"* `[interest]`. The Quantitative Analysis subsection will center around `[interest]` and its relationships with the four traits.

This section comprises four subsections: Design, Quantitative Analysis, Qualitative Analysis, and Reflection. In each of Quantitative and Qualitative Analysis, the method and results of the respective analyses will be examined. In the Reflection, a final overview of the results obtained will be explored within the wider context of the thesis.
The focus of this survey was to obtain a general understanding on how important a role morality plays in any feeling of interest or aversion felt towards InferID. An in-depth exploration of this requires both quantitative and qualitative analyses, and as such both will be present in the thesis. This is to gain a well-rounded understanding of the opinions and associations that participants manifest towards InferID.

Participants

Participants were reached through online group-chats, and asked to anonymously provide their opinions on InferID after reading short paragraphs depicting the technology in a number of use-cases. All participants were anonymous, and none of the participants were removed from the final analysis of the survey.

In total, 60 responses to the survey were collected. Of the 60, 3 were duplicates, bringing the number of valid responses to 57. Of these, 15 were male, 38 female, and 4 other (3 preferring not to say, 1 intersex). Ages of the participants had a mean of 23.9, and standard deviation of 10.2. The youngest participant was 18, and the oldest 72, with an interquartile range (Q1, Q2, Q3) (20, 21, 22). The data thus comprises largely a young population. The participants covered 21 different nationalities in total, with the largest demographic (21 participants) belonging to the Dutch.

Design

See Appendix (a) for a complete list of the questions asked in the survey, and the corresponding variables they are referenced by.
The questionnaire can be split into three distinct sections:

1. The participant's background
2. Questions regarding InferID
3. General privacy and surveillance questions

The first section comprises general questions on the participant's age, nationality and level of confidence when interacting with technology.
The second section has the participant read four paragraphs and answer four questions per paragraph based on its context. The questions are designed to highlight participant's opinions in the four traits.
The questions are straightforward, with subtle changes to the wording and direction in order to prompt the user to rethink their opinion on the subject, given the new context provided in the most recently read paragraph. These are provided in the form of a 5-point Likert scale, where users were asked to rate how precise a given statement was.
The paragraphs are designed to highlight certain *themes* for InferID, such as covertness or omnipresence, or to demonstrate InferID use cases, such as advertising, law enforcement and surveillance. For a complete list see Appendix (b).

As participants are exposed to different themes, it is expected that their evaluation of InferID on the four traits will change. The second section finally asks participants a Likert-scale question [`interest`], an open question [`short_open`] and finally an optional question where they might give any other remarks regarding InferID [`long_open`].

The last section asks five questions, three of which are designed to gauge the participant's knowledge on internet privacy [`cookies`], [`browser_fingerprinting`] and [`adblocker`], and two to gauge their opinion on surveillance [`surveillance1`] and [`surveillance2`].
Gauging participant's knowledge on internet privacy will provide insight into how aware they are of pre-existing surveillance and user-tracking concepts, which may be used to discern variations in both the qualitative and quantitative opinion-based data. The same can be said about their opinions on surveillance, which may explain any large discrepancies within the data.

Quantitative Analysis

In the survey, a number of questions were given as a Likert scale, in the form of a range between 1 and 5, with 3 representing the mid-range value. In the analysis of the data collected, the data was mapped to the range 0-4, such that the mid-range (2) over max (4) = 0.5, or 50%. Throughout this section, all ranged values will be expressed in this manner.
The crux of this subsection concerns itself with the "Four Traits" listed in the introduction to the section and their relationship to the variable [interest]. An in-depth analysis of the four traits is to follow.

Analysis of Four Traits

Participants scored each of the four traits, in each of the four paragraphs, across a wide range of values. Of interest in this paper is noting which of the four traits has an influence on the overall [interest] displayed towards InferID. Before establishing this connection, it is important to examine the relationship between each of the four traits and the four paragraphs they are influenced by. The null hypothesis being that the paragraphs have no effect on the overall scores that participants gave to InferID in these four dimensions.
To test this, an ANOVA was conducted against each trait, grouped by paragraphs eg. [important] by paragraph (1, 2, 3, 4), with a significance threshold **α**=0.05.
The results of the ANOVAs are displayed in Appendix (c).

As displayed, the p-values [PR (>F)] for each trait are less than **α**, indicating that the differences between each paragraph, for all four variables, is significant. Backing this, the F-ratio [F] for all four variables is significantly greater than 1. Thus, the null hypothesis can formally be rejected. The opinion of InferID, in these four traits, changes in response to new information garnered by each paragraph.

In analysing what effect each paragraph had, a series of boxplots for the four traits was generated. Each plot includes four boxplots for each paragraph, grouped by trait. These can be found in the Appendix (d).
Each trait had at least 1 boxplot unaligned from the others, and no paragraph had the same effect across all four traits. From this it can be inferred that the paragraphs had distinct effects on the traits, bolstering the aforementioned argument that exposure to new information presented in the paragraphs produced a measurable effect on participants' opinions. Paragraph 4 in particular appeared to have an effect on the [moral] trait, indicating that InferID may be perceived as more or less moral in specific contexts.

Traits and [interest]

A primary question that the survey wished to examine was how interested users were in using the technology, as it was described in the paragraphs. This was gathered using the question *"I am interested in using this technology"* [interest] (0, not at all) (4, very). The question was presented after the four paragraphs, followed by the two open questions in this survey.

The variable [interest] is hypothesised to give an overall value of desirability towards the system. The trait has a mean of 1.18, and standard deviation of 1.24. The lowest value provided was 0, and the highest 3, with an interquartile range (Q1, Q2, Q3) (0,1,2). 44% of participants rated [interest] as 0, with the second highest demographic (24%), valuing it at a 3. An important question is whether or not [interest] is based on feelings in any of the four traits. For example, are participants [interest]ed in using the technology because it is [accessible]? More to the point, are participants [interest]ed in using the technology because of how they felt in a particular trait, *within* a particular paragraph? To test this, a Spearman Correlation Coefficient was calculated for each trait, in each paragraph, v. [interest], with **α**=0.05. The null hypothesis being that there is a strong correlation between one or all of the traits, with how interested the participant is in using the technology.
The results are displayed in a table in Appendix (e), with [r] referring to the coefficient, and [p], the two-tailed p-value. The Spearman Correlation Coefficient was calculated, as opposed to the Pearson Correlation Coefficient, as the variables were discrete (Likert scale).

Of interest are those traits with a coefficient [r] of magnitude at least 0.5, and whose [p] values are less than **α**=0.05; these have been highlighted. A correlation coefficient [r] of magnitude between 0.5 and 0.7 is considered a moderate correlation, whilst 0.7 and higher is considered a strong correlation (Dancey & Reidy, 2007). From the table, the strongest correlations appear to be those belonging to the [moral] trait, whose minimum value, 0.505, appears only in the first paragraph, and whose highest value, 0.777, correlates strongest overall with interest. This is a strong indication that there is a moral element contributing to the participant's interest in using InferID. Other notable correlations were those of [safe] in the last two paragraphs, which both moderately correlated with [interest], and [important], which correlated moderately in the second paragraph.

From the correlations it can thus be inferred that there is a moral element contributing to [interest]. In order to better gauge participant's opinions towards InferID, and develop upon some of the quantitative findings, a qualitative analysis was performed.

Qualitative Analysis

In this section the types of descriptions participants gave to InferID will be analysed, evaluated and discussed. Further, these descriptions will be analysed in tandem with participants' level of [interest], to garner further insight into the interpretation of [interest], and what characteristics of InferID may have come to define it.

The survey included two open-answer questions, where participants were asked to provide their answers without input constraint: [short_open] and [long_open]. See Appendix (a) for their corresponding questions.
[short_open] was a required question, whereas [long_open] was the only optional question in the survey. Due to its optionality, only 24 participants provided input for [long_open]. It was generally used to provide remarks that were less relevant than [short_open].

Word Frequencies

Building upon the prior quantitative analysis, [short_open] provides further insight into the opinions that individuals had towards the software, which could not be captured in the more restricted inputs that comprised the rest of the form.
Below, a frequency distribution of the 15 most frequent words in [short_open] is displayed (Figure 3).

Figure 3 | Word Frequency In [short_open]

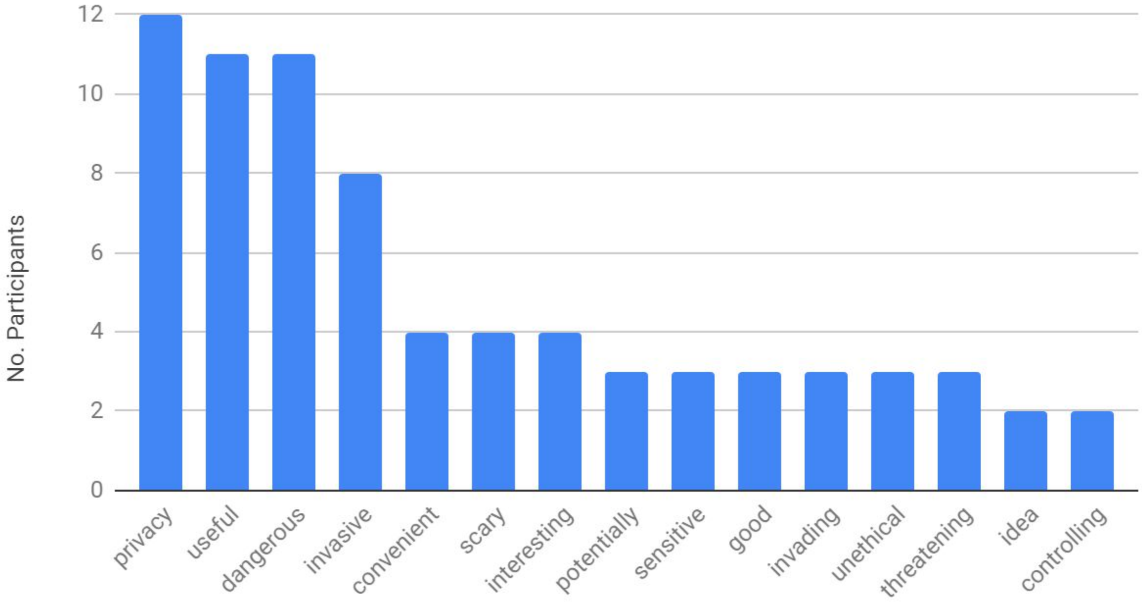

This distribution removed duplicate words for each participant thus it represents the number of participants using a given word. Filler words ("and", "it", "the") were also removed. For a more visual representation, including more words, see the "Word Cloud" in Appendix (f).

From the diagram above, a sense of the associations participants had with InferID can be noted. The word *"privacy"* is most frequent among participants, indicating that it was a salient concept for many. In context, the word was most often used to express concern. Phrases such as *"infringes on people's privacy"* [id=3], *"privacy risk"* [id=24] and *"Breach on privacy"* [id=34] were common.
The words *"useful"* and *"dangerous"* were also used, sometimes in the same sentence ([id=0], [id=15]). A point to note is that many of the words being used indicate a wide array of opinions, and often describe a mix of emotions. The word *"good"* appears as frequently as the word *"unethical"* despite being implicit opposites. The technology is described as *"convenient"*, but also as *"invasive"*, *"threatening"*, and *"scary"*.

[short_open] Grouped By [interest]

An attempt was made to discern any notable patterns in [short_open] responses grouped by [interest]. Finding a pattern could indicate how participants may have interpreted the [interest] question, and adds context to their answers.

Sorting by [interest], whose values ranged from 0-3, a random sample of five participants was taken (where [id] is the unique participant id). These tables can be found in Appendix (g).

Starting with those that gave [interest] a value of 0 (corresponding to *“Not at all interested”*), the words *“unethical”*, *“violating”*, *“creepy”* and *“invasive”* indicate a strong aversion to the use of InferID, which fits the [interest] value provided. Moving up one [interest] point, the word *“invasive”* is again noted, alongside the words *“helpful”* and *“useful”* which denotes a complexity in the opinions held toward the system and its applications. Participant 29, whose [short_open] was *“helpful, invasive”* elaborated further in [long_open] explaining that *“The aspect of installing safety measures to prevent misuse is not discussed”*. The survey did not include any information that would give participants an inclination of the kind of security surrounding InferID, to allow for freedom of interpretation.

Participant 14’s [short_open] introduced a “user” or human element to the argument, essentially treating the system as a tool, but placing mistrust upon the users, who have the potential to abuse its power. This sentiment is echoed by participant 20 [interest=2], in their [short_open], which explicitly stated a lack of trust towards humans.
The samples of [interest=2] introduced significantly more non-uniformity in the answers provided, which ranged from *“innovative”* and *“mostly ethical”* to *“dangerous”* and *“constant surveillance”*. Participant 21’s statement of *“In between”* concisely explains selecting an “in between” option for [interest].

At [interest=3] more discrepancy was perceived in individual answers, which, aside from Participant 8’s short answer of *“Useful”*, all contained both positive and negative descriptions. *“Relevant, threatening, invasive”* [id=11], for example, or the elegantly put *“useful properties, harmful opportunities”* [id=24] describe a mixture of opinions that indicate a complexity in the emotions felt towards InferID.

In summary, outside of [interest=0], whose opinions were almost uniformly negative, the majority of users expressed both a wide array of opinions, but also, internally, a mixture of emotions. This is indicative of InferID’s dual nature, in that it can be perceived as being used for good near as much as it can be seen as dangerous, and invasive.

Reflection

Summary Of Analyses

The results of the analyses provide useful insight into the opinions and concerns that subjects of InferID may have. In the quantitative analysis, strong and moderately strong correlations were observed in the [moral] trait v. [interest], with more moderate correlations in [safe] and [important]. All of these correlations were statistically significant. The strength of the correlations lends credit to the hypothesis that there is a moral aspect contributing to participants’ [interest] in using the technology. The moderate-to-high correlations of [safe] with [interest] indicated that there may be a safety aspect contributing as well. In the qualitative analysis both of the relationships pertaining to [moral] and [safe] were fortified by the [short_open] word frequencies, which featured safety-related words like *"dangerous"*, *"threatening"* and *"unethical"* among the top ranks.

Delving more in-depth into some of the [short_open] responses, clear indications of a moral element surfaced, through phrases like *"potentially immoral"*, and *"ethically sensitive"*. Many respondents indicated an interest into the "privacy" aspect of InferID, with more using the word *"invasive"*, one participant, [id=4], considering it *"Spyware with a slice of tyranny"* and others echoing similar cautionary messages. Participants [id=42] and [id=55] referenced *"big brother"* in their responses, alluding to "mass surveillance" and other privacy-infringing notions associated with George Orwell's "1984", and it's "Big Brother" government.

The majority of [short_open] responses had some negative remarks pertaining to the ethical, safety, and privacy-threatening implications of the technology. This is in line with the distribution of values in [interest], where 68% of participants rated their level of interest below 2 (the 50% line). The study suggests that many of the concerns raised in Section 5, Ethical Implications, were shared among the participants. The salience of "privacy" in the responses was indicative of the significance that participants felt towards InferID's invasive nature.

Not all responses were negative, however. Some responses were wholly positive, such as *"Progressive, intuitive"* [id=45]. One participant felt that InferID was *"Important for the future, we are so many and the internet is such deep and uncharted territory, we need to protect ourselves from it"* [id=22]. These responses were few, but their interest in InferID may indicate that a subset of the population would welcome an InferID-like system. The majority of responses were a mix of positive and negative emotions, and this was apparent in the random samples used in the Qualitative Analysis subsection.

Limitations

Participants

Participants almost entirely consisted of university students, were mostly from a "WEIRD" sample (Henrich et al., 2010), and were mainly born and raised in The Netherlands. This is a very specific sample, and as such is not very representative of the general population. A larger, more diverse population sample would be required to produce reasonable estimates for the general public. Participants were reached out to online, and were not given the survey in a controlled environment; as such, their responses may have been influenced by unknown variables.

Survey

A primary limitation in the survey is the ambiguity of many of the questions used, which may have impacted participants' interpretations, and result in non-uniform understanding of what the question is asking. A prime example would be the crucial variable [interest], whose question *"I am interested in using this technology"*, would have been better formulated as *"I am interested in having this technology integrated into my society and daily life, as it has been described"*. The many interpretations of the statement may account for some discrepancy between [interest] values and [short_open] responses, though as previously mentioned, most of these pairs are in-line with each other.

Another limitation regards the relationship between the four paragraphs and their respective "themes". When first designed, each paragraph was modelled according to their labels, with the expectation that some paragraphs would influence some of the four traits more than others, producing a discernible relationship between the four traits and the paragraph themes. Although the hypothesis that the four traits are influenced separately per paragraph is true (see Appendix (c)), the inclusion of multiple themes across multiple paragraphs has made it difficult to discern any patterns. For this reason, the relationship between traits and themes was dropped from the analysis.

The paragraphs were written by the researcher, and may be imbued with his own biases (though this was not conscious, nor is this the assumption), as such participants may have been influenced by the paragraph in ways not accounted for. The paragraphs may not be reflective of their specific themes, as was the intention; this is up to interpretation.

Further Remarks And Future Research

A number of variables were not mentioned in either of the above analyses, and some of the statistical analyses conducted were removed from their respective sections. In this section, some further remarks regarding some of these will be explored, as well as suggestions for future research.

Included in the survey were questions on the participant's familiarity with cookies and browser fingerprinting, the hypothesis being that they would correlate well with [confidence], another variable dropped from the analysis, and that there may be a relationship between these variables and [interest]. No discernible relationship was found, bar a statistically-significant correlation between [cookies] and [confidence] of 0.582. [confidence] itself was dropped as it, too, did not appear relevant post-data-collection.
Participants were also asked about their use of an adblocker (options: Yes, No, I don't know what that is). Once again no discernible relationship was discovered.
Although the results from this study produced nothing conclusive in the domain of technological confidence v. privacy interests, this would be an interesting follow-up study to pursue. It may be the case that familiarity with technology affects one's interests in privacy, which would be of interest to policy makers wanting to address the needs of different groups of citizens.

Two boolean variables regarding surveillance were included in the survey. Participants were asked whether or not they agreed with a statement for each. The statements were designed to gauge participants' opinions towards mass surveillance used in the name of national security, [surveillance1], and mass surveillance in general, [surveillance2]. Once again these were included with the assumption that some relationship could be discerned between these variables and [interest], but none was found.
No relationship between [age], [sex], [nationality] and [interest] was found.

In summary, a survey was conducted to better understand the opinions that the general population might have towards an InferID-like system. The focus of the survey was to discern whether or not morality was a fundamental trait contributing towards participants' interest in using the technology. The 57 responses were analysed from both a quantitative and qualitative perspective. The analyses concluded that a mixture of opinions were felt towards InferID, though there was a clear "privacy concern" theme present in most of the responses, with few being wholly positive. The hypothesis that there is a moral element contributing to participants' interest in using the technology is supported by strong correlations between `[interest]` and the `[moral-]` variables. On average, participants who rated the system and its uses as immoral, also displayed a low interest in the technology overall.

Future research could incorporate the paragraphs in the survey, but present different paragraphs to different participants to garner a better understanding of the effects that certain "themes" have on participants' perceptions on InferID, and their overall interest. This would be a meaningful endeavour, especially to those wishing to discern what policies should be put in place in order to guide the development of ethical technologies in this domain.

7. Conclusion

This paper explored the question *"Could behavioural biometrics be used to identify individuals without their knowledge or consent, and what moral and societal implications would arise from its adoption?"*. The first four sections of the paper were dedicated to the technological aspects of the system; the "could this happen" aspect of the question. In the latter part of the paper, the "should this happen" aspect was explored.

The paper opened by introducing "InferID" the hypothetical behavioural biometric system of the (near) future. A literature review was conducted in Section 2, which analysed pre-existing behavioural biometric authentication methods to better understand what InferID may be capable of. Behavioural biometrics are currently being explored within the domain of *authentication*, whereas InferID concerns itself with *identification*. The design of an example InferID system as outlined in Section 3a, which used an artificial neural network to identify mobile users using their scrolling behaviour. Various applications for InferID were also addressed (3b), as well as some of the limitations (3c) already present in existing behavioural biometric systems, that may be unresolved for InferID. Section 4 detailed potential mitigation methods that could be used to reduce InferID's capabilities, as an exploration into the system's robustness.
As of yet, a functioning InferID-like system has not been deployed, at least not publicly. Yet the promise of such works as "Touchalytics" and "SilentSense", and advancements in machine learning make it seem plausible. Existing behavioural biometric systems suggest that InferID-like systems exist in the realm of possibility.

In Section 5 the implications of such a system were outlined, discussing both its blessings and its curses. The concerns raised in Section 5 were echoed in the results obtained in

Section 6, which detailed the quantitative and qualitative analyses of a survey conducted about InferID. The dangers that such a system presents and its privacy-violating nature appear to cause an aversive reaction in most of the participants, as the survey revealed. The concept of privacy was a recurring theme in the data, and from the correlations, it was inferred that a moral element influenced individuals' overall interest in using InferID. The paragraphs detailing InferID and its various use-cases had varying effects on the four traits, which suggests that InferID may be considered more or less moral in specific contexts. Future research may explore this further, for the sake of developing appropriate policies to govern the use of such systems.

Technology evolves with unrivaled rapidness, and halting it has proven both difficult and unwise. Societies must be prepared for new technologies to arise, and policy makers need to be at the ready, to ensure that powerful, innovative tech is put to apt use, and is developed for the greater good. Raising awareness for powerful technologies is a step towards developing policies that could prevent this power being abused. InferID and technologies like it may exist, and pose a threat to individual privacy. In a world where mass surveillance is becoming ever-more prevalent, policy makers need to consider the ethical and societal concerns of the public, and the dangers posed by emerging technologies.

8. References

Al-Fannah, N. M., & Mitchell, C. (2020). Too little too late: Can we control browser fingerprinting? *Journal of Intellectual Capital*, *ahead-of-print*(ahead-of-print). https://doi.org/10.1108/JIC-04-2019-0067

Antal, M., & Szabó, L. Z. (2016). Biometric Authentication Based on Touchscreen Swipe Patterns. *Procedia Technology*, *22*, 862–869. https://doi.org/10.1016/j.protcy.2016.01.061

Bo, C., Zhang, L., Li, X.-Y., Huang, Q., & Wang, Y. (2013). SilentSense: Silent user identification via touch and movement behavioral biometrics. *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, 187–190. https://doi.org/10.1145/2500423.2504572

Boerman, S. C., Kruikemeier, S., & Borgesius, F. J. Z. (2017). Online Behavioral Advertising: A Literature Review and Research Agenda. *Journal of Advertising*, *46*(3), 363–376. https://doi.org/10.1080/00913367.2017.1339368

Brown, M., & Rogers, S. J. (1993). User identification via keystroke characteristics of typed names using neural networks. *International Journal of Man-Machine Studies*, *39*(6), 999–1014. https://doi.org/10.1006/imms.1993.1092

Chrome DevTools | Tools for Web Developers. (n.d.). Google Developers. Retrieved 8 June 2020, from https://developers.google.com/web/tools/chrome-devtools

Co, K. T., Muñoz-González, L., de Maupeou, S., & Lupu, E. C. (2019). Procedural Noise Adversarial Examples for Black-Box Attacks on Deep Convolutional Networks. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 275–289. https://doi.org/10.1145/3319535.3345660

Connor, P., & Ross, A. (2018). Biometric recognition by gait: A survey of modalities and features. *Computer Vision and Image Understanding*, *167*, 1–27. https://doi.org/10.1016/j.cviu.2018.01.007

Cornière, A. de, & Nijs, R. de. (2016). Online advertising and privacy. *The RAND Journal of Economics*, *47*(1), 48–72. https://doi.org/10.1111/1756-2171.12118

Crammer, K., & Singer, Y. (2002). On the Algorithmic Implementation of Multiclass Kernel—Based Vector Machines. *Journal of Machine Learning Research*, *2*(2), 265.

Dancey, C. P., & Reidy, J. (2007). *Statistics Without Maths for Psychology*. Pearson Education.

Dargan, S., & Kumar, M. (2020). A comprehensive survey on the biometric recognition systems based on physiological and behavioral modalities. *Expert Systems with Applications*, *143*, 113114. https://doi.org/10.1016/j.eswa.2019.113114

Detecting device orientation. (n.d.). MDN Web Docs. Retrieved 15 May 2020, from https://developer.mozilla.org/en-US/docs/Web/API/Detecting_device_orientation

Firefox Developer Tools. (n.d.). MDN Web Docs. Retrieved 8 June 2020, from https://developer.mozilla.org/en-US/docs/Tools

Frank, M., Biedert, R., Ma, E., Martinovic, I., & Song, D. (2013). Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication. *IEEE Transactions on Information Forensics and Security*, *8*(1), 136–148. https://doi.org/10.1109/TIFS.2012.2225048

Goldfarb, A., & Tucker, C. E. (2011). Online advertising, behavioral targeting, and privacy. *Communications of the ACM*, *54*(5), 25–27. https://doi.org/10.1145/1941487.1941498

Goth, G. (2011). Privacy Gets a New Round of Prominence. *IEEE Internet Computing*, *15*(1), 13–15. https://doi.org/10.1109/MIC.2011.17

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The Weirdest People in the World? In *Working Paper Series of the German Council for Social and Economic Data* (No. 139; Working Paper Series of the German Council for Social and Economic Data). German Council for Social and Economic Data (RatSWD). https://ideas.repec.org/p/rsw/rswwps/rswwps139.html

Horn, R. V. (2016). Cookies, Web Profilers, Social Network Cartography and Proxy Servers—Royal Van Horn, 2004. *Phi Delta Kappan*. https://journals-sagepub-com.proxy-ub.rug.nl/doi/10.1177/003172170408600304

Iqbal, U., Shafiq, Z., & Qian, Z. (2017). The ad wars: Retrospective measurement and analysis of anti-adblock filter lists. *Proceedings of the 2017 Internet Measurement Conference*, 171–183. https://doi.org/10.1145/3131365.3131387

Jenkins, J., Nguyen, Q., Reynolds, J., Horner, W., & Szu, H. (2011). The physiology of keystroke dynamics. *Proceedings of SPIE - The International Society for Optical Engineering*, *8058*. https://doi.org/10.1117/12.887419

Kamkar, S. (n.d.). *Samy Kamkar—Evercookie—Virtually irrevocable persistent cookies*. Retrieved 3 June 2020, from https://samy.pl/evercookie/

Kosinski, M., Bachrach, Y., Kohli, P., Stillwell, D., & Graepel, T. (2014). Manifestations of user personality in website choice and behaviour on online social networks. *Machine Learning*, *95*(3), 357–380. https://doi.org/10.1007/s10994-013-5415-y

Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(15), 5802–5805. JSTOR.

Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*(24), 8788–8790. https://doi.org/10.1073/pnas.1320040111

Lear, S. (2017). The Fight over Encryption: Reasons Why Congress Must Block the Government from Compelling Technology Companies to Create Backdoors into Their Devices Notes. *Cleveland State Law Review*, *66*(2), [i]-476.

Leiba, B. (2012). OAuth Web Authorization Protocol. *IEEE Internet Computing*, *16*(1), 74–77. https://doi.org/10.1109/MIC.2012.11

Li, G., Borj, P. R., Bergeron, L., & Bours, P. (2019). Exploring Keystroke Dynamics and Stylometry Features for Gender Prediction on Chat Data. *2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, 1049–1054. https://doi.org/10.23919/MIPRO.2019.8756740

Liang, F., Das, V., Kostyuk, N., & Hussain, M. M. (2018). Constructing a Data-Driven Society: China's Social Credit System as a State Surveillance Infrastructure. *Policy & Internet*, *10*(4), 415–453. https://doi.org/10.1002/poi3.183

Lusthaus, J. (2012). Trust in the world of cybercrime. *Global Crime*, *13*(2), 71–94.

https://doi.org/10.1080/17440572.2012.674183

Nahin, A. F. M. N. H., Alam, J. M., Mahmud, H., & Hasan, K. (2014). Identifying emotion by keystroke dynamics and text pattern analysis. *Behaviour & Information Technology*, *33*(9), 987–996. https://doi.org/10.1080/0144929X.2014.907343

Nairn, A., & Berthon, P. (2003). Creating the Customer: The Influence of Advertising on Consumer Market Segments: Evidence and Ethics. *Journal of Business Ethics*, *42*(1), 83–99. JSTOR.

Pires, I., Garcia, N., Pombo, N., Flórez-Revuelta, F., & Spinsante, S. (2018). Approach for the Development of a Framework for the Identification of Activities of Daily Living Using Sensors in Mobile Devices. *Sensors*, *18*, 640. https://doi.org/10.3390/s18020640

Sieu, B., & Gavrilova, M. (2020). Biometric Identification from Human Aesthetic Preferences. *Sensors*, *20*(4), 1133. https://doi.org/10.3390/s20041133

Stanton, J. V., & Guion, D. T. (2013). Taking Advantage of a Vulnerable Group? Emotional Cues in Ads Targeting Parents. *Journal of Consumer Affairs*, *47*(3), 485–517. https://doi.org/10.1111/joca.12018

Tiku, N. (2017, May 21). Welcome to the Next Phase of the Facebook Backlash. *Wired*. https://www.wired.com/2017/05/welcome-next-phase-facebook-backlash/

Watts, G. (2008). Google watches over flu. *BMJ*, *337*. https://doi.org/10.1136/bmj.a3076

Yu, J., Zhou, H., & Gao, X. (2015). Machine learning and signal processing for human pose recovery and behavior analysis. *Signal Processing*, *110*, 1–4. https://doi.org/10.1016/j.sigpro.2014.11.008

Appendix

a. Survey Questions

Question	Variable [var]	Type
Section 1		
Age	age	number
Sex	sex	options: ● Female ● Male ● Prefer Not To Say ● Other (*specify)*
Nationality	nationality	text
How confident are you with technology?	confidence	likert
Section 2		
Paragraph 1		
This technology is important	important1	likert
This technology is inconvenient	accessible1	likert
This technology is unethical	moral1	likert
This technology is safe	safe1	likert
Paragraph 2		

This technology is simple to use	accessible2	likert
This technology is uninteresting	important2	likert
This technology is secure	safe2	likert
This technology is unacceptable	moral2	likert
Paragraph 3		
This technology is meaningful	important3	likert
This technology is justifiable	moral3	likert
This technology is threatening	safe3	likert
This technology is accessible	accessible3	likert
Paragraph 4		
This technology is dangerous	safe4	likert
This technology is challenging to use	accessible4	likert
This technology is moral	moral4	likert
This technology is irrelevant	important4	likert
I am interested in using this technology	interest	likert
I would describe this technology as	short_open	text
Do you have any further remarks regarding InferID?	long_open	text

Section 3		
How familiar are you with website cookies?	cookies	likert
How familiar are you with browser fingerprinting?	browser_fingerprinting	likert
Do you use an adblocker	adblocker	options: ● Yes ● No ● I don’t know
Please read the following statement and select the answer closest to your opinion: “Governments and Law Enforcement Agencies should have the right to access any digital device or account, and de-anonymise users online. To prevent terrorist attacks and for the general safety and security of the population.”	surveillance1	options: ● Agree ● Disagree
Please read the following statement, regarding surveillance, and select the answer closest to your opinion: "If you have nothing to hide, then you have nothing to fear."	surveillance2	options: ● Agree ● Disagree

b. Paragraphs And Themes

Index	**Themes**	**Paragraph**
1	● Accessibility ● Personal Security	InferID is a new technology that uses your unique online behaviour to automatically log you into your favourite apps and websites. It keeps your accounts secure and makes it easier to log in than ever before. By using your unique scrolling behaviour, like a fingerprint, InferID can seamlessly keep you logged in without ever asking for a password. This makes it easier for you to access your accounts, and harder for anyone else to log into them. Passwords will no longer be a requirement.
2	● Behavioural Observation ● Identification ● Covertness	InferID analyses your scrolling behaviour while you access your favourite apps and websites, and continuously infers your

	● Personal Security ● Omnipresence ● Accessibility ● Advertising	identity (hence the name InferID). As you scroll, it works in the background, out of your way, to keep you logged in and secure. InferID works across a large number of apps and websites. This means all your accounts can stay logged in no matter what you access. It can seamlessly track your shopping habits, or what videos you like to watch, and help advertisers recommend the right products for you.
3	● National Security ● Omnipresence	InferID is being used by law enforcement across the globe. When law enforcement gain access to a black-market website, they can put InferID in the background, unnoticed. Because InferID already works on so many other sites, it is easy to connect the anonymous users of the black-market with their real-life counterparts — helping law enforcement to identify and eventually take down criminals online. This includes cyber-crimes (hacking) as well as drug / human trafficking, illegal arms trade, etc.
4	● Omnipresence ● Inescapability ● National Security ● Covertness	InferID will become mandatory within the next two years. Every new device will have it running seamlessly in the background, regardless of the app or website. Doing so allows governments around the globe to know who is using a device at any point in time, and how they are using it. Unlawful use of a device will become a thing of the past. All your accounts on the Internet will be tied to your identity.

c. ANOVA Repeated Measures | Traits v. Paragraph

ANOVA Repeated Measures | Traits v. Paragraph				
	F	**Num DF**	**Den DF**	**Pr > F**
important	9.22	3	168	0.00

accessible	11.50	3	168	0.00
moral	28.88	3	168	0.00
safe	26.99	3	168	0.00

d. Boxplots Of Traits

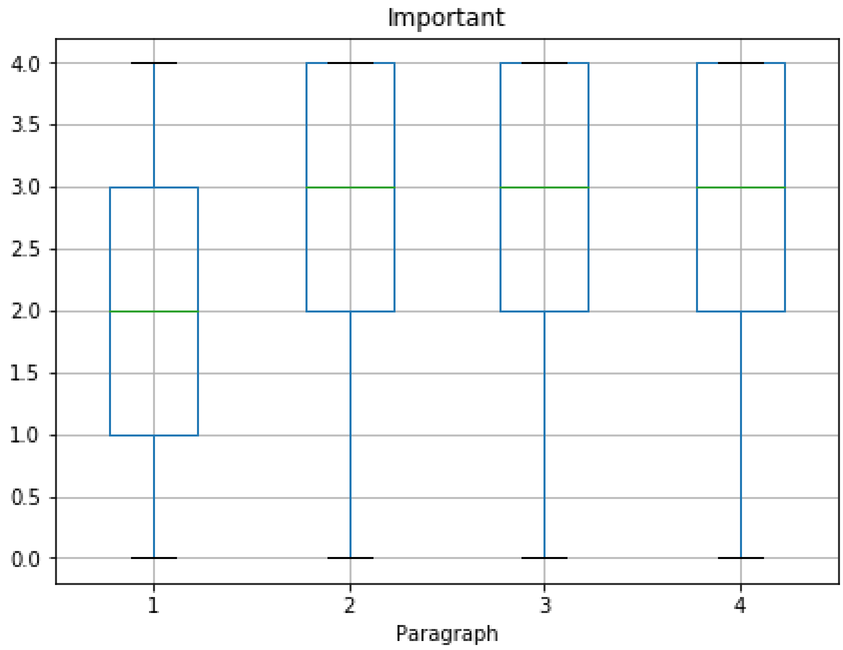

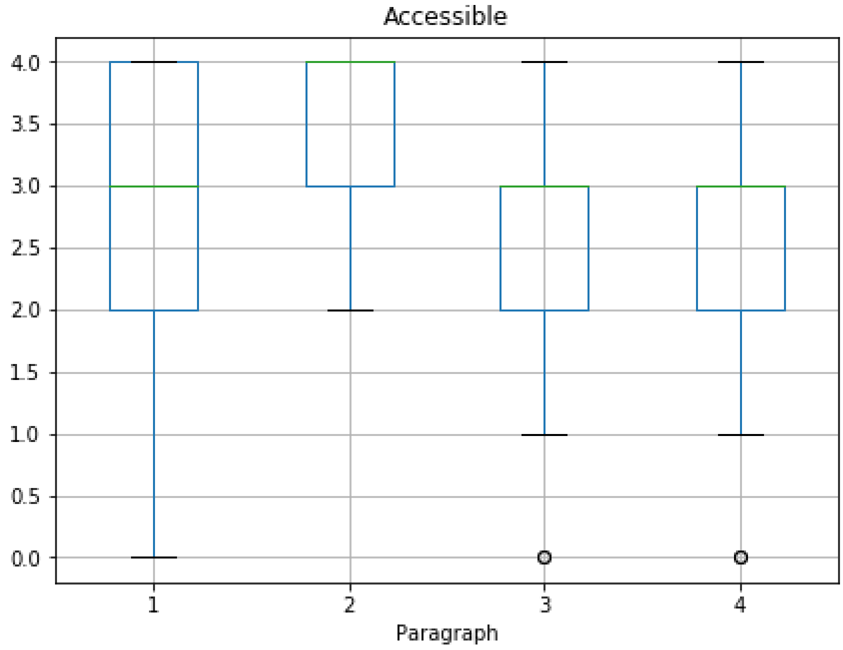

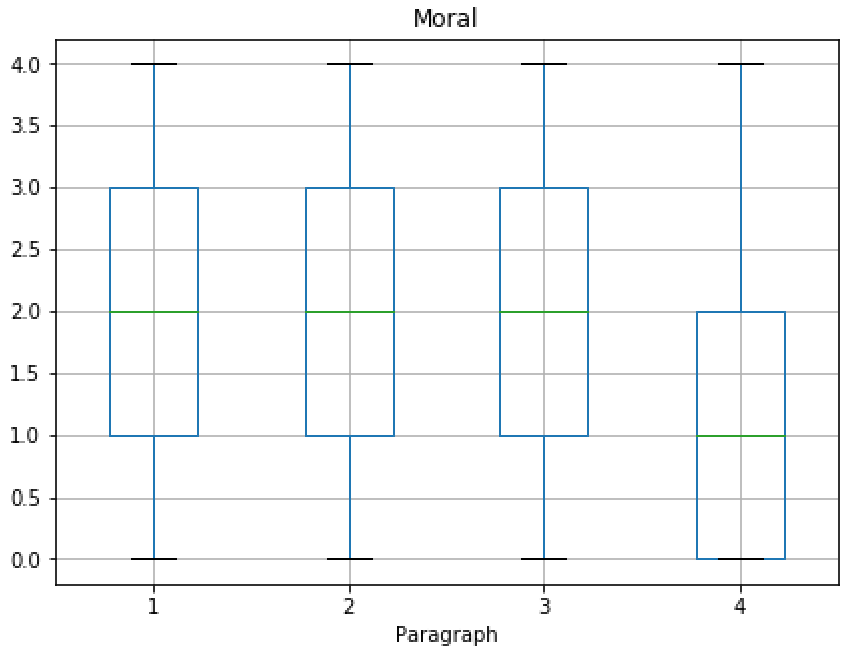
Moral
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
1
2
3
4
Paragraph

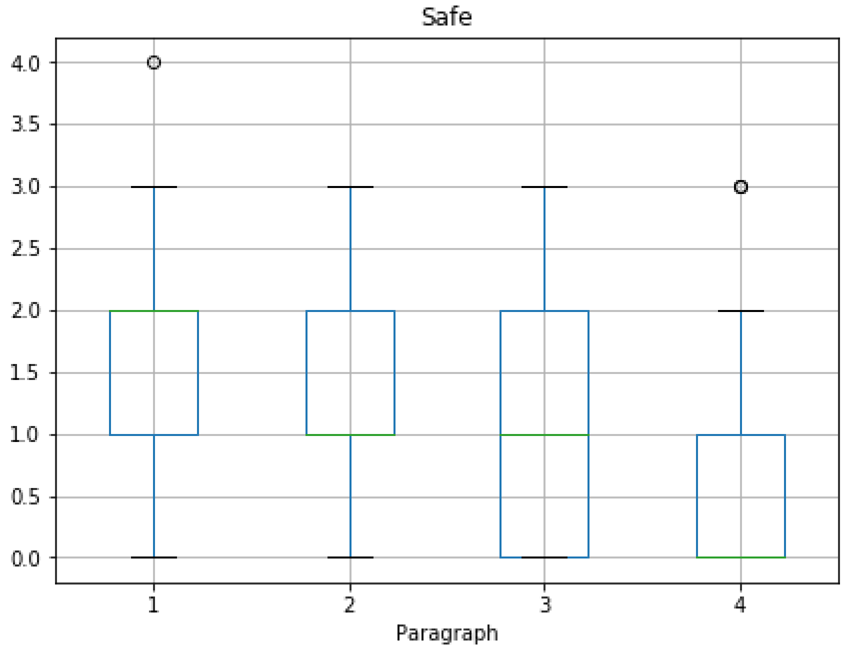
Safe
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
1
2
3
4
Paragraph

e. Spearman Correlations Traits v. [interest]

[abs(r) >= 0.5] & [p < 0.05]

Spearman Correlations Traits v. [interest]

trait	paragraph	r	p
important	1	0.379	0.004
important	2	0.513	0.000
important	3	0.353	0.007
important	4	0.211	0.116
accessible	1	0.327	0.013
accessible	2	-0.057	0.672
accessible	3	0.234	0.079
accessible	4	-0.172	0.201
moral	1	0.505	0.000
moral	2	0.777	0.000
moral	3	0.619	0.000
moral	4	0.736	0.000
safe	1	0.261	0.050
safe	2	0.487	0.000
safe	3	0.589	0.000
safe	4	0.524	0.000

f. Word Cloud

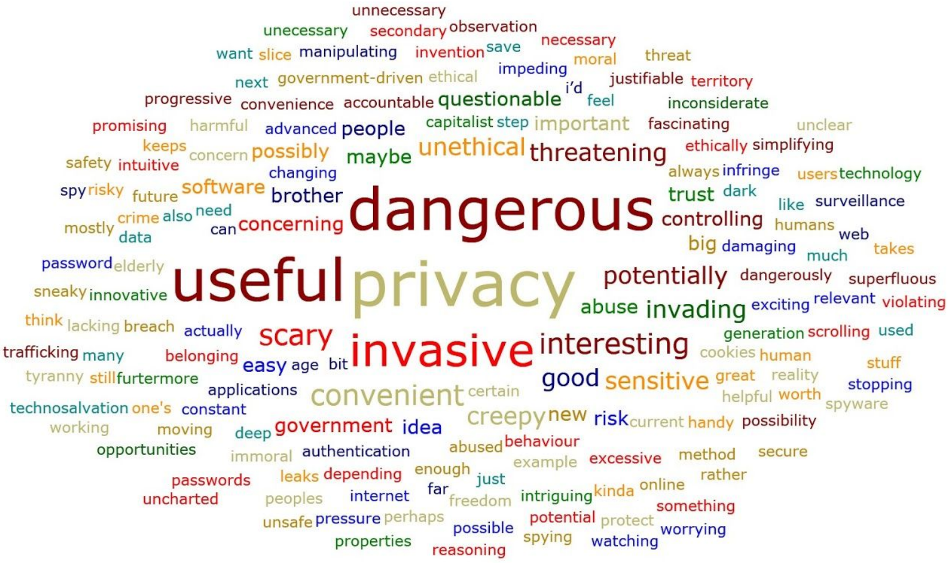

** The size of the word is representative of its frequency in the* [short_open] *texts.*

g. Samples Of [short_open] By [interest]

[short_open] sample where [interest = 0]	
id	**short_open**
53	dangerous, unethical, impeding, violating
25	threatening, dangerous, unethical
18	invasive
23	Worrying, a concern for privacy, intriguing as possibility, concerning as reality
26	Creepy, government-driven

[short_open] sample where [interest = 1]	
id	**short_open**
29	Invasive, useful
43	Helpful, invasive
14	Potentially dangerous, potentially useful, potentially immoral depending on its users
52	Lacking privacy
33	Risky, easy to abuse, useful for certain applications

[short_open] sample where [interest = 2]

id	short_open
48	Not much privacy, constant surveillance. Some use can be justifiable, some can't
20	dangerous, the idea is good i just don't trust humans enough to use it
21	In between. In between moral reasoning, freedom, maybe people won't want this "always working" software because then they would actually feel under observation and pressure
40	Innovative, something belonging to the new age, the next step moving from cookies, ethically sensitive.
9	Useful, mostly ethical, spying, keeps people accountable

[short_open] sample where [interest = 3]

id	short_open
17	interesting, exciting, possibly unsafe
11	Relevant, threatening, invasive
8	Useful,
24	Interesting, useful properties, harmful opportunities, privacy risk
5	Convenient but also invasive