

Balancing Privacy and Awareness for Telecommuters Using Blur Filtration

Carman Neustaedter, Saul Greenberg, and Michael Boyle

University of Calgary, Canada

{carman, saul, boylem}@cpsc.ucalgary.ca

Abstract. Always-on video provides rich awareness for co-workers separated by distance, yet it has the potential to threaten privacy as sensitive details may be broadcast to others. This threat increases for telecommuters who work at home and connect to office-based colleagues using video. One technique for balancing privacy and awareness is *blur filtration*, which blurs video to hide sensitive details while still giving the viewer a sense of what is going on. While other researchers found that blur filtration mitigates privacy concerns in low-risk office settings, we do not know if it works for riskier situations that can occur in telecommuting settings. Using a controlled experiment, we evaluated blur filtration for its effectiveness in balancing privacy with awareness for typical home situations faced by telecommuters. Participants viewed five video scenes containing a telecommuter at ten levels of blur, where scenes ranged from little to extreme privacy risk. They then answered awareness and privacy questions about these scenes. Our results show that blur filtration is only able to balance privacy with awareness for mundane home scenes. The implication is that blur filtration by itself does not suffice for privacy protection in video-based telecommuting situations; other privacy-protecting strategies are required.

Introduction

In this paper, we examine how well different levels of video-blurring safeguard privacy in always-on video links that connect home-based telecommuters with office colleagues. As will be shown, we are interested in not only mundane telecommuting situations, but in high privacy-risk situations that sometimes occur

Cite as:

Neustaedter, C., Greenberg, S. and Boyle, M. (2003) Balancing Privacy and Awareness for Telecommuters Using Blur Filtration. Yellow Series Report 2003-719-22, University of Calgary, Canada. January,

due to the peculiarities of home work. Before going into detail, we offer a brief background on: casual interaction and awareness, the role of always-on video for distance-separated colleagues, the privacy concerns that can arise from telecommuting, and how previous work suggested that video distortion filters (such as blurring) mitigate privacy concerns.

Throughout a typical day, co-workers naturally converse and interact among each other in what is known as *casual interaction*—the frequent and informal encounters that either occur when people serendipitously meet or are initiated by one person (Fish et al, 1993, Hudson and Smith, 1996). Casual interactions foster knowledge and help individuals accomplish both individual and group work (Kraut et al, 1988, Fish et al, 1993). *Informal awareness*—an understanding of who is around and available for interaction—holds casual interaction together by helping people decide if and when to smoothly move into and out of conversation and collaboration (Kraut et al, 1988, Bellotti and Sellen, 1993). Informal awareness is easily gained when people are in close physical proximity, but deteriorates over distance (Kraut et al, 1988, Greenberg, 1996). As a result, casual interaction suffers when co-workers are distributed.

One possible solution for providing awareness between distance-separated collaborators is to use an always-on video link to connect remote locations (Mantei et al, 1991, Fish et al, 1993, Bly et al, 1993, Tang et al, 1994, Bellotti, 1996, Greenberg, 1996, Lee et al, 1997, Greenberg and Kuzuoka, 2000). While always-on video can provide rich awareness, the problem is that it also broadcasts information that individuals judge as privacy sensitive (Bellotti and Sellen, 1993, Bly et al, 1993, Bellotti, 1996, Hudson and Smith, 1996, Bellotti, 1998, Greenberg and Kuzuoka, 2000, Boyle et al, 2000). The goal of many researchers is to find a balance between the rich awareness provided by video-based media spaces and the privacy concerns they raise.

In practice, video media spaces have found some limited success in office situations. Most installations simply ignore privacy issues: risks are fairly low in office settings, and simple privacy safeguards often suffice e.g., people can explicitly switch off the video channel, or turn the camera around to face a wall. Yet, the situation becomes complicated when people choose to work from home as telecommuters while still desiring close contact with colleagues at work. The big problem is that privacy risks increase drastically for the telecommuter. The home is not the office; activities and appearances appropriate in the home are often inappropriate when viewed in an office environment by a colleague. For example, consider these following situations—all derived from actual events reported to us by telecommuters—where the telecommuter habitually uses always-on video to provide a colleague at work with awareness. These situations are all very realistic and threaten the privacy of the telecommuter, as well as others within the home.

Imagine yourself as a telecommuter living the dual role as worker and as home occupant. It is a hot day, and you are going shirtless (or as a female, you are wearing a very skimpy top). Forgetting you are shirtless (because this is not a problem at home), you enter your home office to quickly check your email. The always-on video captures you wearing no shirt. While it is quite appropriate for you to not wear a shirt while at home, the same level of dress may be seen by your colleagues (at the office) as inappropriate.

Our second example results from a telecommuter's unconscious acts, their ease of forgetting that their distant colleague is (virtually) sitting right across from them, and from the lack of feedback that they are actually in a public setting. Imagine yourself again as the telecommuter working at your home computer when you suddenly sneeze. Naturally, you proceed to blow your nose, forgetting that a camera on top of your monitor captures this at a *very* close range. You next begin to pick your nose at great length. Your colleague views the scene over the video link and is disgusted at how inconsiderate you are being.

Family members and friends in the home likely gain no benefit from the video link yet still incur its privacy threat. Now imagine you, as the telecommuter, are working in your home office in the early morning when your spouse (who has just woken up) enters the room wearing only pajamas and gives you a big wet kiss. All this is captured on camera, and your colleague has seen this. Your spouse realizes this and becomes infuriated, telling you never to use the camera again.

Our final example results from the dual purposes typical of most home offices. Imagine what might happen if your home office is also your spare bedroom. One hot day you take a shower in the bathroom next door. You towel off, and then go into the spare bedroom to retrieve a bathrobe in the closet. A few moments after entering the room, you realize that the camera is directed at you. You drop to the floor, crawl to the camera, and knock it off the computer.

In an effort to help mitigate privacy concerns over video links, other researchers have studied *distortion filters*: algorithmic reduction of image fidelity to hide sensitive details in a video image (Zhao and Stasko, 1998, Greenberg and Kuzuoka, 2000, Boyle et al, 2000). Specifically, Boyle et al (2000) studied two distortion filters and how they balance privacy and awareness in mundane and benign office situations, e.g., people working or reading, people chatting, people eating lunch. The two filters were the *pixelize filter* that produces a mosaic of solid rectangles, and the *blur filter* that averages pixel values to produce a blurred effect. They found that each filter offered a filtration level that adequately preserved privacy and still provided awareness for these (office) situations. The blur filter, however, was found to balance privacy and awareness over a wider range of filtration levels than the pixelize filter. Furthermore, these levels were

more heavily filtered for blur filtration than for pixelize, thus they preserved more privacy.

However, Boyle et al (2000) did not study the effects of their distortion filters on situations where one's privacy may be at extreme risk, such as those typified in our telecommuting examples. Consequently, we set ourselves the research goal of determining how well video-blurring safeguards privacy in always-on video links that connect home-based telecommuters with office colleagues. To achieve this goal, we constructed a controlled experiment to test blur filtration with a set of scenes that typify home telecommuting and range greatly in perceived privacy risk. Scenes included: mundane situations, such as working at a computer, moderately risky situations, such as the telecommuter kissing her partner, and extremely threatening situations, such as being shown completely naked. The next section of this paper outlines the study's methodology and includes the specific research questions the study answers. Following this, we discuss the study's results and its implications for design of video-based media spaces.

Methodology

In our study, participants are asked to imagine a scenario where they are the close-working colleague of a telecommuter. Participants then view a series of five video scenes—each blurred at ten different levels of blur—containing the telecommuter. For each blur level, participants answer privacy and awareness questions, described in more detail later. We first describe our hypotheses and variables, followed by our materials and procedure.

Hypotheses

We test three null hypotheses in the study. The first null hypothesis analyses the viewer's ability to extract awareness cues from video scenes at ten different levels of blur. The second null hypothesis analyzes how the ten levels of blur filtration affect the perceived privacy threat within each of the video scenes. The third null hypothesis analyses each scene's effect on the viewer's selection of blur level.

Hypothesis 1: There is no difference in a viewer's ability to determine particular *awareness cues* from ten different levels of blur (from fully blurred to no distortion) applied to five different videos containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

Hypothesis 2: There is no difference in the severity of *perceived privacy threat* to the telecommuter and family members presented by each of the ten levels of blur (from fully blurred to completely clear) applied to five different videos

containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

Hypothesis 3: There is no difference in the viewer's *selection of blur level* as they try to make a particular scene appropriate for viewing by a distant colleague for five different videos containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

Independent and Dependent Variables

The independent variables for the study are scene type (5) \times blur filter levels (10). The dependent variables recorded in a during-test questionnaire (described later) are: a participant's ability to correctly identify awareness cues, a participant's confidence in identifying awareness cues, a participant's perception of the videos' level of privacy threat, and the chosen blur level for safeguarding each video.

Materials: Video Scenes

We recorded five video scenes that vary in the level of risk presented, from scenes we judged to have no risk to those with very high risk. Each scene shows a telecommuter performing a different activity or with a different appearance, where all scenes are recorded from the same point of view, i.e., behind the computer monitor at the same angle (Figure 1). The scenes are sorted by expected perceived privacy risk, from low risk to high risk:

- (1) ***Working at a computer:*** The telecommuter is working at a computer while wearing clothes appropriate for both home and the office. *Low risk.*
- (2) ***Picking one's nose:*** The telecommuter is working at a computer wearing clothes appropriate for both home and the office when he/she begins to pick his/her nose. *Moderate risk.*
- (3) ***Working with no shirt on:*** The telecommuter is working at a computer with no shirt on. *Moderate risk.*
- (4) ***Kissing a partner:*** The telecommuter is working at a computer when his/her partner enters the room, kisses the telecommuter intimately, and leads him/her out of the room. *Moderate risk.*
- (5) ***Changing clothes / Naked:*** The telecommuter enters the room in a robe, is shown completely naked, and then puts on underwear. *High risk.*

Each scene was recorded twice, once with a paid male as the telecommuter and once with a paid female as the telecommuter. The actors were deliberately chosen to be middle-aged individuals with the appearance of working professionals. The final videos were 720 \times 480 pixels at 30 frames per second (fps) DV-format, compared to those used in Boyle et al's (2000) study, which were Intel Indeo™ compressed at 176 \times 144 pixels and 24 fps.

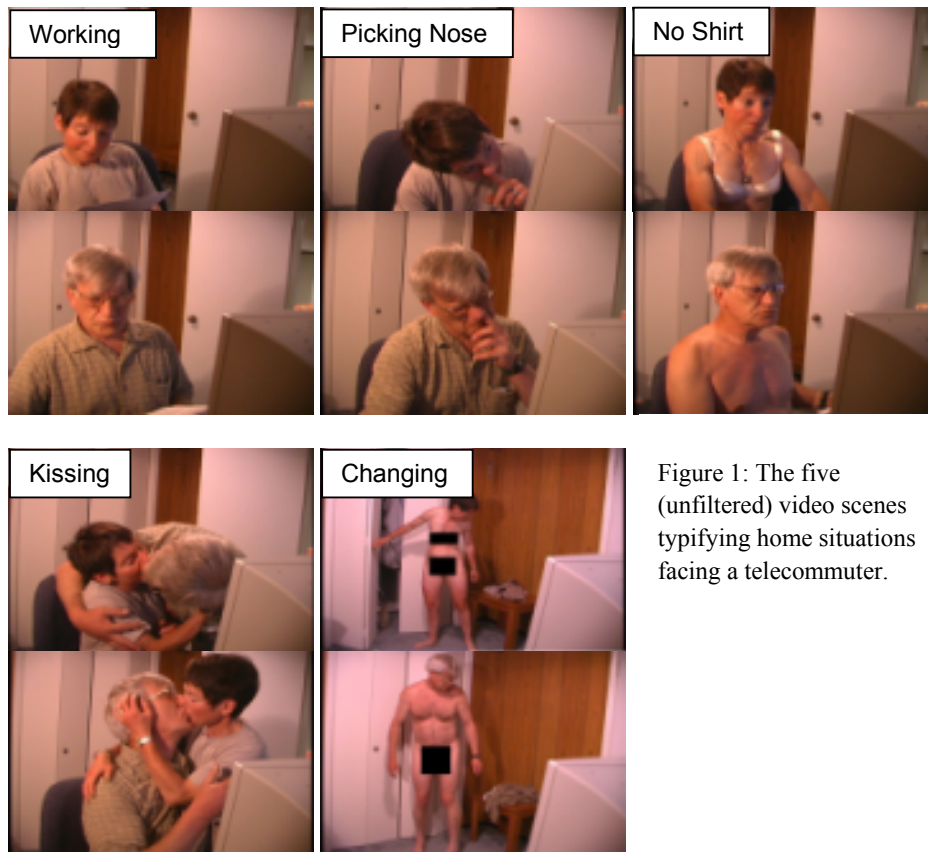


Figure 1: The five (unfiltered) video scenes typifying home situations facing a telecommuter.

Materials: Scenarios Provided to Participants

To set the context of a *home media space*—defined as an always-on video media space used within a home setting—we asked participants to imagine they were the colleague of the telecommuter shown in the videos, nicknamed Larry for males and Linda for females. Furthermore, they had known Larry (or Linda) for more than a year and shared with them a close working relationship. Larry (Linda) had recently decided to work from home and a video link would be used to maintain this close working relationship.

Materials: Blurred Video Scenes

The ten video scenes (five male, five female) were also pre-processed to create a set of videos at each of the ten different blur levels to be evaluated (Figure 2). We used the same algorithm to blur our images as Boyle et al (2000) and our blur levels are roughly equivalent. Our distortion algorithm computes a filtered pixel's value as the unweighted average of itself and its neighbors. The larger the neighborhood, the greater the distortion. This is a typical method for smoothing (i.e., blurring) an image. The algorithm is applied on a frame-by-frame basis.




Figure 2: The blur levels evaluated in the study (the 10th level, not shown, is the unfiltered scene).

Materials: Questionnaires

A pre-test questionnaire gathered demographics. In our results, we discuss only gender. A during-test questionnaire asked participants about each of the blur levels for all video scenes. The questionnaire was web-based and used two 17" CRT displays: the left display showed a video scene, while the right display presented questions about the video. The awareness and privacy related questions asked for each of the blur levels are shown in Figure 3 (top). Similarly, Figure 3 (bottom) shows the set of questions used for each scene after the participant viewed all the blur levels for it. These questions ask the participant to choose a blur level that would make the scene appropriate for a colleague to view. A post-test questionnaire gathered each participant's opinion of balancing privacy and awareness using blur filtration, and asked participants if they would use an open video link in an office if it was blurred and also at their home if it was blurred.

A final question asked participants to perform a forced sort of five pictures (one for each video scene, printed on standard 21.59 x 27.94 cm pieces of paper) according to how risky they felt each scene was to their privacy if they were the person in the scene. Participants were then asked to place the sorted pictures on a "line of privacy risk" that was 300 cm long: one end represented low risk, the other end represented high risk. Participants were told that they could leave as much space between pictures as they liked, but no two pictures could overlap. The "line of privacy risk" is used in a post-hoc assessment/validation of our original rating of each scene's privacy risk.

Mediating Home Privacy

1. Describe what you see in the video:

Who can you see?

Unsure

Confident

If you can see a person, what is he doing?

Unsure

Confident

If you can see a person, what is he wearing?

Unsure

Confident

What else can you see?

Unsure

Confident

2. How available is Larry for you to talk to right now?

It's a bad time.

It's a great time.

Unsure

Confident

Why?

3. Given what you can see at this blur level, how threatening is this scene to Larry's privacy?

Not Threatening

Very Threatening

Why?

4. Given what you can see at this blur level, how threatening is this scene to Larry's family members?

Not Threatening

Very Threatening

Why?

View the next blur level

Mediating Home Privacy

5. Please choose the level of blur that would make this scene appropriate for Larry to see over the video link while at the office. Remember, you want to stay in close contact with Larry.

Full Blur

No Blur

Turn Camera Off

6. If anything, what are you trying to mask in the video by blurring it or turning the camera off?

View the next video scene

Figure 3: During-test questionnaire (extracted from the study): privacy/awareness questions for each blur level (top), choosing a blur level (bottom).

Participants

Participants were ten females and ten males, all experienced computer users, holding a range of professional occupations and ranging in age from 21 to 55 years old. Participants were also balanced for telecommuting experience.

Method

Each male participant was shown all five video scenes where the telecommuter was male, while each female participant saw scenes using the female as the telecommuter. After completing the pre-test questionnaire, participants were given the telecommuting scenario. They were then asked to role-play, where they were first told they were at the office and that they would look at each scene in

turn in order to determine whether or not Larry/Linda was available for interaction.

- (1) Participants viewed one of the five video scenes at the first fully blurred level (e.g., Blur level 1 in Figure 2).
- (2) As they viewed each blur level, they answered awareness and privacy related questions (Figure 3, top).
- (3) They repeated steps 1 and 2 for the same scene at each of the remaining blur levels. This always progressed from fully blurred to completely unfiltered, and answers from previous blur levels remained visible so the participant could simply modify his or her answers.
- (4) They were then asked to imagine themselves as the telecommuter in each scene and were now themselves being watched by their colleague (Larry/Linda) at the office.
- (5) They chose a blur level for the scene and gave a reason (Figure 3, bottom). At this point, participants were able to view all blur levels at their discretion.
- (6) Upon completion of the first video scene, steps 1-5 were repeated for each of the remaining four video scenes.
- (7) After completing all five video scenes, they answered the post-test questionnaire, and performed the forced sort of all scenes.

The first scene shown to participants in Step 1 was always our most benign control scene containing the working telecommuter (Figure 1). We used this scene first to offset the chance that a participant may become “ultra-conservative” if they first saw a risky scene and thus rate later less risky scenes as being more threatening than normal. The viewing order for the remaining four scenes was randomized. Participants did not see video scenes where the telecommuter was of the opposite sex because it was felt that imagining yourself as the opposite sex for a portion of the questions may be quite difficult and could confound the results.

Results

Results are divided into four sections. First, we validate our original risk assessment of each scene, where we compare it with the results of the forced sort. Second, we address our three null hypotheses by analyzing our results. Finally, we determine people’s willingness to actually use blur filtration within a home media space, as captured on the post-test questionnaire.

We should mention that our original data analysis divided our participants into telecommuters/non-telecommuters. However, our analysis showed little difference between these two groups. For simplicity and clarity, we exclude this distinction in the following discussion and figures unless absolutely necessary.

Perceived Privacy Risk of Scenes

To discover how people perceived the privacy risk of each scene, we had them perform a forced sort of representative non-blurred pictures of each scene along a “line of privacy,” with one end indicating no risk and the other high risk.

There was reasonable consistency in participant responses: we saw only six distinct orderings, and even those did not differ much. Figure 4 shows these orderings: 6 of the 20 participants gave the first sequence, 5 the 2nd, 3 the 3rd and 4th, 2 the 5th, and 1 the 6th (these total numbers are further separated into male/female in the Figure).

All participants placed Working as least risky (column 1), while 18 of the 20 had Changing as the most risky scene (column 5). The two male dissenters felt Kissing was more risky than Changing. The major difference between the orderings is the placement of the middle three scenes. For 5 of the 6 orderings, No Shirt, Picking Nose, and Kissing were the middle three scenes, albeit in varying positions. We can ascribe part of this variation to gender differences, i.e., how males rated male actors with No Shirt vs. how females rated female actors with No Shirt.

As a whole, we feel that participants’ ordering of scenes confirms our original assessment of each scene’s risk: Working is low risk, Picking Nose, No Shirt, and Kissing are moderate risk, and Changing is high risk.


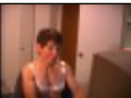





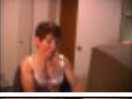












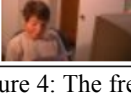

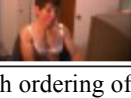

Ordering of Scenes (most frequent to least frequent)					Frequency		
					Male	Female	Total
1					5	1	6
2					1	4	5
3					1	2	3
4					1	2	3
5					2	0	2
6					0	1	1

Figure 4: The frequency of each ordering of scenes found in the forced sort by males and females. Male participants used the male equivalences of the scenes shown.



Figure 5: The mean placement of scenes according to risk, from low risk (0 cm) to high risk (300 cm), during the forced sort.

Figure 5, showing mean placement values, and an ANOVA suggest the scenes can be ranked into four categories of risk. First, almost all judged the Working scene (Figure 5, far left) as very low risk: it is close to the 25 cm rating, and the standard deviation is relatively small. The next category collects No Shirt and Picking Nose into a low-moderate risk rating ($p < 0.01$). Kissing has a somewhat greater moderate-high risk rating ($p < 0.01$). All judged Changing (far right) as very high risk ($p < 0.01$); images were positioned around 275 cm, with little deviation.

Determining Awareness

For each level of blur, participants were asked to write what they could see in the scene, rate how available the person was for conversation, and indicate the confidence they had in their guesses (Figure 3, Questions 1 and 2). We took this information and separated it into four awareness categories:

- (1) *activity*—the main activity found in the scene
- (2) *person*—who was in the scene
- (3) *appearance*—what the person(s) in the scene was wearing
- (4) *availability*—how available the telecommuter is for interaction right now

We now pose a series of questions to be answered by our observations.

At what blur levels did people correctly identify awareness cues? Cues from each of the four awareness categories (listed above) were generally identifiable between blur levels 3 and 5. Figure 6 plots the mean blur levels at which participants were first able to correctly identify categories of awareness cues for each scene. We judged correctness for the activity/person/appearance categories by verifying that the participants' descriptions matched what was actually happening in the scene, regardless of their confidence in their response. Because availability is a subjective measure, we judged an availability response as correct when the participant indicated they were quite confident (3 or greater) in their response (Figure 3, Question 2). Figure 6 suggests that participants were first able to identify the actor's activity, followed by the availability of the person in the scene, then the appearance of the person in the scene, and then who was actually in the scene.

Does scene type affect the blur level at which people begin to correctly extract awareness cues? No, the scene type did not affect the blur level at which people begin to correctly extract awareness cues. A two-factor ANOVA (scene type (5) \times awareness categories (4)) shows that there is no significant difference in the blur levels people used to first identify awareness information across the five scene types ($p = 0.06$). That is, people were able to identify the activity at a particular blur level regardless of the scene type. This is not unexpected, as each scene was created from the same fixed camera, and the information revealed (or hidden) is generally of a similar 'size' within the image.

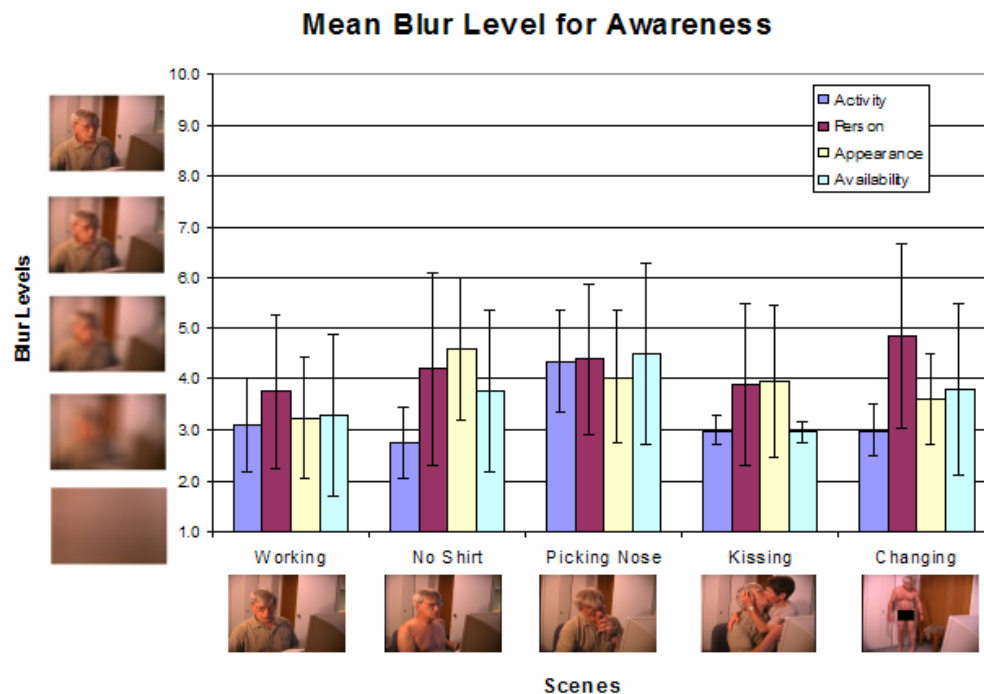


Figure 6: The mean blur level at which participants were first able to identify awareness cues for each scene.

Did people’s ability to correctly extract availability information from a blurred scene depend on the awareness category? The same ANOVA analysis suggests there is a significant difference in blur levels found for each of the awareness categories ($p < 0.05$). That is, only particular categories of information could be determined at particular blur levels. A post-hoc analysis shows that people determined activity at slightly blurrier scenes before they could determine either the person or their appearance ($p < 0.01$); no significant differences were found between the other awareness items. Because these actual differences are small, for practical purposes we can say that people’s ability to correctly determine awareness information is independent of the category of information they are looking for.

How confident were people in their awareness responses? Participants were not very confident in their initial answers (even when they were correct) and in most cases did not become confident until fidelity increased another 2 or 3 more levels. Figure 7 shows the confidence participants had in their ability to determine awareness cues. This mean represents the average confidence that participants had in identifying all awareness components: activity, person, appearance, and availability. We use this to represent the amount of awareness presented by each of the blur levels, as their confidence reflects their belief that they were correctly interpreting the scene. A two-factor ANOVA (scene types (5) \times blur levels (10)) confirms that for all scenes, there is a significant difference ($p < 0.05$) only in the amount of awareness presented by the blur levels.




Figure 7: The mean level of awareness confidence found at each blur level (1-low confidence to 5-high confidence).

In summary, all these results suggest *there is a difference in a viewer's ability to determine particular awareness cues over different blur levels across different scene types*. In particular, our results show that people begin perceiving all categories of awareness cues between blur levels 3 to 5, and that scene type does not make a difference to this result. Furthermore, we've shown that their confidence in their guesses increases with fidelity. With these results we reject our first null hypothesis.

Privacy Threat

For each level of blur, participants were asked to rate how threatening the scene was for the telecommuter and family members, given what they could currently see (Figure 3, Questions 3 and 4). In this section, we first describe the privacy threat to telecommuters, followed by the threat to family members.

Does the perceived threat to the privacy of the telecommuter differ by blur level? Yes, blur level affects the perceived privacy threat to the telecommuter. The mean privacy threat indicated at each blur level is shown in Figure 8. At levels 1 and 2, participants perceived little to no threat for all scenes. After this, the perceived threat increased with fidelity (correlations are all > 0.86). This increase occurs dramatically between blur levels 3 and 5 (the region indicated in the figure), and levels off by blur level 7. A two-factor ANOVA (scene type (5) \times blur levels (10)) verifies that the privacy threat between different blur levels does differ significantly ($p < 0.05$). Figure 8 clearly shows that these differences typically start between blur levels 2 and 3, and increase steadily until blur level 5. We see this result even for the Working scene (which remains mostly non-threatening), suggesting that people associated added threat with greater image fidelity, even in non-risky scenes.

Does the privacy threat to the telecommuter differ by scene? The same ANOVA also verifies that there is a significant difference in the threat for telecommuters between scenes ($p < 0.05$). A post-hoc analysis of overall mean privacy threat ($p < 0.01$) shows the scenes may be partitioned into three categories of threat. The low risk category consists of the Working scene. A moderate risk category includes the No Shirt and Picking Nose scenes. A high risk category holds the Changing and Kissing scenes.

What, if anything, made each scene threatening to the telecommuter? Participants usually associated threat with the person's particular activity or appearance. As fidelity increased these acts and their details became clearly visible and thus more threatening. Several participants also commented that they felt the scenes would be more threatening if they were viewing a colleague of the opposite sex.

Does blurring affect the privacy threat to family members? Despite the fact that a family member appeared in only one scene, participants still found the scenes to present some level of threat for family members. This threat is similar to that posed to the telecommuter: single factor ANOVAs ($p < 0.05$) performed on a scene-by-scene basis reveal no significant differences between the threat to family members and the threat to the telecommuter, except in the Picking Nose scene. There are two obvious distinctions, however: the mean threat at a given blur level is generally lower for family members than it is for the telecommuter, and the Kissing scene posed the highest risk to family members, while the Changing scene posed the highest risk to the telecommuter.

What, if anything, made each scene threatening to family members? Participants' responses were quite similar to those given for the telecommuter, i.e., threat was associated with the visibility of the person's risky activity or appearance. Curiously, people rated the Changing scene as very threatening to family members, even though no family member is ever present in the scene. The most common reason given by participants for this rating concerns the potential for threat: at any time a family member could walk into the room, and the fact that one wasn't there now was almost moot. This reason was given despite the fact that our question (which was accompanied by verbal explanation) specifically asked participants to rate the threat based on what could be seen currently, i.e., people had a tendency to infer what could happen even when instructed not to. A second, less common reason given was that participants felt that the family members may suffer the consequences—e.g., embarrassment or ridicule—should the telecommuter's reputation be affected by a privacy violation.




Figure 8: The level of privacy threat (1-low threat to 5-high threat) to the telecommuter at each blur level. The rectangle highlights blur levels 3 to 5, shown to provide awareness.

In summary, all these results suggest that *the perceived threat to the privacy of the telecommuter and family members differs between scenes and increases with fidelity*. We can reject the second null hypothesis. In particular, our results show that *only blur levels 1 and 2 made all scenes non-threatening*. The results also allow us to partition the scenes into three categories of risk: low (Working), moderate (Picking Nose and No Shirt), and high (Kissing and Changing).

Choosing Blur Levels

Participants were asked to imagine themselves as the telecommuter (i.e., Larry or Linda) and then, for each scene, choose a blur level (from 1 to 10) that they felt would make the scene appropriate for their colleague to view (Figure 3, Question 5). They also had the option to ‘turn the camera off,’ which we codified as a blur level of 0.

What blur levels did participants choose to make a scene appropriate for a colleague to view? The results vary with risk category (found in the previous privacy analysis) but do not differ in statistically significant ways by gender. Figure 9 plots our results, where the y-axis shows the mean selected blur levels chosen by participants for each scene. As one would expect, people chose more revealing blur levels for the low-risk Working scene (mean = 6.6, s.d. = 1.4) than for higher risk scenes, e.g., Changing (mean = 1.3, s.d. = 1.4). The results from a single-factor ANOVA looking for differences by scene ($p < 0.01$) show that the responses to this question partition the scenes into the same three risk categories we found in other analysis.




Figure 9: The mean blur levels chosen by participants for each scene. Blur level 0 represents choosing to turn the camera off.

We were curious as to if gender made a difference. A two-factor ANOVA (gender (2) \times scene type (5)) found that there is no statistically significant difference between the blur levels chosen for a particular scene by males vs. females ($p = 0.083$).

	Male (n=10)	Female (n=10)	All (n=20)
Working	0 %	10 %	5 %
No Shirt	0 %	40 %	20 %
Picking	0 %	30 %	15 %
Kissing	20 %	50 %	35 %
Changing	30 %	60 %	45 %

Table 1: The percent of participants who chose to turn off the camera.

When did people choose to turn off the camera? Nearly half of all participants chose to turn the camera off for the riskiest scene, yet only one turned it off for the least risky scene. That participant was adamantly opposed to using video at home and turned the camera off for every scene. Table 1 summarizes the proportion of participants who felt no blur levels were adequate for a scene and chose to turn the camera off (i.e., blur level 0) broken down by gender. For every scene, more female participants turned the camera off than male participants, and a two-factor ANOVA (gender (2) \times scene type (5)) showed the propensity to turn the camera off does in fact differ in a statistically significant way according to gender ($p < 0.05$). Ignoring this gender difference, we can see in the ‘All’ column of Table 1 that the five scenes break down into roughly the same three risk categories determined in other analyses.

In summary, these results show that *participants choose more distorted blur levels or more participants choose to turn the camera off altogether in order to make a video scene appropriate for a colleague to view as the risk to privacy posed by a scene increases*. That is, we can reject the third null hypothesis. Perhaps more importantly, we see that as the risk posed by a scene increases, more people abandon the blur filter in favor of turning the camera off altogether, and that nearly half of participants turn the camera off in order to make high risk video acceptable.

Willingness to Use Blurred/Unblurred Video

In the post-test questionnaire, we asked participants how willing they would be to use video in their own home to connect to a colleague they work closely with (1-unwilling to 5-willing). The mean willingness for all participants to use *unblurred* video was 1.9 (s.d.=1.0), while *blurred* video was 3.3 (s.d.=1.3). These values are significantly different ($p < 0.05$). We also checked to see if there were significant differences between male and female responses, but none were found.

Participants were also asked what they liked and disliked about using blurred video to balance privacy and awareness. Common *likes* included: being able to show availability while masking sensitive details, having the ability to control one’s privacy, and being able to easily stay in contact with others. Common

dislikes included: not being able to easily determine availability from blurred video, not knowing what the other person thinks they are seeing, and having to decide how much to blur and to alter this blur level for various scenes. Several participants said that they felt there was no balance between privacy and awareness: at the point where they could tell what was going on, they didn't feel the person's privacy was adequately being preserved. One participant also indicated a concern that blurred video could be unblurred by the viewer. As mentioned previously, one person was adamantly opposed to using video.

	Participants (n=20)
Office - yes	13
Office - no	7
Home - yes	9
Home - no	11

Table 2: Participants who would/ would not use blurred video in an office and at home.

When asked if they would—given the opportunity—actually use blurred video in an *office*, 65% of participants said they would (Table 2). Those who wouldn't use blurred video from an office said they preferred other means of gaining awareness, such as email, instant messaging, phone, or simply just walking over to see a person. They also commented that they felt their personal security would be violated when using blurred video, as the balance between privacy and awareness simply wasn't there.

Participants were then asked if they would—given the opportunity—actually use blurred video from *home*, 45% said they would (Table 2). Most of those who said they would use blurred video at home imposed caveats and restrictions: they wanted to choose the room where the camera was located; they wanted a mirror facility to know what was being captured; and, they wanted control over the blur level and whether or not the camera was on. Several also commented that they would simply leave the room to do private things that they would not want their colleagues to see. Those who said they would *not* use blurred video at home explained that they would find it intrusive, that it would violate their personal security, and that they felt blurring did not balance privacy and awareness. They also said that they saw the home as a place where they could go to achieve solitude from their colleagues. They felt that conventional mechanisms—e.g., email, instant messaging, or phone—are adequate means for gaining awareness. One participant said that she would be fine with using blurred video at home, but didn't feel her husband would want it.

Discussion

This paper set out to evaluate blur filtration for its effectiveness in balancing privacy and awareness for video-based telecommuting situations. In particular, we wanted to know whether or not blur levels existed that could provide adequate

awareness, while still preserving privacy. Aggregating the results across all scenes tested, we found that awareness cues were first identifiable between blur levels 3 and 5, while privacy is only preserved for blur levels 1 and 2. It is clear that these blur levels do not overlap; thus, *there are no general-purpose blur levels which can balance privacy and awareness in any scene.*

If we analyze this on a scene-by-scene basis, we see that the Working scene, representing a mundane home situation, is the only scene where privacy preserving levels overlap the awareness range. Thus, we can see that blur filtration is only able to balance privacy and awareness for mundane home situations. Our results also show that participants prefer to use blur levels that preserve as much privacy as possible, while still allowing awareness. This suggests that a balance for mundane home situations would occur between blur levels 3 and 5.

Conclusion

The study's results confirm those found in Boyle et al's (2000) study. For benign situations like working, there are blur levels that are able to balance privacy and awareness. However, blur filtration is not capable of balancing privacy and awareness for all home-based telecommuting situations; other privacy-protecting strategies and technologies are required.

The results of our study have important design implications for a home media space. First and foremost, it is clear that a home media space is not suitable for everyone; media space participants must possess a strong desire to be a part of the space. For those who choose to participate, a high degree of control over privacy is desired, e.g., controlling where the camera is located, how much filtration is used, and when the camera is turned on/off. Along with this comes a need for feedback of how much privacy is being attained. This allows people to make informed decisions about their privacy. It is important to consider that in most homes multiple people exist and privacy expectations may vary between them. Home media space designs must address the privacy concerns of both the telecommuter and others in the home. A final consideration is the effect of combining two separate cultures—an office culture and a home culture—into a single space where real-world assumptions related to privacy need not always hold true. A successful home media space design must cope with varying privacy expectations.

Acknowledgments

We are grateful to NSERC and Microsoft Research for their partial funding of this research. We also thank the two actors who performed in our video scenes.

References

- Bellotti, V. (1998) 'Design for Privacy in Multimedia Computing and Communications Environments, in Technology and Privacy: The New Landscape', Agre and Rotenberg eds., MIT Press, pp. 63-98.
- Bellotti, V. (1996), 'What you don't know can hurt you: Privacy in Collaborative Computing', *Procs of HCI '96*, Springer, pp. 241-261.
- Bellotti, V., and Sellen, A. (1993) 'Design for Privacy in Ubiquitous Computing Environments', *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93)*, Kluwer Academic Publishers, Milan, pp. 77-92.
- Bly, S., Harrison, S. and Irvin, S. (1993) 'Media spaces: Bringing people together in a video, audio, and computing environment.' *Communications of the ACM* 36(1), pp. 28-46, ACM Press.
- Boyle, M., Edwards, C. and Greenberg, S. (2000) 'The Effects of Filtered Video on Awareness and Privacy', *Proceedings of the CSCW'00 Conference on Computer Supported Cooperative Work* [CHI Letters 2(3)], ACM Press.
- Fish, R.S., Kraut, R.E., Rice, R.E., and Root, R.W. (1993) 'Video as a Technology for Informal Communication', in *Communications of the ACM*, vol. 36, no. 1, pp. 48-61, ACM Press.
- Greenberg, Saul. (1996), 'Peepholes: Low Cost Awareness of One's Community', *ACM SIGCHI'96 Conference on Human Factors in Computing System*, Companion Proceedings, pp. 206-207.
- Greenberg, S. and Kuzuoka, H. (2000) 'Using Digital but Physical Surrogates to Mediate Awareness', *Communication and Privacy in Media Spaces*. *Personal Technologies*, 4(1), January.
- Hudson, S.E., and Smith, I. (1996) 'Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems', *Proc of the Conference on Computer Supported Cooperative Work (CSCW'96)*, Cambridge, MA.
- Kraut, R., Egido, C., and Galegher, J. (1988) 'Patterns of contact and communication in scientific observation', *Proc ACM CSCW '88*, pp. 1-12.
- Lee, A., Girgensohn, A., Schlueter, K. (1997) 'NYNEX Portholes: Initial User Reactions and Redesign Implications', *Group '97*, pp. 385-394, ACM Press.
- Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T., and Wellman, B. (1991) 'Experiences in the use of a media space', *Proc. of CHI '91 Human Factors in Computing Systems*, pp. 203-209, ACM Press.
- Tang, J.C., Isaacs, E., and Rua, M. (1994) 'Supporting Distributed Groups with a Montage of Lightweight Interactions', *Proc. of the ACM Conference on Computer-Supported Cooperative Work, CSCW '94*, pp. 23-34, ACM Press.
- Zhao, Q.A., and Stasko, J.T. (1998) 'Evaluating Image Filtering Based Techniques in Media Space Applications', *Proc of the Conference on Computer Supported Cooperative Work (CSCW'98)*, pp. 11-18, ACM Press.