

Dark Patterns in Proxemic Interactions: A Critical Perspective

Saul Greenberg¹, Sebastian Boring², Jo Vermeulen³, Jakub Dostal⁴

¹Interactions Lab, University of Calgary, Calgary, Canada

²HCC Group, University of Copenhagen, Copenhagen, Denmark

³Hasselt University – tUL – iMinds, Expertise Centre for Digital Media, Diepenbeek, Belgium

⁴School of Computer Science, University of St Andrews, Scotland, UK

¹saul@ucalgary.ca, ²sebastian.boring@diku.dk, ³jo.vermeulen@uhasselt.be, ⁴jd67@st-andrews.ac.uk

ABSTRACT

Proxemics theory explains peoples' use of interpersonal distances to mediate their social interactions with others. Within Ubicomp, *proxemic interaction* researchers argue that people have a similar social understanding of their spatial relations with nearby digital devices, which can be exploited to better facilitate seamless and natural interactions. To do so, both people and devices are tracked to determine their spatial relationships. While interest in proxemic interactions has increased over the last few years, it also has a *dark side*: knowledge of proxemics may (and likely will) be easily exploited to the detriment of the user. In this paper, we offer a critical perspective on proxemic interactions in the form of *dark patterns*: ways proxemic interactions can be misused. We discuss a series of these patterns and describe how they apply to these types of interactions. In addition, we identify several root problems that underlie these patterns and discuss potential solutions that could lower their harmfulness.

Author Keywords

Dark patterns, anti-patterns, proxemic interactions.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION

Authors of human-computer interaction papers concerning innovative design ideas tend to forward their central idea in a positive – often highly idyllic – light. True critical perspectives are rarely offered. When they are, they tend towards a few cautionary lines in the discussion, or relegated to future work where its actual use would be examined. The problem is that many of our new innovations involve designing for ubiquitous computing situations that are extremely sensitive to intentional or unintentional abuse (e.g., privacy, distraction

and intrusion concerns). Rather than wait until some future field study of our technology (where it may be too late to address emerging concerns), we should consider the 'dark side' of our technologies at the outset.

The particular innovation we are concerned with is *proxemic interactions*, which was inspired by Hall's Proxemic theory [13]. The theory explains people's understanding and use of interpersonal distances to mediate their social interactions with others. In proxemic interactions, the intent is to design systems that will let people exploit a similar 'social' understanding of their proxemic relations with their nearby digital devices to facilitate more seamless and natural interactions [12]. This is especially important as we become immersed in ubiquitous computing ecologies, i.e., where we carry and are surrounded by myriads of devices, all potentially capable of interacting with one another. Examples include: mobile devices that understand their spatial relations to mediate information exchange between nearby devices [19, 22]; large displays that sense people's position relative to them, where they dynamically adjust what is shown and how people can interact with them [28, 17, 20]; public art installations that respond to the movement and proximity of people within its sphere to affect what is shown [26]; application areas such as home media players that monitor the distance and orientation of its viewers to dictate what is shown [2], and information visualizations that tune their visuals to people's position relative to them [16]. The literature also includes more general essays about the role of proxemics, such as how it can address well-known challenges in Ubicomp design [21].

Yet it is clear, at least intuitively, that there is a dark side to proxemic interactions. For example, the systems above rely on sensing people and their devices within the surrounding environment. Indeed, [12] describe several sensed dimensions that would be valuable to system design: *distance*, *orientation*, and *movement* of entities relative to one another, the *identity* of these entities, and *contextual information* about the location. While their purposes are honorable, such sensing – as well as the inevitable inaccuracy of interpreting and translating that information into action – immediately raises concerns by experts and non-experts alike about privacy, errors, distraction and intrusion. In addition, dystopian visions of the future hint at abuses of such technologies – a

well-known example is the movie *Minority Report* that illustrates how a character is selectively bombarded by targeted advertisements as he moves in a public space.

In this paper, we revisit the idea of proxemic interactions, where our goal (and contribution) is to present a critical perspective – the dark side – of this technology. Our method is to articulate potential *dark patterns* indicating how we think this technology can be – and likely will be – abused, and *anti-patterns* in which the resulting behavior occurs as an unintended negative side effect. To avoid being overly broad, we focus our scope somewhat to people’s proxemic interactions with large (and mostly public) displays, although we do illustrate other examples as needed.

DARK PATTERNS AND ANTI-PATTERNS

Architect Christopher Alexander introduced the notion of *design patterns*: a documented reusable and proven solution to an architectural design problem [1]. Design patterns are typically derived by examining existing solutions to design problems (which may include ‘folk’ solutions) and generalizing them. Design patterns were later advocated as a way of describing common solutions to typical software engineering problems [11], and to interaction design problems [5]. They are usually structured as a *name*, a *problem* that explains it, a *solution* that describes how the problem is solved, and *consequences* of applying the pattern [11].

A *dark pattern* is a special kind of pattern, defined as:

“a type of user interface that appears to have been carefully crafted to trick users into doing things [where these user interfaces] are carefully crafted with a solid understanding of human psychology, and they do not have the user’s interests in mind.” – Brignull et al. [7]

Brignull et al. created a web-based library of dark patterns concerning *intentionally* deceptive e-commerce practices [7,6]. Their specific goal was to recognize and name these practices so that people would be aware of dark patterns in an interface, and to shame the companies using them. For example, they describe a ‘hidden cost’ pattern that “occurs when a user gets to the last step of the checkout process, only to discover some unexpected charges have appeared”, illustrated by how several named companies use it.

Highly related to dark patterns are *anti-patterns* that indicate a design failure or non-solution [18], or an otherwise bad design choice. While dark patterns are intentional, anti-patterns are designs that *unintentionally* result in a negative experience or even harm [31].

In the remainder of this paper, we combine the notion of dark patterns and anti-patterns somewhat more broadly. We articulate not only possible deceptions and misuses of proxemic interactions (dark patterns), but also problems that may appear even when the designer has reasonable intentions (anti-patterns). Unlike true patterns that are based on analyzing a broad variety of existing solutions, we draw from a broad variety of domains to derive our patterns. We consider the dark

side of emerging uses of proxemics in commercial, experimental and research, fictional portrayals of such technologies foreshadowed by the popular literature and cinema, and our own reflections of where misuses could occur. We also revisited the dark pattern library [7] to see if and how they could be applied to proxemic interactions (possibly as variations). That is, our patterns are a mix of those that describe existing abuses and those that predict possible future ones. Our pattern articulation process is somewhat akin to affinity diagramming, where we clustered the examples we saw and named them as patterns. Unlike an empirical analysis (which is infeasible given the youth of this area), we use these examples to speculate on what could be.

We do not differentiate whether a particular pattern is dark vs. anti: as our examples suggest, the difference between the two often arises from the designer’s intent rather than a particular design feature. That is, the same pattern – depending on the designer’s intent – can be viewed as either a dark pattern or an anti-pattern.

While the novelty of proxemic interaction makes pattern elicitation somewhat of a thought exercise (albeit grounded in existing examples where possible), we believe this approach to be appropriate for forecasting – and ideally mitigating – the dark side of our future technologies *before* actual deceptive patterns become widespread. We now turn to our patterns. Afterwards, we discuss how many of our patterns arise from several foundational issues.

1. THE CAPTIVE AUDIENCE

The person enters a particular area to pursue an activity that takes a given time, and that does not involve the system. The system senses the person at that location, and begins an unsolicited (and potentially undesired) action based on the fact that the person is now captive.

Unlike desktop computers, technology can be spatially located in an environment to leverage a person’s expected routines. When done for beneficial purposes, the technology enhances or supports what the person normally does at that location – one of the basic premises of embodied interaction [9]. The captive audience pattern instead exploits a person’s expected patterns and routine for its own purposes, where the system knows that the person cannot leave without stopping what they otherwise intended to do.

Novo Ad (www.novoad.com), for example, produces advertising mirrors that display video ads on mirror-like screens ranging in size from 21–52”. The Novo Ad web site on its Advertising Mirror page states:

“the system serves as a mirror screen which identifies figures standing in front of it and switches itself automatically on. At start-up the screen displays a 6 second long ad on a full screen, which is later reduced to ¼ of the screen”.

Novo Ad identifies public washrooms as one of the prime locations for their displays, and even created a promotional video showcasing their technology in a woman’s washroom (YouTube ID: PXwbacfAwnY) as sketched in Figure 1. The



Figure 1. Novo Ad screenshot (sketch by David Ledo).

woman becomes the captive audience, as her primary task is to use the sink and mirror for grooming. The video ad, which starts on her approach, is the unsolicited system action. Other captive locations listed by Novo Ad include dressing rooms and elevators.

Captive Media, a British company, takes this one step further (www.captivemedia.co.uk). They estimate that a man using a urinal is captive for ~55 seconds. They place screens directly above the urinal (Figure 2, left), and use proximity and ‘stream’ sensors “to detect the position of a man’s stream as he pees” (Figure 2, right). This information is then used to activate advertising-sponsored pee-controlled games as illustrated in Figure 2. Later, we discuss how Captive Media has various positive policies in place to mitigate dark pattern effects, where it tries to balance the benefits to product user (by game play) and the client (by effective advertising).

Some ATMs employ the captive audience pattern in a particularly effective way, at least for the advertiser. By displaying advertisements when customers are waiting to receive cash or have their bankcard returned, they exploit that the captive customer cannot leave or divert their attention without risking loss of the desired transaction or even one’s bank card.

As another example, BBDO Düsseldorf and Sky GO developed a device that can transmit audio advertisements to commuters resting their head against the train window. Commuters suddenly hear an advertisement that is inaudible to other passengers and sounds like a voice in their head. The transmitter works by sending high-frequency vibrations through



Figure 2. Captive Media screenshot (YouTube ID: XLQoh8YCqo4#t=44), © Captive Media, with permission.



Figure 3. Scene from Black Mirror (sketch by David Ledo).

the window, which are then picked up in the commuter’s inner ear using bone conduction.

‘15 Million Merits’, the 2nd episode of the dystopian Black Mirror Channel 4 television series, also includes several examples of the captive audience pattern. It portrays a future where each person’s bedroom is built out of display walls that are always on when that person is present (Figure 3). They can only be turned off temporarily when a person makes a payment, or by leaving the room.

2. THE ATTENTION GRABBER

The person happens to pass by the field of view of a strategically located system. The system takes deliberate action to attract and keep that person’s attention.

Attracting attention of a passerby is an exceedingly common strategy used by anyone selling a product or service: the goal is to turn the passerby into a customer. Carnival barkers, greeters in establishment doorways, aggressive street peddlers – all verbally address a passerby to try to get them to enter into a conversation and ultimately into a sales transaction. Establishments use storefronts and windows to advertise their wares. Flashing lights and myriads of public signage and billboards (some electronic and digital) commonly compete for the passerby’s attention.

Proxemics-aware public devices are perfectly poised to grab attention of passersby. Like barkers and greeters, they can sense the passerby as an opportunity, as well as gauge how well their attention-getting strategies are working by how the person responds. For example, turning to face the device, or stopping, or approaching the display all suggest that a person’s attention is momentarily acquired.

The dystopian future depicted in the movie ‘Minority Report’ contains a scene that popularized this scenario. Multiple advertising walls detect the protagonist John Anderton moving through a crowded hallway. All walls vie for his attention in a visual and audio cacophony. The ad wall for Guinness Draught, for example, shouts his name along with a directed message: “John Anderton, you could use a Guinness right about now!”




Figure 4. The Nikon D700 Billboard (sketch by David Ledo).

An example of an existing simple but compelling public display in this genre is the *Cheil Worldwide Nikon D700 Guerilla-Style Billboard* (Figure 4). Located in a busy subway station in Korea, it displays life-size images of paparazzi that appear to be competing for the passerby's attention. When the passerby is detected in front of the billboard, lights flash (as in Figure 4) to simulate flashing cameras. The red carpet leads to a store that sells the type of cameras being used.

Within advertising and marketing, this pattern is commonly referred to as **AIDA**, an acronym for: attract Attention, maintain Interest, create Desire, and lead customers to Action [27]. Wang et al. [30] extended AIDA to proximity-sensing digital displays by their Peddler Framework, itself an extension of the Audience Funnel [23]. The framework covers six interaction phases a person may be in, all of which can be inferred by the proxemics measures of distance, motion, and orientation. Each phase indicates increasing (or decreasing) attention and motivation of the passerby [30].

- Passing by* relates to anyone who can see the display.
- Viewing & reacting* occurs once the person shows an observable reaction.
- Subtle interaction* happens if the person intentionally causes the display to react to his movement and gestures.
- Direct interaction* occurs when the person moves to the center of the display and engages with it in depth.



- Digressions and loss of interest* occur when a person either looks away from the display, or starts moving away from it.
- Multiple interactions* occur when the person re-engages with the display.
- Follow-up actions* happen after interactions with the display are completed.

Wang et al. then illustrate a proxemics-aware public advertising display for selling books [30]. It exploits the phases above to attract and retain the attention of a passerby. For example, the initial attention of a passerby is attracted by rapid animation of a pictorial product list; once the passerby looks at the display, the animation slows down to become readable (Figure 5, left). If the person approaches the display, various products are featured by growing in size. If the system detects him looking or moving away, it tries to regain the passerby's attention using subtle animation, e.g., by shaking particular displayed products (see 5, right), and by displaying other potentially interesting products.

Commercial interest in attention-grabbing systems are increasing. For example, Apple's iBeacon is an experiment that recognizes a person (via her iPhone) at specific locations in an Apple store. Notifications about a particular nearby product are then sent and displayed on that person's phone.

While the above examples illustrate how proxemic displays can grab attention in an entertaining and perhaps subtle manner, they can also be obnoxious. An earlier version of the Peddler system [30] displayed flashing graphics and shouted out loud audio messages to the passerby. The more the display was ignored, the more insistent it became. The Black Mirror episode mentioned previously includes an extreme example of a fascist Attention Grabber pattern within the context of a Captive Audience pattern: the display wall shown in Figure 3 detects when the person is trying to shut out the displayed information by sensing if that person's eyes are closed, or turned away. If so, it plays increasingly annoying sounds and messages to force the person to look at the content.



Figure 5. Proxemic Peddler. Left: Attention-attracting animation slows if passerby gazes at display. Right: Product graphic shakes to re-attract attention if person turns away [30]. © Wang, Boring & Greenberg, with permission.

3. BAIT AND SWITCH

The system baits the viewer with something that is (from the viewer's perspective) desirable, but the system then switches it to something else after the person directs his or her attention to it and moves closer.

Brignull et al. characterize this pattern as follows [7]:

“The user sets out to do one thing, but a different, undesirable thing happens instead. This is one of the oldest tricks in the book, and it is very broad in nature...”

Consider the case where a public display has gained a viewer's attention because the viewer is in fact interested in the 'bait' being displayed (e.g., an apparently incredible offer). The viewer 'opts-in' by approaching the display. In turn, the display recognizes the viewer's interest and offers further enticing details concerning its content. The viewer's attention becomes increasingly focused. Once the viewer is fully drawn in, the system then switches to something else. A typical 'switch' would be to an inferior or more costly product purportedly because the initially advertised product is no longer available. Another switch may require the viewer to sign up to some otherwise unwanted service before the viewer can proceed (which could also become a security issue). Yet another switch is the introduction of other content (i.e., unexpected advertising) in this process.

A compelling (and in this case useful) bait-and-switch example was developed by Amnesty International, where they created a bus-stop display that detects when people are looking at it (menstrualpoetry.com/postimages/amnestyinternational_happensnobodywatchingbig.jpg; also YouTube id: DQI_pnuNskQ). When no-one's gaze is directed at it, it displays a scene showing domestic violence, which is viewable out of the corner of one's eye (Figure 6a). Yet when a person turns to look at the display directly, it changes into a photo of the couple pretending to be happy (Figure 6b). A slight delay is introduced so that people can get a glimpse of the switch-over. This example is also relevant to the Attention Grabber pattern.

Bait and switch also exists in other proxemics-aware systems that do not use public displays. Consider public wireless networks such as those at airports. They detect travelers within its range, and offer the bait of what appears to be free-of-charge wireless. Yet once a traveler is apparently connected, the network may require the traveler to give up information by signing into some service, or the offered 'free' service may be so slow that the alternate higher quality pay service is the only realistic offering.

4. MAKING PERSONAL INFORMATION PUBLIC

As the person enters a particular area, the system makes that person's personal information publicly visible.

One of the appeals of proxemic interactions is to make personal information readily available on nearby devices. Vogel et al's original work on ambient displays [28] illustrated how




Figure 6. Amnesty International Eye Tracking (sketch by David Ledo).

a public ambient display reveals both public and personal information as a person approaches it. Personal information includes calendars, notifications, and directed messages, which can then be manipulated by that person.

Their system is intended to be helpful. Yet the basic issue is that other onlookers can see that personal information. Vogel et al. tried to mitigate this by describing how the person's body could physically shield personal information presented directly in front, and how the person could hide information through an explicit gesture [28].

The previously mentioned scene from 'Minority Report', with its myriads of advertising walls, make passerby's private information public as a byproduct of their clamor for attention. We see the Guinness advertising wall (amongst others) publicly identify the protagonist by shouting out his name. In that scene, another advertising wall for a credit card visually displays both the protagonist's name and personal information about him (that he has been a member since 2037).

Making personal information public could be an intentional design goal rather than an unintended side effect. An example is the guerilla-style bus stop display produced for the Fitness First health club chain in Rotterdam (Figure 7). The bench nearby the display contains a weight sensor, where the commuter's weight is then publicly displayed on the bus



Figure 7. The Fitness First health club display. © Dutch ad agency N=5, with permission.

stop's wall. Its purpose is purportedly to motivate people to join the health club by intentionally publicizing their weights and offering them a price reduction fee of their weight in Euro's in return.

5. WE NEVER FORGET

In day-to-day life, proximity is an ephemeral phenomenon. The proxemic relationship between parties dissolves as soon as they separate. In contrast, systems can tag any proxemic interactions as indicating a permanent, persistent (and undesirable) relationship that is never forgotten.

The 'we never forget' pattern occurs when systems maintain a history of peoples' past proxemic connections, where that history is used to re-establish connections, to trigger information exchange, and/or to recreate prior contexts (e.g., showing the last-displayed information). When used beneficially, the idea is to remember details that make it easy to pick up where one has left off. Unfortunately, this might be completely inappropriate in a different context.

For instance, mobile devices – when brought into range of other devices – typically remember any entered credentials (such as a passphrase) that allow both to connect to one another. This can be a tremendous convenience: when that device comes back into range, those credentials are reused automatically to re-establish the connection, minimizing user load. Remembered WiFi hotspots automatically reestablish network connections when a device returns to a location, while Bluetooth pairings ease device to device interconnections, such as how a person's mobile phone is linked to a hands-free system in that person's car. Similarly, various interaction techniques trigger pairings and information exchange when proxemics-aware devices are brought close together, e.g., by bringing mobile devices together [14,15,20].

On the other hand, this approach can fail for several reasons. First, people may do a one-off connection to a device they otherwise do not control or trust (e.g., a one-time transaction with a public display). If that person happens to pass by that other device at a later time, there is no reason for that connection to be re-established (particularly if there is some risk involved). Second, security is compromised. If (say) one's

mobile phone is stolen, the thief may be able to explore nearby locations to see if he or she can access other devices or networks without entering any credentials.

Third, circumstances change even with trusted devices. For example, a person that previously used a conference room display to show some personal photos on his phone to visiting friends could have these photos reappear inappropriately on the display while walking past it with her work colleagues. Or, consider the case of cell phones paired to one's Bluetooth car system, where it automatically displays incoming calls and redirects audio to the car's speakers. We can easily imagine what could happen on a family trip when an incoming call from one's secret lover is broadcast for all to see on the radio consul and the lover's greeting heard if accepted. As another example, a manager and an employee may be working physically close together, where they pair their laptops to work on a project report. A week later, the manager and employee sit next to each other in a meeting discussing the team's progress. As their laptops get close to each other again, the manager's laptop automatically shares the currently opened document, which, in this case, is a sensitive spreadsheet with the wages of all team members.

Fourth, a person may be unaware that he or she is again sharing his or her device's data with another person that they had previously shared with. This absence of reciprocity (if you share with me, I should know that I share with you) is a known problem in groupware, where one of the parties may be unaware that one's data is being shared with others. To remedy this, such systems should provide awareness of other users and their actions [4]. When proxemic connections are established, the system needs to inform its users about what information is being shared and when, and to whom this information is made available (who is making a connection? [24]). Likewise, users need to know what will happen to their information once it is shared [4], and what happens once the connection is destroyed.

Finally, credentials obtained in one setting may be remembered by the system and inappropriately applied to other settings. This 'one login for all' is an increasingly common practice in other systems, such as Facebook or Google. The danger, of course, is that a person who has established a single proxemic connection to (say) a particular display may not want that connection to occur when they happen to pass by other associated displays.

6. DISGUISED DATA COLLECTION

Information gathered to provide a certain service is abused to build a rich user profile, without the consent of users.

Systems that track proxemic relationships have access to large quantities of data about the behavior of their users. Public advertising displays that track the user's distance, location, orientation, and movement are a goldmine for marketers, who can exploit this information to figure out which ads users are looking at and for how long. Fortunately, personal risk is somewhat mitigated as long as the person is not

equipped with technology that can be tapped to uniquely identify him (e.g., broadcasting cell phones, RFID chips, smart cards).

Unfortunately, many public displays rely on some form of computer vision to track proxemic relationships. Given that the installation has access to images of its users anyway, it is entirely plausible to use image analysis to try and uniquely identify users. Systems such as these would make the targeted advertisements from Minority Report a reality. Indeed, there are already commercial systems – such as the Aware-Live Technologies Look product (to be found at: <http://www.aware-live.com/>) – that use computer vision to identify characteristics of its users such as gender, approximate age, and a classification in marketing segments (e.g., Generation X). Similar to the AIDA model mentioned earlier, Aware-Live’s mantra is “recognize [demographics], analyze [to make intelligent decisions] and engage [to interact with customers in a precise manner]”.

Similarly, free WiFi services can collect a person’s location inside stores by tracking the signal strength and IP of their device to different WiFi hotspots. For example, Euclid Analytics offers services that measure walk-by traffic, visit duration, and even brand loyalty. If the store offers the WiFi service, it can potentially track their browsing behavior via web server proxies (<http://euclidanalytics.com/>).

These and other data collection approaches can be combined to build an even richer user profile. Indeed, this would allow systems to exploit the user’s proxemic history, thereby leveraging the ‘We Never Forget’ pattern. Just like so-called ‘loyalty’ cards track a person’s shopping behaviors, the user’s location could be tracked when they walk past different advertising displays and locations, where the personal profile is both constructed by and shared between these systems, thereby allowing information to ‘chase’ the moving person.

7. THE SOCIAL NETWORK OF PROXEMIC CONTACTS OR UNINTENDED RELATIONSHIPS

The system tracks your proxemic relations with others and constructs a social network on the assumption that you are somehow socially related, when there is no relationship.

Proxemics assumes that increasing social engagement (and thus a social relationship) is typically accompanied by decreasing physical distance, mutual orientation, etc. That is, social engagement leads to people adjusting these factors to their mutual benefit. Proxemic interaction systems do this somewhat backwards. They assume that some sensed phenomena (decreasing physical distance, mutual orientation, etc.) *signals* a social relationship, i.e., it treats the sensed phenomenon as causal. This assumption is not always correct. In real life, strangers may approach and even glance at each other, but no social relationship exists between them. Moreover, not all relationships are reciprocal: while one person believes they have a relation to another, the other may not reciprocate at the same strength, if at all.

The assumption that *all* proxemic interactions imply a social relationship is problematic for a variety of reasons. Perhaps the most worrisome is that the underlying system may be trying to infer one’s social network from proxemic events between two people, where strangers are included. This scenario is not at all farfetched. In 2013, Edward Snowden revealed the US National Security Agency’s controversial practice of tracking phone call metadata records (the number dialed, a cell phone’s location, time and duration of call, etc.). They used this information to compile sophisticated social network diagrams of Americans, ostensibly to identify and target terrorist networks. Even if one accepts this practice, innocent parties may be inadvertently included as ‘false positives’ in one’s social network, perhaps due to erroneous (wrong numbers) or innocuous calls.

It would be just as straightforward to create an equivalent social network by sensing one’s proximity to others. These too could easily include unintended relationships. For example, matches between location and time information in cell phone metadata records can be used to determine those people in the same proxemic vicinity. Eagle et. al. compared observational data from mobile phones with self-report data, and concluded that they could accurately infer 95% of friendships based on observational data alone [10]. This also means that 1 in 20 are *not* friendships (i.e., they are false positives). Other technologies can provide even more accurate data of one’s proximity to another and thus record that as a potential relationship, e.g., facial recognition systems identifying co-located people in a public place, or passing by the front of an identity-sensing device (such as a large display).

Once created, the social network could be used for a variety of dark purposes. Authorities could exploit the social network to identify potential ‘suspects’ by their inferred association to an unsavory character. Marketers could use that social network to identify a potential target audience by their association with a known demographic fitting that profile. Spammers and phishers could exploit it for their own deceptive purposes. In all cases, the agencies involved may not care that ‘false positives’ are included, where they may be treated as collateral damage or simply as noise.

While algorithms could perhaps detect and minimize the number of false positives, the social network will always include some unintended relationships.

8. THE MILK FACTOR

The proxemics system forces you to move through or go to a specific location in order to get a service.

The rules of proxemic interactions, which we use in our everyday lives, can be misused to force people to move to or from a specific location. In non-computer scenarios, this can be seen in the design of supermarket spaces. Products that are purchased frequently (e.g., milk, or bread) are located in distant areas of the store. Thus, shoppers are forced to walk through isles with goods, which leads to increased visibility of promoted items and impulse purchases.

Proxemic interaction systems can force people to position themselves in specific places by limiting access to functionality to particular locations. For example, all zone-based proxemic displays invoke certain types of interactions at specific distances. While most research systems do this with good intentions (e.g., [28,2,30,8,25]), all require its user to stand within specific boundaries.

This can have unintended consequences. MirrorSpace [25] is a video conferencing system that mitigates privacy: images are blurred when the person is far away and only become sharp and identifiable when they stand close to the display. If the person needs to be in another corner of the room while talking over the link, they lose fidelity. In the video player by [8], the motivation for limiting visibility of information is to allow some viewers to watch a film with subtitles from a position on the right side of the sitting area, while simultaneously allowing others to watch the same film without the subtitles on the left side of the sitting area. This forces people to sit in particular locations if they want to see a movie a particular way (vs. sitting on the floor).

A commercial example that exploits people having to go to a specific location is the *Design Studio S* vending machine in Japan (see www.design-ss.com/products/2010/09/01/vending-machine.html?ctg-jp). When potential customers are far, the vending machines show advertising images tailored to the season, time of day and temperature. However, to see what drinks are available for purchase, the potential customer must approach the area in front of the vending machine, which only then shows a drinks menu. However, at that point the vending machine uses its camera to covertly perform a computer vision based analysis of the nearby customer to establish their approximate age and gender, as in the *Disguised Data Collection* pattern. This data is used to “subtly” offer targeted drinks selections. Demographic and sales data is uploaded (without consent) to the company’s servers for further analytics and marketing use. This is a clear example of a dark pattern: the customer cannot even see the range of drinks for sale, which forces them to move close enough to the machine to make covert data collection possible. This also raises the spectre of price discrimination (already practiced on some web sites), where an identified demographic may be charged differently depending on their ability to pay.

DISCUSSION

We acknowledge that the patterns we discussed are a sampling rather than an exhaustive list. Even so, they suffice for reflection. From our sample, we were able to identify several common root problems that can be exploited as dark patterns, or that promote side effects that then lead to anti-patterns. In the following (1) we discuss these problems in more detail, (2) denote how they apply to the aforementioned patterns, and (3) aim to identify a code of conduct where applicable.

Opt-in / opt-out choices are particularly problematic in proxemics-aware systems. The overall problem is that a person implicitly opts-in simply by entering a space and approaching the proxemics-aware entity, regardless of whether the

person actually intends to opt into the situation. Currently, opt-out requires the person to leave that area, which may not be a reasonable choice for them (e.g., as in the *Captive Audience*). Opt-out may further inflict uncertainty about what will happen to traces of bygone interactions (e.g., trails of personal information on public displays, as in *Making Personal Information Public* and *We Never Forget*).

There is a clear trade-off. Implicit opt-in strategies are popular because they both simplify interaction (from the user’s perspective) and increase engagement (from the vendor’s perspective). Yet their high potential for misuse is problematic (e.g., as in the *Disguised Data Collection* or *Unintended Relationships* pattern). At the very least, proxemic interaction systems must have a way to opt-out if interaction is not desired. Leaving the space, while simple, may not always be a practical option. Explicit user actions are also possible, such as invoking a particular gesture to opt out [17], or turning off services on personal devices. Yet these require both learning and extra work.

Physical space is imbued with dual meanings. People’s practices and expectations of the physical space can be quite different from the meaning and practice applied by the technology. This means that a person may approach a location for one reason, but as a consequence they are exposed to the system exploiting their approach for another reason (e.g., simply wanting to walk past a display as in the *Attention Grabber* pattern).

In many of the discussed patterns, a user’s context plays an important role. For example, being surrounded by many commuters in a subway may form highly *Unintended Relationships* simply due to the close proximity of others.

One possible solution is to gather *more* contextual information to better infer whether a person is using the physical space as is, or whether they actually have an interest in the system. For example, an Attention Grabber can sense a person’s speed to determine whether they are in a hurry, and thus let them pass by undisturbed. Unintended Relationships can be avoided by comparing its collected data to other data sources that mine friendship data, such as social network data. Of course, this introduces other concerns.

Ownership of the physical space is ambiguous. A person looking for a quiet corner may consider that space as temporarily their own, but if this happens in a public area, their presence can still be exploited. Yet a public display may consider the installation space around itself as its own, where any person (and the devices they carry) in that space becomes fair game. While people have social rules that dictate what happens when interacting in private, personal, or public space, technology can easily violate those rules (e.g., an obnoxious display invades a person’s privacy with targeted advertising as in the *Attention Grabber* pattern).

We believe it is crucial to define who *owns* the space around a proxemics interaction system. This is particularly true for public spaces that people perceive as owned by them. Yet,

the definition of a public space is somewhat vague. Consider the urinal in the *Captive Audience* pattern: a company running public restrooms may own this space, but the person using it would consider it a private enclave. Ultimately, there has to be some control and rules for *who* is allowed to do *what* in a given space. At the very least, the system must make it clear (e.g., by its visuals, or by marking) that it has taken a certain amount of space for its own use.

Attention is inherently sought after in proxemic interactions. The gradual engagement design pattern [20] suggests that proxemic interactions gradually reveal information as entities approach one another (e.g., as in the *Making Personal Information Public* pattern). Whether done subtly or blatantly (as in the *Attention Grabber* pattern), attention of the person is demanded – even if that person has no intent to interact with the system.

The problem is that a user's context (and his or her willingness to pay attention to the system respectively) again plays an important role. That is, people should be able to move through a space with a proxemic system installed without being affected by it if they (maybe explicitly) opted out of being part of the system.

Accidental proxemics occurs when people unintentionally enter what could be interpreted as a proxemic relationship. They may approach and even orient themselves towards something with no real intent of engaging with the system. Yet inferences of such a relationship leads to problems, such as engaging people without consent in the *Captured Audience*, and the accidental sharing of private information in the *We Never Forget* pattern. If the approach is due to another reason (e.g., just walking past a display), it becomes relatively hard to discriminate that action from an intentional opt-in to use the system. Bellotti et al. describe this as one of the typical challenges in context-aware systems, where it is difficult for users to know when their actions are being attended to [3].

When proxemic systems interpret *any* approach action as the start of a proxemic relationship, users cannot enter a space without triggering the system (similar to the *Midas Touch problem* [29]). For example, smart keys for various cars now allow one to automatically unlock and lock the doors of a car when approaching or leaving the car. However, the person cannot physically verify that the doors are locked, as approaching the car again will unlock them.

Accidental proxemics is a particularly nasty variation of opting in vs. opting out. Similar to the other root problems, avoiding accidental proxemics is difficult if intention is sensed implicitly. No matter how carefully done, the system will sometimes get it wrong.

Ideally, proxemic interactions systems must strike a delicate balance between implicit and explicit interaction, and by making users aware of what is happening [17]. While the so-

lution is to intervene and override the proxemic system's behavior if it does not correspond with their intentions, it demands that they do extra explicit work.

CONCLUSIONS

In this paper, we reconsidered the vision of proxemic interactions through a critical lens. We identified potential *dark patterns* and *anti-patterns* demonstrating how proxemic systems can abuse people either intentionally or unintentionally. Based on these patterns, we discussed several common root problems and speculated on potential solutions. Solutions are at best tentative, but we hope that they could evolve into a code of conduct taken into account by designers, with the goal of both lowering the risk of intentional abuse and unintentional design flaws.

Unfortunately, this may be easier said than done. At least two parties are involved in proxemic interaction systems: the party deploying the system vs. the system's users. Both may have quite different intentions and desires. For example, if the goal of system stakeholders is to acquire a person's attention, the actual users may have little chance of opting out. Thus legislation may play a role, as it has in other cases of a mismatch in interests. For example, governmental authorities have (to some extent) enforced rules to better protect users from the excesses of online e-commerce systems, and to limit spammers and phishers.

Another and perhaps much better solution is to consider proxemic interactions systems design from a mutually beneficial perspective. No company wants to be viewed negatively, and we expect good companies will pay attention to how potentially dark patterns can be turned into a good (or at the worse a neutral) pattern that balances benefits to the company, client and product user. This already happens in the advertising industry, where the best ads provide value to its viewers (e.g., humor, engagement, interest, etc.) as part of its service. Indeed, several of our examples already do this. The Nikon D700 Billboard of Figure 4 is an example of an entertaining and novel guerilla ad that invokes curiosity; its cost is also small – the red carpet suggesting the direction to the store selling the displayed cameras. Proxemic Peddler (Figure 5) uses subtle rather than aggressive visuals to strike a balance between how it senses and reacts to people's attention vs. loss of interest [30]. The Amnesty International campaign in Figure 6 presents an important (albeit disturbing) social message to the public.

Captive Media offers a particularly positive example of institutional awareness. Their representative told us how their policy deeply considers mutual benefits. They consider Captive Media's urinal (Figure 2) primarily as an interactive entertainment device into which they incorporate some advertising: it enlivens a time that is otherwise boring, and they see it serving as a social ice-breaker (e.g., bars, hotel events). During actual urinal use, they advocate that the majority of time is for game play, with only a fraction of the time for advertisements – that is, the cost of play is small. When acting as an ambient display, they advocate that it must display

entertaining content at least 50% of the time. They also advocate outfitting only a portion of the urinals in a bathroom with their display, thus allowing people to opt out. Finally, they collect experience reactions, where they found that Captive Media users are overwhelmingly positive.

Yet this is still early times. Even if proxemic interactions systems were designed to avoid abuse, problems will inevitably cause user frustration, likely due to well-known issues in implicit interaction [3]. This remains a grand challenge.

ACKNOWLEDGEMENTS

This paper was inspired by Aaron Quigley and a break-out session at the Dagstuhl Seminar on Proxemics in Human Computer Interaction (#13452). Intellectual contributors to the early stages of this work included Petra Isenberg, Fabrice Matulic, Thomas Pederson, Stacey Scott and paper authors. Additional grant funding by NSERC, SMART and AITF.

REFERENCES

Web site descriptions, quotes, images, and videos referenced in this paper are based on material retrieved in January 2014.

1. Alexander, C. (1977) *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
2. Ballendat, T., Marquardt, N. & Greenberg, S. (2010) Proxemic Interaction: Designing for a Proximity and Orientation Aware Environment. *Proc. ACM ITS*, 121-130
3. Bellotti, V., Back, M., Edwards, W. K., Grinter, R. E., Henderson, A., & Lopes, C. (2002) Making Sense of Sensing Systems: Five Questions for Designers and Researchers. *Proc. ACM CHI*, 415-422.
4. Bellotti, V. & Edwards, K. (2001) Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human-Computer Interaction*, 16(2-4), 193-212.
5. Borchers, J. (2001) A Pattern Approach to Interaction Design. *AI & Society*, 12:359-376, Springer.
6. Brignull, H. (2011) Dark Patterns: Deception vs. Honesty in UI Design. *Interaction Design, Usability*, 338.
7. Brignull, H., Miquel, M. & Rosenberg, J. Dark Patterns Library. <http://darkpatterns.org>. Retrieved Nov. 2013
8. Dostal, J., Kristensson, P. & Quigley, A. (2013) Multi-view Proxemics: Distance and Position Sensitive Interaction. *Proc ACM. Pervasive Displays*, 1-6.
9. Dourish, P. (2004) *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press
10. Eagle, N., Pentland, A. & Lazer, D. (2009) Inferring Friendship Network Structure by Using Mobile Phone Data. *Proc. Nat'l Academy of Science*, 106:36, 15274-8.
11. Gamma, E., Helm, R., Johnson, R. & Vlissides, J. (1994) *Design Patterns: Elements of Reusable Object-Oriented Software*. Pearson.
12. Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R. & Wang, M. (2011) Proxemic Interactions: The New Ubicomp? *ACM interactions*, 18, 42-50.
13. Hall, E.T. (1966) *The Hidden Dimension*. Doubleday, N.Y.
14. Hinckley, K. Synchronous Gestures for Multiple Persons and Computers. (2003) *Proc. ACM UIST*, 149-158.
15. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., & Smith, M. (2004) Stitching: pen gestures that span multiple displays. *ACM Proc. AVI*, 23-31.
16. Isenberg, P., Dragicevic, P., Willett, W., Bezerianos, A. & Fekete, JD. (2013) Hybrid-Image Visualization for Large Viewing Environments. *IEEE TVCG*, 19(12).
17. Ju, W., Lee, B. & Klemmer, S. (2008) Range: exploring implicit interaction through electronic whiteboard design. *Proc. ACM CSCW*, 17-26.
18. Koenig, A. (1995) Patterns and Antipatterns. *J. Object-Oriented Programming*, 8 (1): 46-48.
19. Kortuem, G., Kray, C. & Gellersen, H. (2005) Sensing and visualizing spatial relations of mobile devices. *Proc. ACM UIST*, 93-102
20. Marquardt, N., Ballendat, T., Boring, S., Greenberg, S. & Hinckley, K. (2012) Gradual Engagement between Digital Devices as a Function of Proximity: From Awareness to Progressive Reveal to Information Transfer. *Proc. ACM ITS*, 31-40.
21. Marquardt, N. & Greenberg, S. (2012) Informing the Design of Proxemic Interactions. *IEEE Pervasive Computing*, 11(2):14-23.
22. Marquardt, N., Hinckley, K. & Greenberg, S. (2012) Cross-Device Interaction via Micro-mobility and Formations. *Proc. ACM UIST*, 13-22.
23. Michelis, D. & Müller, J. (2011) The Audience Funnel: Observations of Gesture Based Interaction with Multiple Large Displays in a City Center, *Int. J. HCI*, 27 (6),
24. Pierce, J. S., Mahaney, H. E., & Abowd, G. D. (2003). Opportunistic Annexing for Handheld Devices: Opportunities and Challenges. *Proc. HCIC*.
25. Roussel, N., Evans, H. & Hansen, H. (2004) Proximity as an Interface for Video Communication. *IEEE Multimedia*, 11(3):12-16.
26. Snibbe, S. & Raffle, H. (2009) Social Immersive Media: Pursuing Best Practices for Multi-User Interactive Camera/Projector Exhibits. *Proc. ACM CHI*, 1447-1456.
27. Strong, E. *The Psychology of Selling and Advertising*, McGraw-Hill NY, (1925).
28. Vogel, D. & Balakrishnan, R. (2004) Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. *Proc. ACM UIST*, 137-146.
29. Velichkovsky, B., Sprenger, A., & Unema, P. (1997). Towards Gaze-Mediated Interaction: Collecting Solutions of the "Midas Touch Problem". *Proc. INTERACT*, 509-516.
30. Wang, M., Boring, S. & Greenberg, S. (2012) Proxemic Peddler: A Public Advertising Display that Captures and Preserves the Attention of a Passerby. *Proc. ACM Pervasive Displays*.
31. Zagal, J., Bjork, S. & Lewis, C. (2013) Dark Patterns in the Design of Games. *Proc. Foundation of Digital Games*, <http://www.fdg2013.org/program/>