Computationally Mediated Pro-Social Deception

Max Van Kleek
Dept. of Computer Science
University of Oxford, UK
emax@cs.ox.ac.uk

Dave Murray-Rust
School of Informatics
University of Edinburgh, UK
d.murray-rust@ed.ac.uk

Amy Guy
School of Informatics
University of Edinburgh, UK
Amy.Guy@ed.ac.uk

Kieron O'Hara
Web and Internet Science
Univ. of Southampton, UK
kmo@ecs.soton.ac.uk

Nigel R. Shadbolt
Dept. of Computer Science
University of Oxford, UK
nigel.shadbolt@cs.ox.ac.uk

ABSTRACT

Deception is typically regarded as a morally impoverished choice. However, in the context of increasingly intimate, connected and ramified systems of online interaction, manipulating information in ways that could be considered deceptive is often necessary, useful, and even morally justifiable. In this study, we apply a speculative design approach to explore the idea of tools that assist in *pro-social* forms of online deception, such as those that conceal, distort, falsify and omit information in ways that promote sociality. In one-on-one semi-structured interviews, we asked 15 participants to respond to a selection of speculations, consisting of imagined tools reifying particular approaches to deception. Participants reflected upon potential practical, ethical, and social implications of the use of such tools, revealing a variety of ways such tools might one day encourage polite behaviour, support individual autonomy, provide a defence against privacy intrusions, navigate social status asymmetries, and even promote more open, honest behaviour.

Keywords

Deception; disinformation; speculative design; autonomy; privacy

Categories and Subject Descriptors

H.5.m. [**Information Interfaces and Presentation (e.g. HCI)**]: Miscellaneous

1. INTRODUCTION

Most people like to consider themselves to be quite honest in their communications with friends, family and acquaintances. However even honest people routinely modulate what they share, omitting or even falsifying information in order to reduce social friction, avoid confrontation, diffuse awkward situations, or to save face [14, 15]. Hancock et. al. introduced the term *butler lies* to refer to a common use of simple lies to manage communications, such as to smoothly exit from an unwanted conversation [33]. Online, the notion of who our 'friends' are has become increasingly blurred and difficult to define. In such settings, people commonly navigate different social spaces, projecting and varying self-presentation according to the ways they want to be perceived by each [40].

Whilst part of tailoring one's presentation to an audience is the ability to carry out some level of *deception*, with personal communications, there is an implicit expectation of authenticity [7]. However, online, the need to navigate multiple and uncertain audiences means that we may constantly vary our self-presentation. Authenticity becomes a social construct derived from the social context and how we wish to be perceived by a given audience [12]. We may be deceiving, at least to some extent, nearly constantly without even being conscious of it.

The use of deception as a technique for system designers has been discussed previously within the HCI community. For example, manipulation of users' mental models of systems in ways that benefit both systems' designers and end-users were documented by Adar et al. [1]. Ambiguity, often promoted through deception, gives people space for flexible interpretation [29], and to tell stories they need to in order to preserve face and reputation [7, 10]. However, the complexity of modern social software means that a growing cast of actors have to be considered, both human and computational, as targets, confederates, dupes and adversaries for any action.

Here, we are interested in exploring the complex contexts in which deception might take place, to consider not just cases where the system lies to a user [1] or computer mediated communication where one user lies to others, but situations where systems lie to each other about users; where a user needs to lie to one audience but not another; where tools or systems might protect a person from disclosure to other systems or tools. As Nissenbaum puts it:

> Those who imagined online actions to be shrouded in secrecy have been disabused of that notion. [...] We have come to understand that even when we interact with known, familiar parties, third parties may be lurking on the sidelines, engaged in business partnerships with our known parties. [44]

The actors involved now include not just the people who are being immediately addressed, but others who are peripheral or incidental to the interaction as it occurs. Many systems include silent 'lurkers', who observe without speaking. Others will discover and read conversations later, outside the context of their production. Beneath the visible surface of the communications tools people use, a growing series of invisible actors mine the interaction data which occurs on their platforms, and others use the results of this mining. Many of these actors are computational systems of increasing power, sifting, sorting, re-purposing and inferring from the full

CHI'16, May 07-12, 2016, San Jose, CA, USA

DOI: http://dx.doi.org/10.1145/2858036.2858060

spectrum of communicative data.

The contributions of the paper are as follows:

- A summary of recent work on deception in HCI, with a focus on its use in systems and tools;
- An expansion upon previous models of *computer-mediated social deception* with new configurations, in which tools conduct or facilitate deception towards other people/systems/tools;
- A description of a speculative design experiment in which reflections on fictional tools for social deception were elicited;
- A characterisation of the practical, ethical, moral, & social perspectives on the use of such tools, along with design guidelines for future tools employing deception in social contexts.

2. BACKGROUND

Deception has long been studied, both within and outwith the HCI community. Traditionally, deception has been cast in a negative light [11], to be used only if no other option is available. In the 1980s, however, communications researchers began to investigate the positive aspects of lying, in particular *white lies*—socially acceptable lies which cause little or no harm to the recipient [16]. In 1992, McCornack cast deception as an understandable response to complexity: "[r]esearchers studying deception recently have begun to argue that deceptiveness is a message property that reflects a kind of functional adaptation to the demands of complex communication situations" [41]. People then manipulate the information which they share as a necessary part of participation in society. This has led to recent work on the positive aspects of deception in human computer interaction, in particular how 'butler lies' are used to ease social situations [33], and how systems can deceive their users for beneficial reasons [1].

Several different taxonomies of lying and deceptive behaviours have been proposed [16, 21, 38]; Anolli et al. examined a family of deceptive miscommunications, including self-deception and white lies [6]. They look at *omission* of relevant information, *concealment* using diversionary information, *falsifaction* and *masking* with alternative, false information. Of particular interest is their claim that "a deceptive miscommunication theory should be included in a general framework capable of explaining the default communication", that is that deception should not be seen as a psychologically different activity than 'normal' communication. This tallies with the earlier approach of McCornack [41] who situates deceptive messages within the spectrum of *information manipulation*. This, combined with the lens of Gricean maxims, allows for an explanation of deceptions where some of the truth is told, but information which the speaker knows is relevant to the listener is omitted or obscured [31].

Motivations for lying have also been extensively studied in social psychology. Turner et al.'s taxonomy included *saving face*; guiding social interaction; avoiding tension or conflict; affecting interpersonal relationships; and achieving interpersonal power [50]. Camden et. al. [16] develop a detailed categorisation of lies to do with basic needs, managing affiliation with others, self-esteem and miscellaneous practices such as humour and exaggeration. A recent study of online behaviour found that the most common self-reported motivation for online lies was either to make one's life seem more exciting, or to downplay personal difficulties. Responses also included avoiding harassment and a range of creative endeavours alongside more clearly adversarial deceptions [34].

Another strand of research borrows from information warfare, to look at the possibilities for *disinformation*. Disinformation tactics are most useful when a channel of information cannot be completely closed, but can be rendered useless by being filled with incorrect, but plausible, assertions in order to lower its overall signal-to-noise ratio [51]. The intended target of the lie may not be the official recipient of the message: lies can be directed at those who are eavesdropping on the communications channel or surveilling the participants [5]. Techniques used include *redaction* to remove parts of the message, *airbrushing* to blur parts of the message and *blending* to make the message similar to other plausible messages, as well as other forms of *information distortion* [5].

2.1 Ambiguity, Distance, Social Privacy

These properties of communication channels—the transparency, and the amount of context which is conveyed relate to notions of distance. Birnholz et. al [10] look at different aspects of ambiguity in setups ranging from radically co-located to physically separated teams. They found that people who were co-located manage the release of information in order to maintain a sense of autonomy. Ambiguity was used to allow the hearer to believe a particular story, with social constructs forbidding intrusiveness being leveraged to maintain the space for ambiguity—for example, a norm against 'screen-surfing' and looking at a colleagues monitor allows a flexible explanation of exactly what one is working on.

Aoki and Woodruff [7] pick up on a need for ambiguity within personal communication systems, not for explicit lies, but to allow the participants space in which to construct mutually agreeable stories. If one's online activity—or read receipts—are visible, the kind story about being too busy to reply becomes problematic. This impinges on our ability to carry out *face-work*, and project desired images. Gaver et. al. examine different types of ambiguity [29], of information, context and relationship, and suggest avenues for their use in HCI—the completely unambiguous "Seen at 12:57pm" could be altered in many ways to soften it and allow more space for interpretation.

Burgoon et. al delineates four different dimensions of privacy: *physical*, being free from surveillance and intrusions into one's space; *social* or *interactional*, controlling the 'who, what, when and where' of encounters; *psychological*, freedom to introspect, analyse and so on, and freedom from persuasive pressures; and *institutional*, the ability to control who gathers what information about oneself and under what circumstances [15]. Raynes-Goldie [46] finds that while young people are happy to abandon institutional privacy to pragmatism, the social aspects of privacy remain tightly held.

The *social* aspects of privacy relate to what DeCew terms *expressive* privacy—a freedom from peer pressure and an ability to express one's own identity [20]. Nissenbaum's contextual integrity [44, 45] seeks to understand appropriate sharing, looking at the ways in which flows of information are governed by norms, which may easily violated as technological systems repurpose and share data.

2.2 Pervasive Surveillence and Privacy Tools

We are rapidly moving into a world where information about nearly every aspect of our lives is becoming sensed, recorded, captured and made available in digital form. Data is captured and shared voluntarily, as tools invite ever more intimate participatory surveillance [4]. While the abundance of information traces has unlocked a wide range of new kinds of applications (eg. [3] [19]), the creation and potential for disclosure poses new threats to individual privacy and autonomy. The overall lack of transparency by manufacturers regarding how they are capturing and handling personal information has created a heightened sense of unease among many, in addition to the potential threats dealing with their unintentional

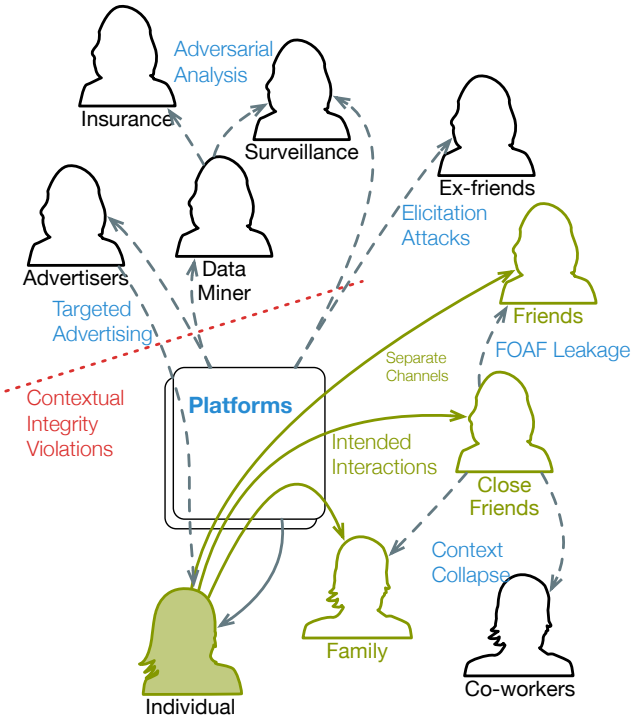

Figure 1: Common information transmission vectors in a standard computer-mediated communication setting. Platform(s) represents carriers/service providers enabling the communication; blue arrows denote channels controlled by the individual, while dashed arrows denote invisible channels out of their control. Communication with a particular group of people may end up being shared with others through context collapse [12], or through leakage to friends of friends. Ex-friends can use social engineering to elicit data which was not shared with them, or carry out cloaking attacks [39]. Data may be mined and analysed, often violating contextual integrity [44], and repurposed for use in advertising, surveillance and so on.

disclosure or misuse [27, 42, 26].

Many tools — indeed, entire research fields — have been dedicated to researching tools that carry out various kinds of digital deception for the purpose of protecting the individual's privacy. Without aspirations of comprehensiveness, we mention some here. Tools for masking identity are currently available for all levels of the software stack, from tools like *tor* for masking the origin and destination at the network level [22], to privacy-enhancing features at the browser level. Such browser features include *Do Not Track* [49], user-agent spoofing, and tracker and cookie-blocking capabilities [25]. At the application level, anonymous e-mail remailers [32], anonymous e-Cash and cryptocurrencies [17], and anonymous secure file sharing systems [47] have started to support certain activities offering guarantees of privacy under specified conditions.

3. DECEPTION IN MEDIATED SOCIAL SITUATIONS

One of the striking aspects of deception is how little it changes with the advent of computationally mediated communications. The added distance may allow people to lie more, and justify to themselves more easily [43], but many of the motivations and techniques remain similar.

However, one of the key differences is in the context in which deception takes place. Mediated communication brings an opportunity for many different structures of deception, for several reasons:

- Imagined audiences [40] and understanding of publics in digital space are increasingly complicated.
- The individual may wish to provide false information to the communications platform where the interaction is taking place, for reasons including privacy, mistrust of the platform provider, or dislike of targeted advertising.
- The individual may wish to manipulate the secondary data which is derived from their actions, such as controlling the summary data given to their insurance company.
- Deceptions can work in either or both directions: platforms may deceive some or all of their users, autonomously or due to the will of their designers and commissioners.
- People often communicate with platforms through some intermediary, such as an app on a mobile phone. These intermediaries can deceive the platform on behalf of the user, especially about what information is being automatically collected (eg. through sensors).
- As well as being targets of lies, others can be enlisted to lend credence to statements, for instance supporting alibis, agreeing that the network is down at the moment, and so on.

Some of these actors are shown in Figure 1, and based on this, Figure 2 shows some structures, along with references to systems which embody each configuration.

4. STUDY DESIGN

In this study, we wished to explore various positive uses for future tools that employ computer-mediated deception. Specifically, we wished to identify ways that such tools might positively facilitate the maintenance of sustained, positive social relationships with others, especially in complex social environments, both in online and hybrid online/offline settings. With respect to deception, we adopted McCornack's definition from information manipulation theory [41], which includes both the introduction of false information, as well as selective information disclosure, such as (but not limited to) for purposes of creating ambiguity, or selective identity (i.e. omitting information conflicting with one's desired presentation).

4.1 Materials, Method and Recruitment

In order to start to understand the practical, ethical and social implications of the use of computationally-mediated deception in social settings, we organised a study to elicit perspectives and experiences from people from a variety of backgrounds, around some of the deception configurations imagined in 2. Drawing inspiration from critical design [9], which takes a critical theory [18] approach to speculative design, we first generated a series of speculative design proposals [24] in the form of realistic depictions and descriptions of imagined, "near future" privacy tools. These fictional privacy tools, with accompanying descriptions, which will

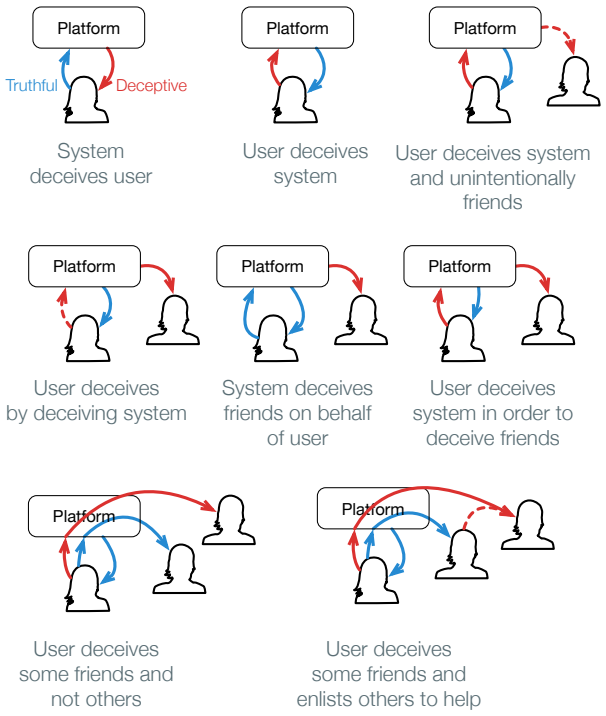

Figure 2: *Structures of deception* - Depiction of possible deception configurations involving 1 platform and up to two parties.

henceforth be referred to as *vignettes*, were then showed to participants in semi-structured interview settings.

We recruited participants via Twitter, open Facebook groups, and word-of-mouth through personal connections. Those interested first answered demographic questions covering age, gender, employment status and frequency of use of social media. To ensure diversity in participants, fifteen (aged 18+) were selected in a way that maximised saturation on the attributes collected. Interviews were conducted in person and via video chat. At the start of interviews, participants were asked an opening question, "How do you feel about your privacy online?" which was used to their general attitudes and sensitivity towards privacy online. Then, twwo interview questions were asked of each vignette, first, whether the individual would consider using a tool like the one described (and why or why not), and second, whether their interactions online would change if they found out their friends were using tools like the one described. Finally, participants were encouraged to share thoughts or personal experiences that they were reminded of by the vignette.

Audio from sessions was recorded, transcribed and anonymised for identifiers of people, places and entities. Inductive thematic analysis was carried out on the transcripts by analysing and coding them for themes, by three researchers independently. Themes were then compiled, combined into a single pool, and discussed to derive a final coherent set of themes. Then, related themes were clustered into groups. We organise the result section according to these groups.

4.2 Designing the Vignettes

We identified two main axes to guide us when crafting the speculative vignettes. The first was the degree to which machines mediated the deception; from tools that simply facilitated, otherwise manual acts of deception, to those that entirely automated it, and

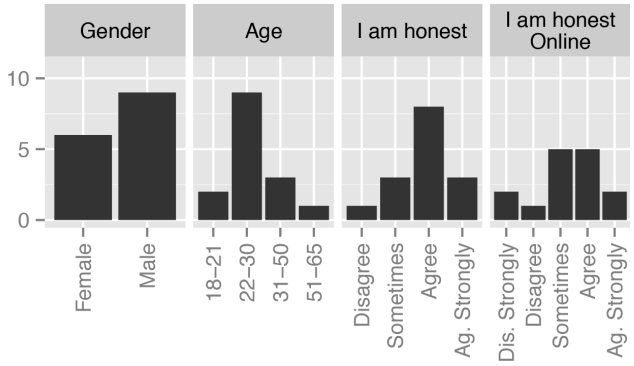

Figure 3: Participant demographics and responses to questions pertaining to self-perception of honesty on and offline.

including those that helped people in the post-hoc maintenance of past deceptions. The second axis, inspired by Gaver's *conceptual design proposals*, seeking to explore the "balance between concreteness and openness: [...] specific enough to evoke intuitive reactions, yet indefinite enough to encourage imaginative extensions" [28]. With respect to realism, we wanted to aim for tools that would be realisable in the near future, inspired by AugerâĂŹs *speculative designs*: "speculative designs exist as projections of the lineage, developed using techniques that focus on contemporary public understanding and desires, extrapolated through imagined developments of an emerging technology" [8].

With these axes and guidelines, we generated two dozen candidate ideas, and selected five that met the above criteria, were the most plausible, and that best covered the space spanned by design axes just described. To break ties, we preferred simpler scenarios, to encourage participants to focus on implications rather than the tools themselves. This process resulted in the following final five vignettes:

Social Steganography (Figure 4) inspired by danah boyd's studies of networked teens [13] that used in-group codes to discuss activities so that they were inscrutable to their parents. Here, the steganography is performed automatically: a trusted set of people see the 'real' message, while everyone else sees an 'innocent', socially plausible message.

lieCal (Figure 5) creates fictitious appointments based on common diary structure, to automate the process of deploying butler lies. Friends can be enlisted to give weight to the lie, and corroborating evidence is posted on social networks.

lieTinerary (Figure 6) draws on Merel Brugman's *Same Same But Different*, enables the pre-curation of a fictitious trip or fictional event attendance through pre-scheduled, coordinated posts across multiple social media platforms.

lieMoves (Figure 7) is a fictional service for mobile phones that replaces the user's actual location with data from user-selectable and customisable deception strategies: blurred (low-grain), superposition of locations, past replay, or "typical" herd-behaviour or individual simulation.

lieMapper (Figure 8) shows the interconnectedness of communication channels. Extending Facebook's 'this post will go to X people' functionality, it works across multiple networks to visualise all those within one's friend networks likely to hear about a particular piece of information.

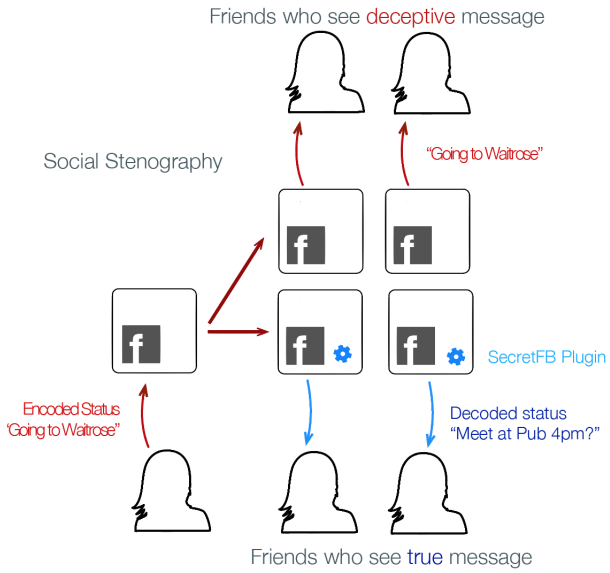

Figure 4: *Social Stegonagraphy*: Diagram illustrating a social steganography tool for microblogging/SNS sites that hides "real" messages behind other, plausible status messages but allows certain people to recover the true meaning.

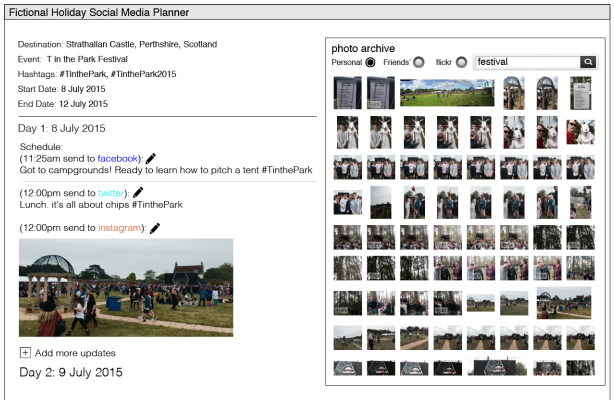

Figure 6: *lieTinerary*: Fictional tool to create a narrative of going somewhere (on holiday) or attending an event, along with images and social media posts to be sent out at preset times to corroborate the story.

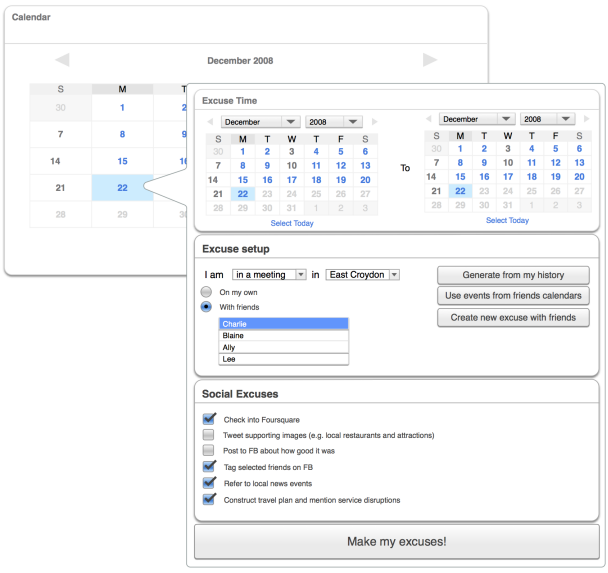

Figure 5: *lieCal*: Fictional interface for a tool which automatically generates excuses on behalf of the user, optionally including friends in the deception and strengthening alibis by posting on social media.

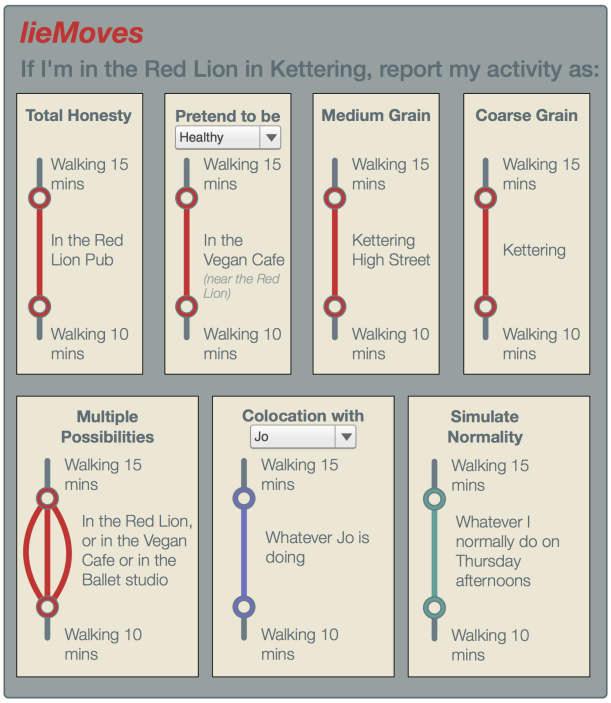

Figure 7: *lieMoves*: A fictional smartphone service for letting people obfuscate their location using various strategies, including blurring, substitution, past-replay and impersonation.

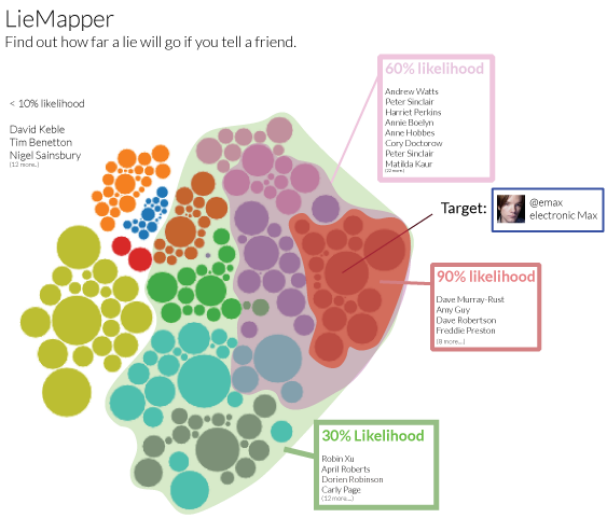

Figure 8: *lieMapper*: Fictional tool for predicting the flow of information (e.g. a lie) across a person's social network starting from a single friend.

Participant ID	Age/Gen	Employment	Privacy
p1	31-50 f	student, parttime	unc
p2	22-30 m	fulltime	prag
p3	31-50 f	freelance	prag
p4	18-21 m	student,parttime	prag
p5	22-30 m	fulltime	unc
p6	22-30 m	st, freelance	fund
p7	31-50 f	student	fund
p8	51-65 f	student	prag
p9	22-30 m	student	prag
p10	18-21 m	student, fulltime	prag
p11	22-30 f	student	unc
p12	22-30 f	student	prag
p13	22-30 m	student	prag
p14	22-30 m	fulltime	unc
p15	22-30 m	student	prag

Table 1: List of participants by ID with age range, gender (male, female, other), employment status (student, full time, part time, freelance, unemployed), and Westin Privacy Scale category (unc=unconcerned, prag=privacy pragmatist, fund=privacy fundamentalist).

5. RESULTS

Assuming they reported truthfully, the 15 participants we selected covered most of the major attributes in our demographic categories (see Figure 3). One notable exception is that all participants identified as either male or female, and almost half of the participants were males aged 22–30. We did not collect information on race, sexuality or any other attributes which might be used to identify marginalised groups.

11 participants self-reported using social networks several times a day, and all but one believed that half or less of their real world activity was represented on social media. 11 agreed or strongly agreed that they saw themselves as honest, but only seven agreed or strongly agreed to seeing themselves as honest online. Nearly half agreed that they thought their friends were honest.

Pertaining to attitudes towards privacy, 13 reported being at least *somewhat* concerned about their privacy online. Based on our categorisation of participants according to responses to the opening question, slightly over half fell into the Westin category of *privacy pragmatists*, while two fell into the category of *privacy fundamentalists*, and the remaining four were *unconcerned* about privacy. (High inter-rater agreement was achieved for this category; Fleiss's $k = 0.624$ for 3 raters and $N = 15$ participants.) These results show that in comparison to Westin's large survey of the American public [36], which had a respective breakdown of 55%-25%-20%, we had relatively few privacy fundamentalists among our participants, and slightly more of those in the unconcerned category. However, a meta-survey of privacy indices show that our proportion is comparable to more recent results [37].

In the following sections, we first present detailed case studies of three participants (P8, P9, and P13) to illustrate how individuals' attitudes towards privacy influenced their answers to some of the vignettes. We follow these descriptions with a presentation of themes derived from all participants.

5.1 Case study: Privacy and people (P8)

P8 is a former gradeschool teacher who has returned to university to get her Ph.D. She started using social media ten years ago when she was still working at the school, and her role as a teacher strongly shaped how she managed her exposure online. Specifically, her role influenced her caution in disclosing too much personally identifying information, but acknowledged that disclosure itself was important for fostering relationships and participation online.

> "When I was a teacher, I was very careful about what I said about teaching in school because at that point I'm not just 'me', personally; I'm also 'me' as a teacher, representing that school I was working at. Since I've stopped being a teacher, I unlocked my twitter feed, but still try not to post too much personal stuff online. But really, if you don't share some personal information then you miss out on so much interaction stuff, so it's a real balancing act."

When discussing *lieTinerary*, she described discovering that her ex-partner was fabricating extravagant holidays after their breakup in order to make her jealous.

> "[H]e wants me to think, "Oh, I should have stuck with him - he's having a really good life!". So there were pictures he was putting up [on Twitter] which were supposedly where he was on holiday, but of course once you know how to scrape people's Twitter data, you could see all of his posts were made in the UK. And at that point it became really obvious that that's what he was doing, so that made me smile."

She described wanting greater controls to be able to block said partner from getting around creating new profiles to look at her information:

> "I do know that, if he really wanted to he could easily set up another account. So in the end, although he's blocked [on Twitter] I don't assume he can't see what I'm saying; I assume that he can, and that's another reason that I'm a bit careful with what I say. So I wish it was easier, to stop people from being able to see what you're doing – how that would happen I don't know – but that would be really helpful."

5.2 Case study: Honesty and self-image (P9)

P9 is a 22-year-old recent graduate who is concerned about his privacy and the security of data he gives out online, in particular due to a distrust of the guarantees given by companies and platform providers.

> "I'm quite an honest person, [...] like if I was on a forum and I was talking to someone I'd tell the truth. But if a company were to ask me for my number or my name – I won't bother. "

Preferring honest approaches, P9 described that he would feel bad using tools that deceived other people directly, but would feel especially reluctant to use tools that left digital evidence of such acts, which would potentially serve as a later reminder:

> "Well, um, I imagine [lieCal] would be useful because it would give me an excuse if I wanted to do something, but I would probably feel worse [. . .] because it would serve as a reminder that I lied – like I'd go back into my calendar and go, shit, that that day I'd lied, and that day I'd lied as well – whereas I can just repress it otherwise. "

He was confident, however, that there were many people online that he knew of who would consider using tools like *lieTinerary* to promote themselves online, including pretending to go to exclusive events:

> "Well they might use [lieTinerary] to come across as fashionable or trendy — they might put up a post like "oh yeah I'm at London Fashion Week" when they're not really [...] I could say I'm at Glastonbury for the weekend, and immediately my cool points would go up. "

P9 believed that to a certain extent such fabrication was widespread already, and small acts of playing one's self up were already widespread.

> "I know people who have paid for likes and followers and stuff and they hashtag everything to death because they're so desperate for attention [...] I used to go on YouTube, Facebook I used to watch all these cool people build up nice communities [...], whereas there are lots of people nowadays who just want quick success and they'll take all of these cheap, cheating routes. "

5.3 Case study: Privacy and technology (P13)

P13 is a postgraduate student in his mid-twenties; technologically savvy and uses social networking sites every day. He is acutely aware of the volumes of data being collected through his web use, but finds himself weighing up the practicalities of taking steps to preserve his privacy with his immediate communication needs, often concluding that "life's too short" to act on his discomfort around third-party software.

> "I say what I'm doing on my Facebook because otherwise no-one will ever talk to me [...] I try and use small bits of privacy enhancing stuff, to whatever extent they actually work [...] So in the past I've had Facebooks where they're not tied to my... my lying even extended to that and all the information on them was fake. Nowadays I tend not to do that because the net effect of that is no-one talks to you. "

He takes steps to manage who sees his data on social media, by segregating his audience by platform, choosing who to share which aspects of his self with, and using privacy settings built into the platforms themselves. Sometimes this leads him to obtain information by proxy:

> "I don't connect to my mum's stuff and I don't want to connect to her stuff [...] but I wanted to find something out and so I remember asking my sister to look it up for me. "

He is also resigned to data leakage, and being surveilled, by both the government and advertisers.

> "I don't think I'm under any illusions about web stuff. If it's out there, it's out there. If someone wants to find it and knows the information or ways to get the information then they can get it. It's annoying, but it's a fact of life. "

This does not stop P13 from providing false information to services whenever he has the opportunity, under the impression that the data many services ask for is superfluous. He speculated that tools could be useful to generate more believable false data on his behalf.

> "So for instance airport wifi. I spend a large amount of my time in airports. So I think I'm listed as John Smith or... [...] So mostly it's whenever these anonymous websites want some personal information that they don't tend to have, then I tend to lie [...] But I always sort of wonder, should I be able to generate this? "

In general, he was concerned about the social risks of using tools to aid online deception,"especially when you can do this social ways, just going, oh I forgot to use the Google calendar again," but was also skeptical about how much he could trust the tools themselves.

> "If [social steganography] was something that I could run on my computer and I'd have it disconnected from the network then maybe. "

Despite his concerns, P13 expected that he would follow the status-quo if many people began using these tools, and expressly supported other people's right to use them, reasoning that the more people did so, the more effective they would become. However, he also anticipated that the output of the tools may be prone to detection and thus rendered useless.

> "You could imagine someone attacking these kinds of things and trying to start to write distinguishers for when is this posted by a human or is this posted by a social media bot. "

5.4 Effort and Complexity

A common reason why participants wouldn't use these tools related to the amount of effort required to use them. P8 observed that the effort-of-use barrier is a challenge even for tools already available today, and how platforms were potentially exploiting the lack

of adoption of these tools to their advantage:

> "The thing I've noticed is that people will always do the easiest [thing]. That's why nobody encrypts. I don't. You know, for all my concerns about privacy, I don't encrypt anything, [...] very few people take the extra security steps they can because it's convoluted because you have to make the effort to do something different. And the minute you ask people to do that, they'll just take the easiest route. And providers like Facebook and Twitter and Whatsapp and all the apps out there know that, and that's why it's so easy for them to collect data - they know people will just take the easiest route." *(p8)*

However, for some vignettes the extra effort was seen to pay off as an opportunity. For instance, in response to *Social Stenography*, P6 contemplated that by broadcasting different status updates to distinct subsets of his friends on Facebook, he could control multiple identities simultaneously:

> "I think essentially at this point you are projecting two identities simultaneously and you really would want to manage both. [...] it almost becomes twice the task. But the really interesting thing would be if different groups all had different keys - so you'd send a single status but they'd all see different ones. That would be sort of be neat, [to be] projecting multiple identities at once, because you can't really do that offline. Finally, technology would give us a chance to BETTER control our identities! " *(p6)*

A second aspect that was mentioned was not the direct effort of use, but that indirectly required to stay on top of the wake of deception left by using such tools. In some settings, participants noted specific compensatory measures that would be required to prevent being found out, and noted the complexity and effort of these measures.

> "If I used a tool like this and said I had been in meetings but then actually NOT logged the hours against the project, what the meeting was about or anything like that, it would make my accounting for my own time very hard. [...] It wouldn't balance, it wouldn't add up! So, in a corporate environment I actually think that's more problematic because I might end up losing my job over it." *(p7)*

5.5 Availability of Alternate Strategies

The most common reason given for not needing to use a tool was the availability of an alternative approach in the situation(s) in which the fictional tools were imagined to be most useful. The most common such strategy was simply *omitting* or *suppressing* information they did not wish to share; this strategy was used for the *Social Steganography* scenario, and *lieMoves*. The second most common strategy was the use of *alternate channels* and *access control features*. For instance, participants mentioned Facebook and Google+'s features for limiting the scope of a particular message as an alternative to using a steganography approach, as well as direct messages to individuals. P13 discussed the use of encrypted channels to both help control scope of a message and control for unwanted leakage by platforms.

In some cases, participants identified that alternate strategies were imperfect, and sometimes the fictional tool offered a better solution. For example, the alternate strategy of suppressing location leakage by turning location tracking off, was perceived as worse than *lieMoves* by both P6 and P9, because in suppression meant that apps that needed location (such as mapping applications, train schedule applications and location-based chat services) would not work.

There were fewer alternative strategies given for the other vignettes; "simply being honest", and in particular "blocking off time" was given as a common strategy for situations where *lieCal* would be useful (P4, P8, P9).

5.6 Privacy and Control

Several participants cited potential benefits to privacy control and management. The leaking of location information was a concern; six participants reported keeping location services on their smartphones turned off by default for reasons such as to prevent apps from sending their location to third-parties without their consent.

> "[lieMoves] would mostly catch out apps that were taking my location without even asking, because if I want to tell the truth when I think it matters, I can still do that, but those that are just spying on me gets crap! And that appeals, because they shouldn't be able to collect in the first place! " *(p6)*

P8 asked whether *lieMoves* was available for use, because she wanted it immediately to keep Google from tracking her.

> " I want to install it immediately and keep using it for the rest of my life! I wouldn't have any ethical worries about it because I wouldn't be lying to anyone, I would be lying to Google, and that's exactly what I want to do! Because they shouldn't have this information in the first place, so giving them wrong information is perfect. As I said, can I have this today, please? " *(p8)*

She added that some of the tools might bring to users a heightened awareness of how their information was being shared. However, respondents who knowingly shared a lot didn't see a reason to be concerned about the spread of their data; for P14 this was due to his social standing.

> "If we lived in a more totalitarian police state and I were genuinely afraid then I would understand, but then again I am a straight white male so... I don't really have much to fear. " *(p14)*

Others felt they had no choice but to share data, or that people didn't understand well enough how the services they used operated.

> " People can't make value judgements about the systems they interact with because they don't understand them well enough yet, especially what's going on behind the scenes. They don't actually feel the need to deceive system and platforms because they don't even know they're being spied upon. " *(p6)*

5.7 Authenticity and Self-Image

Participants reflected on how the data they shared affected other people's perceptions of them, as well as their perceptions of others on social media. P11 (in agreement with P1, P2, P3, P6, P7, P8, P9, P12 and P15) assumed that her friends engaged in "image-shaping" by "being quite selective or trying to present a particular kind of persona", and described an occasion when a contact's online presentation was at odds with what she knew to be happening

offline.

> " People will always seem like they're having a really good time and post about how great everything is but then you talk to them and things aren't actually quite how they're made to be portrayed on social media. [...] So like one of my friends, her sister was just posting about her one year anniversary of getting married, and how brilliant it was, and they were both posting about the presents they got for each other. Within a month they were separated [...] I know more about that from talking to my friend personally, but in terms of what's presented online to a different audience, to a much wider audience, that was not what was going on. " *(p11)*

P12 described a friend who, unable to withhold information or resist questions from an inquisitive audience, made up stories about her life to satisfy them, thus creating a persona.

> " 'Cos of the following that some fanfiction gets, she gets asked a lot of personal questions and she doesn't want to feel rude so she just lies, so she answers these very personal questions so she feels connected to her audience but she deliberately lies 'cos she finds it sometimes a bit invasive. " *(p12)*

P8 and P15 similarly mentioned deception used to protect privacy without alienating people. In contrast, others saw total openness in their sharing as important for presenting their "authentic" selves on social media, and thought less of those who they perceived to be engaged in deliberate image-shaping.

> " I wouldn't be friends with people who would be lying all the time or who make up stuff just for attention. [...] I know people who have paid for likes and followers, but if I found out that there was someone I was quite interested in doing this [...], the faith I put in them or the fact that I was being very genuine would take a bit of a hit. " *(p9)*

5.8 Polite Social Signalling and Autonomy

Though sometimes in conflict with attempts at authenticity, a number of respondents echoed the sentiment that degrees of deception are crucial for maintaining a well-functioning society.

> " I think that not telling people – everyone, everything – is a central aspect of being kind in the world. " *(p15)*

> " It's about empowerment – little lies, like "I'm just too tired and you're quite a taxing person" could be the truth but that's a bit mean, and you didn't want to say that! versus "oh no sorry I have plans with my boyfriend" which might be a lie, but it's nice. " *(p6)*

> "Often you lie to save people's feelings or – to stop someone finding out about a surprise party. Like there are really nice reasons to lie, and if you could help people make nice lies safer, that would be awesome! " *(p14)*

P6 commented that this could be a subtle method of signalling violations of personal privacy online:

> "The idea of being able to put massively sarcastic calendar appointments just so that, when someone looks at my calendar to see what I'm doing, they know I don't want them to know, and they should just stop asking. " *(p6)*

Such methods were also viewed as a form of civil empowerment; a way of giving people freedom to block off time (*lieCalendar*) or send a message (*Social Steganography*) in situations where the honest approach would be awkward due to shyness, introversion, or differences in social positions, e.g. having to contradict a superior in a social or professional environment.

> "Somebody younger, less experienced, less confident might find that this is a nice, straightforward way of blocking time out for themselves and feeling good or comfortable about it. Because it can be quite difficult saying "no, I'm not free" to someone senior. " *(p8)*

5.9 Ethics and Morality

Finally, many of the participants volunteered their views on ethical or moral reasons of why they would or would not use these tools in specific ways. Perspectives varied in general and according to the vignette presented.

The technology vignettes could be seen as ethically neutral, with the ethics coming from the manner of their use:

> " If your intention is to use these tools to harm someone, then that's the individual's own decision to make and you can decide for yourself whether that's morally right or wrong. But simply using the tools themselves doesn't imply you're going to do something that is harmful or morally wrong. " *(p5)*

> " Ultimately, this is just like any other cracking software: you leave it to the user to decide what to do with it. You're not responsible for their moral actions, or at least that's what the developers say. " *(p4)*

However, in some cases, there was such a strong implication between the design of the tool and the kinds of lies which it facilitates that the morality of the tool became the morality of the action:

> "Well as someone who's considered murdering people before, this is exactly how I would do it. I would create a fake social media presence so I could go off and do something illegal or even . . . I could commit adultery, I really can't see much of a practical application for ethically good things. . . " *(p14, discussing lieTinerary)*

To P6, whether deception was moral contextually dependant on whether the recipient had a legitimate need for the truth and why.

> " If someone has a right to know something for some reason [...] then lying to them there is more problematic than if they didn't have a right to ask you, or to be looking for that information. [...] that's their own fault; they should have know they shouldn't have looked. " *(p6)*

Some participants suggested that they would need a really good reason to use deception tools. P14 felt that a better alternative to having to lie was to get out of situations in which one felt the need to lie.

> " And if you're in a situation where you have to lie to people about where you are, then that's a situation you need to get out of cos that's a creepy situation [...] The only time I can see this being good is like if you're in an abusive marriage and you're going to a divorce

lawyer in secret. " *(p14)*

There was often a moral distinction made between friends and platforms as the targets of deception. While a majority (11) reported taking issue with deliberately deceiving friends, there was also widespread consensus on wanting not to deceive a general audience on social media. A notable exception to this was a feeling that lying to platforms is not dishonest.

> "well if I'm talking to my friend I always tell the truth; I think I'm quite an honest person, and I don't really discuss anything with any other people in that sort of way which I'd lie about,. . . I don't think lying to Facebook is unethical [. . .], because it's not affecting any of your friends or anyone on your list, so it has no effect – so you're not really lying to anyone? [...] I don't trust these companies enough, to be honest, with the information I supply them. " *(p9)*

P6 took the position that lying to platforms should be the moral choice, even part of one's civic duty.

> " I think lying to Facebook is to be encouraged! [platforms] spend so much effort in deceiving users into thinking they're doing one thing when they're doing another, that giving users some control seems fine. Its sort of like the debate whether minorities can be racist against white people – like, whether the power imbalance seems to negate any meaningful argument, certainly when it comes to lying to services. " *(p6)*

6. DISCUSSION

6.1 Morality of Deception

Our participants, like the majority of people, like to think of themselves as being generally honest, but this has a nuanced relationship with their stated behaviour. There was a common feeling that deceiving platforms and corporations was acceptable, or even a moral imperative. Nomenclature was significant: casting activities as 'lying' provoked responses which paid more attention to the ramifications of being found out, and a greater sense of ethical violation. Hiding information was generally seen as acceptable, as was partitioning information for different audiences, especially in the context of avoiding unwanted attention. Politeness was often cited as a valid reason for performing white lies, a variety of kindness.

Akerlof and Schiller's account [2] focuses on deception from the point of view of corporations, and therefore helps explain the existence of situations in which our survey subjects were motivated to deceive. In the information economy, data subjects are beguiled, misled or strongarmed into giving away more data than is required for the service they wish to access. However, perhaps because their focus is wider than the information economy, Akerlof and Schiller fail to consider the possibility of the individual creating counter-asymmetries by manipulating the data they provide to corporations. Their recommended counter-measures are all intended to support truthfulness - standards-setting, reputation, regulation. Yet these all require concerted action; deception is a strategy open to the individual.

6.2 Promoting Social Honesty

One viewpoint is that mendacious impulses are indicative of a problematic situation: that fixing the socio-technical context would remove the need to deceive, and the community could become more socially honest. Systems requesting excessive information frequently provoked anger, and a feeling that feeding back fictitious information was justified. One lens for designers to engage with this issue is Grice's conversational maxims. Typically, these are used to define one side of a social contract: the quantity, quality, relation and manner of information production. A complimentary view applies them to requests for information. This accounts for many of the indignant quotes we recieved—systems were asking for *too much* information, or *irrelevant* information. Providing clarity here, relating information demands to the current context, limiting information to the that which is necessary can guide designers towards upholding the platform's end of the social contract. Our *lieMapper* vignette asked how far through our social networks personal information was likely to diffuse, alerting the user to social information violations; similarly, when designers illuminate the hidden pathways which our data takes—or doesn't—it provides a grounding on which trust can be built.

Legal identities, and the problems which they cause, highlight the multifacted aspects of life, whether online or off. The general trend is towards a collapse of context, the joining of identities across sites and networks, but the attitude that people should be happy to connect all of their identities together in this way is an expression of social privilege. Tools exist to aid the management of multiple personas, typically used by astroturfing organisations [35, 30]. As a provocation, what would design for multifaceted life look like? Are there ways to support participants in plural presentation, helping them to understand and maintain their context bounds, rather than attempting to force a homogenisation. How can we support radical self expression and support marginalised groups? What about systems which acknowledge that there are parts of users lives which they don't want to share publicly, but they still need to express them and connect with similar people? Designing for contextual authenticity rather than imposing singular identity pushes back against marginalisation.

6.3 Memory, Safety, and Plausible Deniability

It was clear from responses that being reminded of one's lies can be upsetting, especially for people who consider themselves honest. On one hand, this might suggest systems might automatically remove, or reduce the visibility of, digital traces that could serve as reminders of one's past deceit. The recent growth in messaging apps that automatically delete messages after a single viewing [23] might, in fact, be related to this perceived design need. On the other hand, visibility of such actions can lead people towards greater honesty—knowing how often one was deceptive could clearly be a powerful push towards veracity.

A second major theme pertained to effort, both of use and potentially of maintenance, post-deceit. It was clear that any tool that required more time and effort than what they were used to was preceived as not very useful. But having to explicitly act at all was also viewed negatively; that is, having to engage with a tool in order to carry out a deception, such as with *lieCal*, was viewed less favourably than something that could do it automatically, such as *lieMoves*.

An additional problem with tools that require explicit action is that they leave little space for plausible deniability. Since explicit action is needed, it becomes often difficult to justify that such an action was taken accidentally or unintentionally (assuming the individual was of sound mind). If we had instead imagined tools that *deceived by default*, the possibility that deception was unintended, but that the individual was busy or simply forgot to make the system tell the truth would still be plausible. For example, a deceive by default variation of *lieCal* might automatically fill the person's cal-

endar for the following week or month, allowing its user to quickly identify and replace these false appointments with real ones they want others to know about. Such designs would additionally support many of the goals of *privacy-by-design* [48].

Another significant barrier to the use of such tools pertained to safety and discoverability, the first pertaining to ensuring that deceptive actions would not have unintended consequences, while the second pertained to effort and actions that would be necessary to ensure deceptive actions would not be discovered. Such concerns suggest that there is a potential space for future tools that are able to support *safe deception*, both in terms of highlighting potential hazards and towards mitigating the burden of managing active lies or their effects. Tools such as *lieMapper* that provided situational awareness about social information flow could help individuals tell certain, especially *nice lies* (as described by P8), safely.

7. CONCLUSION

Deception is a long-established strategy for informational self-determination, and it is not a surprise to see the practice in online behaviour. The study reported here is a necessary prolegomenon to the deep study of deception, and establishes interesting lines of enquiry which mark out a descriptive vocabulary, a potential design space, and even the beginnings of a sketch of a bottom up morality in this area.

Nissenbaum outlined the importance of contextual integrity for online design, the idea that individuals bring a set of expectations and meanings to their online interactions that are often derived from offline analogues, appropriately or otherwise. A designed interaction that leaves no space for someone to present themselves creatively for non-malevolent purposes fails to preserve contextual integrity, and would consequently produce an asymmetry of understanding between user and system of which the user may be unaware.

Particular strategies and opportunities for deception were common to many of our subjects, who were often concerned with the balance between the moral injunction against lying, and their own interests. Mitigating factors were sought - for example, if the interlocutor in the interaction is non-human (a platform, for instance), or if the interaction provided an opportunity for malign action (e.g. could be used by a stalker), or if the interlocutor did not have a good reason for wanting the data, then these were seen as justifications for using deception for protection. Morally, there are of course issues with this – in particular, whether deceivers are free-riding on the efforts of a truthful majority. Deception is a successful strategy for self-protection, but presumably the deceiver also wants the benefits of the interaction, which may not be forthcoming if interactions with other agents also produced false data. However, the moral calculation that our subjects were compared their own interests and the legitimacy of the interests of whoever demanded the data.

This suggests design principles which could be tested in future work. Those providing services for data need to identify, respect and avoid the factors which lead users to deception. The act of deception creates a situation in which data minimisation is in the interests of the platform – the less that it asks for, the more likely it is to be trusted, and the less likely the deception strategy is to be invoked. In particular, contextual integrity is preserved if users are able to represent themselves differently in different contexts, and it is clear to them that the more data that is demanded, the easier it is to resolve these personas. Similarly, there is a set of deceptions, such as butler lies, which are adapted to specific communication situations, and facilitating these will also help transfer and preserve expectations in the digital context. Systems which facilitate deception will have both positive and negative potential. Most obviously, their wide uptake would reduce trust in data generally. On the other hand, it is clear from our survey that for most people, deception is a last resort, and the majority self-image is one of general honesty so that deception would demand ad hoc justification. The rather more calculated invocation of a deception system might, if such attitudes were widespread, be a step too far. Framing the objective of the system will be key – for example, classifying such systems as privacy-enhancing, rather than deceiving, might increase their acceptance. On the other hand, software that maintained a consistent, false, record of events might remove the burden from users of understanding that their behaviour is deceptive, and make it easier to deceive. Such divergent potential outcomes require investigation.

8. ACKNOWLEDGEMENTS

This work is supported under *SOCIAM: The Theory and Practice of Social Machines*. The SOCIAM Project is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/2 and comprises the Universities of Oxford, Southampton, and Edinburgh.

9. REFERENCES

[1] Eytan Adar, Desney S. Tan, and Jaime Teevan. 2013. Benevolent Deception in Human Computer Interaction. *Chi 2013* (2013), 1863–1872. DOI: http://dx.doi.org/10.1145/2470654.2466246

[2] George A. Akerlof and Robert J. Shiller. 2015. *Privacy in context: Technology, policy, and the integrity of social life*. Princeton University Press.

[3] Harm Akker, Valerie M. Jones, and Hermie J. Hermens. 2014. Tailoring Real-time Physical Activity Coaching Systems: A Literature Survey and Model. *User Modeling and User-Adapted Interaction* 24, 5 (Dec. 2014), 351–392. DOI:http://dx.doi.org/10.1007/s11257-014-9146-y

[4] A Albrechtslund. 2008. Online social networking as participatory surveillance. *First Monday* 13, 3 (2008), 1–11. http://firstmonday.org/ojs/index.php/fm/article/viewArticle/2142

[5] James M Alexander and Jonathan M Smith. 2010. *Disinformation : A Taxonomy*. Technical Report. University of Pennsylvania Department of Computer & Information Science.

[6] Luigi Anolli, Michela Balconi, and Rita Ciceri. 2001. Deceptive Miscommunication Theory (DeMiT): A New Model for the Analysis of Deceptive Communication. *Say not to say New perspectives on miscommunication* FEBRUARY (2001), 73–100. http://books.google.com/books?hl=en&lr=&id=PsiLjRHr1JQC&oi=fnd&pg=PA73&dq=Deceptive+Miscommunication+Theory+(+DeMiT+):+A+new+model+for+the+analysis+of+deceptive+communication&ots=GkJjOpQJfZ&sig=DfUVh93xYcd2WThfB5_w3vU7nfo

[7] Paul M. Aoki and Allison Woodruff. 2005. Making Space for Stories: Ambiguity in the Design of Personal Communication Systems. (2005), 10. DOI: http://dx.doi.org/10.1145/1054972.1054998

[8] James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1 (2013), 11–35.

[9] Jeffrey Bardzell, Shaowen Bardzell, and Erik Stolterman. 2014. Reading critical designs: supporting reasoned interpretations of critical design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1951–1960.

[10] Jeremy Birnholtz, Graham Dixon, and Jeffrey Hancock. 2012. Distance, ambiguity and appropriation: Structures affording impression management in a collocated organization. *Computers in Human Behavior* 28, 3 (2012), 1028–1035.

[11] Sissela Bok. 1978. Lying: Moral choice in public life. *New York: Pantheon* (1978).

[12] danah boyd. 2002. *Faceted id/entity: Managing representation in a digital world*. Ph.D. Dissertation. Massachusetts Institute of Technology.

[13] danah boyd. 2014. *It's Complicated: the Social Lives of Networked Teens*. Yale University Press.

[14] David B Buller and Judee K Burgoon. 1996. Interpersonal deception theory. *Communication theory* 6, 3 (1996), 203–242.

[15] Judee K Burgoon, Roxanne Parrott, Beth A Le Poire, Douglas L Kelley, Joseph B Walther, and Denise Perry. 1989. Maintaining and restoring privacy through communication in different types of relationships. *Journal of Social and Personal Relationships* 6, 2 (1989), 131–158.

[16] Carl Camden, Michael T. Motley, and Ann Wilson. 1984. White lies in interpersonal communication: A taxonomy and preliminary investigation of social motivations. *Western Journal of Speech Communication* 48, 4 (1984), 309–325. DOI:http://dx.doi.org/10.1080/10570318409374167

[17] Elizabeth Anne Casale. 2015. *Cryptocurrencies and the Anonymous Nature of Transactions on the Internet*. Ph.D. Dissertation.

[18] Dwight Conquergood. 1991. Rethinking ethnography: Towards a critical cultural politics. *Communications monographs* 58, 2 (1991), 179–194.

[19] Sunny Consolvo, David W. McDonald, Tammy Toscos, Mike Y. Chen, Jon Froehlich, Beverly Harrison, Predrag Klasnja, Anthony LaMarca, Louis LeGrand, Ryan Libby, Ian Smith, and James A. Landay. 2008. Activity Sensing in the Wild: A Field Trial of Ubifit Garden. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1797–1806. DOI:http://dx.doi.org/10.1145/1357054.1357335

[20] Judith Wagner DeCew. 1997. *In pursuit of privacy: Law, ethics, and the rise of technology*. Cornell University Press.

[21] Bella M DePaulo, Deborah A Kashy, Susan E Kirkendol, Melissa M Wyer, and Jennifer A Epstein. 1996. Lying in everyday life. *Journal of personality and social psychology* 70, 5 (1996), 979.

[22] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. *Tor: The second-generation onion router*. Technical Report. DTIC Document.

[23] Maeve Duggan. 2015. Mobile Messaging and Social Media 2015. *Pew Research Center* (2015). http://www.pewinternet.org/2015/08/19/mobile-messaging-and-social-media-2015/

[24] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT Press.

[25] Peter Eckersley. 2010. How unique is your web browser?. In *Privacy Enhancing Technologies*. Springer, 1–18.

[26] Mauricio S Featherman, Anthony D Miyazaki, and David E Sprott. 2010. Reducing online privacy risk to facilitate e-service adoption: the influence of perceived ease of use and corporate credibility. *Journal of Services Marketing* 24, 3 (2010), 219–229.

[27] Elizabeth Fife and Juan Orjuela. 2012. The privacy calculus: Mobile apps and user perceptions of privacy and security. *International Journal of Engineering Business Management* 5, 6 (2012), 7.

[28] Bill Gaver and Heather Martin. 2000. Alternatives: Exploring Information Appliances Through Conceptual Design Proposals. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. ACM, New York, NY, USA, 209–216. DOI: http://dx.doi.org/10.1145/332040.332433

[29] William Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a resource for design. 5 (2003), 233–240. DOI: http://dx.doi.org/10.1145/642611.642653

[30] Nicholas Gilewicz and François Allard-Huver. 2012. Digital Parrhesia as a Counterweight to Astroturfing. *Online Credibility and Digital Ethos: Evaluating Computer-Mediated Communication: Evaluating Computer-Mediated Communication* (2012), 215.

[31] Herbert P Grice. 1970. *Logic and conversation*. na.

[32] Ceki Gülcü and Gene Tsudik. 1996. Mixing E-mail with Babel. In *Network and Distributed System Security, 1996., Proceedings of the Symposium on*. IEEE, 2–16.

[33] Jeff Hancock, Jeremy Birnholtz, Natalya Bazarova, Jamie Guillory, Josh Perlin, and Barrett Amos. 2009. Butler lies: awareness, deception and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 517–526.

[34] Max Van Kleek, Dave Murray-Rust, Amy Guy, Daniel A Smith, and Nigel R Shadbolt. 2015. Self Curation , Social Partitioning , Escaping from Prejudice and Harassment : the Many Dimensions of Lying Online. In *WebSci2015*.

[35] Eugenia Kolivos and Anna Kuperman. 2012. Consumer law: Web of lies-legal implications of astroturfing. *Keeping Good Companies* 64, 1 (2012), 38.

[36] David Krane, Laura Light, and Diana Gravitch. 2002. Privacy on and off the Internet: What consumers want. *Harris Interactive* (2002).

[37] Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. Privacy indexes: a survey of Westin's studies. (2005).

[38] Svenn Lindskold and Pamela S Walters. 1983. Categories for acceptability of lies. *The Journal of Social Psychology* 120, 1 (1983), 129–136.

[39] Shah Mahmood and Yvo Desmedt. 2012. Your Facebook deactivated friend or a cloaked spy. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2012 IEEE International Conference on*. IEEE, 367–373.

[40] a. E. Marwick and d. boyd. 2010. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 1 (July 2010), 114–133. DOI: http://dx.doi.org/10.1177/1461444810365313

[41] Steven a. McCornack. 1992. Information manipulation theory. *Communication Monographs* 59, 1 (1992), 1–16. DOI:http://dx.doi.org/10.1080/03637759209376245

[42] Miriam J Metzger. 2004. Privacy, trust, and disclosure: Exploring barriers to electronic commerce. *Journal of Computer-Mediated Communication* 9, 4 (2004), 00–00.

[43] Charles E Naquin, Terri R Kurtzberg, and Liuba Y Belkin. 2010. The finer points of lying online: e-mail versus pen and paper. *The Journal of applied psychology* 95, 2 (2010), 387–394. DOI:http://dx.doi.org/10.1037/a0018627

[44] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington law review* 79, 1 (2004).

[45] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

[46] Kate Raynes-Goldie. 2010. Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook. *First Monday* 15, 1 (2010).

[47] Vincent Scarlata, Brian Neil Levine, and Clay Shields. 2001. Responder anonymity and anonymous peer-to-peer file sharing. In *Network Protocols, 2001. Ninth International Conference on*. IEEE, 272–280.

[48] Peter Schaar. 2010. Privacy by design. *Identity in the Information Society* 3, 2 (2010), 267–274.

[49] Omer Tene and Jules Polenetsky. 2012. To Track or Do Not Track: Advancing Transparency and Individual Control in Online Behavioral Advertising. *Minn. JL Sci. & Tech.* 13 (2012), 281.

[50] Ronny E Turner, Charles Edgley, and Glen Olmstead. 1975. Information control in conversations: Honesty is not always the best policy. *Kansas Journal of Sociology* (1975), 69–89.

[51] Wikipedia. 2015. Disinformation — Wikipedia, The Free Encyclopedia. (2015). https://en.wikipedia.org/wiki/Disinformation [Online; accessed 15-Sept-2015].