# The Age of Avatar Realism

## *When Seeing Shouldn't be Believing*

**BY LAUREL D. RIEK AND ROBERT N.M. WATSON**

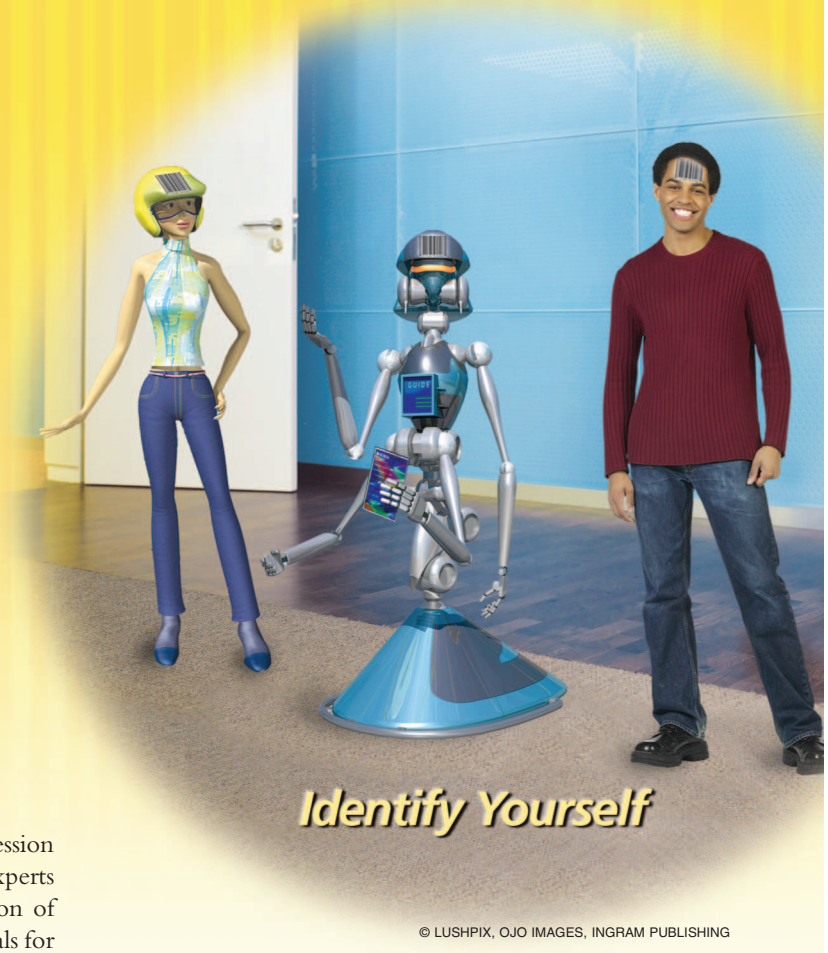




Recent advances in affect recognition, facial expression synthesis, and speech synthesis are giving nonexperts access to tools for performing live manipulation of audio and video streams, as well as control signals for telepresence robots. We apply the vocabulary of security protocol analysis to this new environment, considering how changes in technology call into question assumptions of trust that seem natural to users of video teleconferencing and robot telepresence.

In the current economic climate, many organizations are turning to telepresence to reduce travel costs, while maintaining their global presence. In a recent report released by market research firm ABI, telepresence software, hardware, and services reached US$567 million and is expected to become a US$2.7 billion industry by 2015 [1]. These figures only reflect point-to-point video conferencing, but several new advances in robots that support telepresence may also be entering the market soon. Such robots include mobile robots with liquid-crystal display (LCD) screens, depicting the remote collaborator [2]–[4] as well as teleoperated, lifelike android systems [5], [6].

Users of such photorealistic communication systems reasonably expect that their remote representation is faithful and accurate. However, like any other communication system, telepresence systems are subject to security vulnerabilities. Both platform end-hosts (such as the robot or video phone) as well as the communication channel may be vulnerable to attack. Such attacks might occur because of supply chain vulnerabilities, nonexistent or poor cryptography, stolen keys, or compromised infrastructure.

Once the platform or channel is compromised, an attacker can control what messages are sent or received. In the terminology of protocol attacks, we are concerned with message integrity, confidentiality, and availability; attackers may modify messages, improperly obtain their contents, or prevent the system from operating at all. However, in the context of video communications and telepresence, these abstract concerns take on new and subtle implications: modifications to the channel may not be immediately (or at all) obvious to the end user, as recent improvements in technology allow the realistic modification of both verbal and nonverbal communication signals in real time. This may allow the malicious modification of communication, or even the complete impersonation of a participant.

At the more overt end of the spectrum, an adversary (Eve) may spoof communications entirely, assuming an otherwise legitimate party's identity (Alice) for the intent of a conversation. For people with a close relationship to Alice, this kind of identity theft may be tricky to maintain because of the contextual information and subtle differences in her behavior, in particular, nonverbal behaviors that violate their expectations [7]. However, for people who are occasional acquaintances or strangers to Alice, this type of attack could pose a major problem.

More subtly, Eve might inject facial expressions and gestures into a conversation between Alice and another party (Bob), such as furrowing Alice's eyebrows to make her look angry. Or, Eve may choose to inhibit Alice's expressions, such as reducing the intensity of her smiles. Eve can also augment or inhibit Alice's tone of voice, or indeed even the

*Digital Object Identifier 10.1109/MRA.2010.938841*


1070-9932/10/$26.00©2010 IEEE

***In the current economic climate, many organizations are turning to telepresence to reduce travel costs, while maintaining their global presence.***

words she says. In the field of computer security, attacks such as these are referred to as a man-in-the-middle attack, in which a third party interferes with the expected execution of a protocol.

The technology to facilitate such alterations of expressions was once only in the hands of experts, such as film studios, and required weeks or months of compute time to produce a realistic result. However, with recent advances in affect recognition [8], facial expression synthesis [9], and speech synthesis [10], nonexperts are becoming increasingly enabled to carry out such attacks. As the technology becomes more ubiquitous, this problem will only get worse.

In this article, we will describe the state of the art in terms of affect and speech synthesis systems, explore the range of attacks that telepresence systems pose, and suggest areas of future research that may help defend against such threats.

## Background

Using technology to make people appear to say things they did not is by no means a new idea; Bregler et al. [11] introduced Video Rewrite in 1997. This system interjects novel speech and accompanying visemes (facial and oral movements that accompany voiced phonemes) into a video. At that time, this technique only required the manual labeling of 26 frames. Today neither video nor manual labeling is needed, and far more than just mouth movements can be readily synthesized. In fact, SitePal recently launched three-dimensional (3-D) photoface [12], which allows one to create a talking, moving, and realistic 3-D facial avatar from a single still photograph. The avatar can be easily manipulated to make expressions, move, and generate plausible visemes to either prerecorded or synthesized speech (see Figure 1).

***Figure 1.*** *An original image of Jimmy Wales and three fake facial expressions made using SitePal 3-D photoface. A believable, fully animate avatar can be created on this site from just a still-image. (Original photo taken by Manuel Archain, Creative Commons ShareAlike © 2008 by Jimmy Wales.)*

For virtual avatar communication, Bailenson and Blascovich [13] introduced the possibility of people systematically self-filtering their behaviors or their appearance by suppressing or amplifying various communicative signals. Indeed, they suggest that this raises ethical concerns regarding such "transformed social interactions," in that it may become impossible to detect how accurate and faithful an avatar representation is, thus impeding people's ability to assess communication honesty.

In terms of altering telepresent communication in real time, Boker et al. [14], [15] developed a system that permits the real-time manipulation of the face and voice of a realistic, resynthesized avatar. The avatars were resynthesized using active appearance models (AAMs) [16], which are statistically based techniques that describe shape deformations of the face. While AAMs traditionally require many manually labeled training images, recent advances in the technique require hardly any manual labeling at all [17], [18].

In the experiments, Boker et al. describe naive subjects who participated in six conversations over video teleconference. Three of these conversations were with a female confederate, and three were with a male confederate. However, as far as the naive participants were concerned, they had conversations with six different people, for they saw six completely different avatars and heard six different voices. The female confederate was portrayed on occasion as a male avatar and the male confederate as a female avatar. None of the subjects reported they were speaking with fewer than six people, and none realized those people were represented by a computer animation. Furthermore, in a follow-up experiment, the researchers found that 91.9% of the time subjects rated female avatars as female even when the confederate was male and male avatars as male even when the confederate was female [15].

In the area of speech synthesis, a number of significant advances have recently taken place in terms of increased naturalness, better prosody support, and reduced training data requirements. Taylor [10] describes such advances in detail, but one particularly noteworthy recent idea is personalized text-to-speech engines. CereProc [19] specializes in this kind of synthesis, and with only 4 h of audio can recreate someone's voice, as they recently did for Roger Ebert [20].

In robotics, a range of systems exist that can support remote telepresence and the transmission of a wide range of human communicative cues. On one end are androids that appear and behave nearly exactly as their remote operator does and convey facial expressions, head gestures, and speech with high fidelity (e.g., the Geminoid [5]). On the other end of the spectrum are robots that look nothing like their remote operators and convey only limited cues, such as gaze direction and speech (e.g., the WowWee Rovio [21]). Based on previous work on people's expectations of a robot's capabilities from its appearance [22], we expect a similar trend applies to telepresent robots in terms of their ability to transmit nonverbal cues.

Thus, as a robot becomes more realistic (thus resembling the remote collaborator), many people's expectations will be that it is faithfully transmitting the nonverbal cues of its operator (see Figure 2).

Because people judge the authenticity of others based on how they look and sound, these new advances that support real-time manipulation of human communication cues make telepresence increasingly vulnerable to malicious attack.

## Protocol Attacks in a Telepresent World

Since the 1970s, security research literature has considered both the construction and analysis of security protocols; however, despite extensive discussion, security protocols defy simple definition. Many security protocols are cryptographic in nature, and cryptography provides the basis for reasoning about their properties. Such protocols involve two or more parties who wish to reliably draw conclusions about the success of a protocol, perhaps for the purposes of mutual authentication, safe communication of confidential information, or distributed decision-making. Common primitives for constructing protocols include digital signatures, cryptographic hashes, and encryption/decryption using shared secrets or public/private key.

Protocol literature describes not only a large number of protocols but also a large number of attacks against those protocols. Vulnerabilities may arise not only because of the incorrect use of cryptographic primitives or poor reasoning by a protocol's designer but also as a result of poor understanding of the properties and assumptions of a protocol by its users. These vulnerabilities may allow an attacker to violate the integrity or confidentiality guarantees of a protocol, or prevent the protocol (or services built on the protocol) from operating, and hence deny availability. Protocol attacks target the integrity, confidentiality, and availability of communications and may employ techniques such as eavesdropping, injecting messages, modifying messages, and replaying messages. Detailed descriptions of cryptographic primitives, protocols, and attacks can be found in Anderson [25] and Ferguson et al. [26].

However, security protocols can also be social in nature, leading in turn to a set of attacks known as social engineering. For the purposes of these protocols, end points are not computers on networks passing messages but people participating in the social protocols of organizations and society. Attacks on social protocols are often much more straightforward than cryptographic attacks—it is easier to trick a human into revealing his or her password than to break even off-the-shelf cryptography protecting it! A detailed consideration of social engineering, as well as its combination with technical attacks on protocols, is provided by Mitnick and Simon [27].

Established protocol security vocabulary for both cryptographic protocol and social engineering attacks lends itself well to describing robot telepresence and video teleconference stream attacks. In the telepresence scenario, as with social engineering attacks, we consider both people and computer systems to be participants in the protocol and reason about the conclusions reached by the participating humans as critical protocol results. To this end, telepresence system attacks may

***In robotics, a range of systems exist that can support remote telepresence and the transmission of a wide range of human communicative cues.***

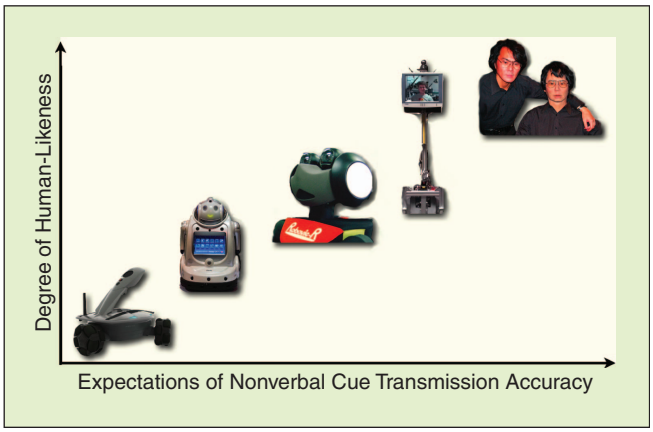


***Figure 2.*** *As the level of human-likeness increases, people's expectations of the nonverbal cue transmission accuracy will also likely increase. Robots from left to right: WowWee Rovio [21], Yujin Robot iRobi [23], Robovie [24], Willow Garage Texas Robot [3], and Geminoid H1-1 [5]. (Courtesy of Wikimedia.)*

be considered another tool in the tool chest of the social engineer, blending social and technical elements.

For the purposes of this discussion, and adopting the parlance of security protocol literature, we consider the interactions of three actors: Alice and Bob, well-meaning participants in security protocols, and the malicious Eve, who wishes to circumvent the protocol by exploiting vulnerabilities. In security protocols, cryptographic techniques are combined with communications to allow participants in protocols to exchange information securely or provably accomplish common objectives. Typical attacks on security protocols exploit incorrect assumptions about the elements or results of a protocol; for example, they may rely on an unsafe use of a cryptographic primitive, or a failure to include fresh information in the exchange to prevent replay attacks. In the context of telepresent communications, the realism of the communication channel and the use of social techniques substitute for attacks on cryptographic primitives and incorrect assumptions of the participants.

In a two-party scenario, Alice engages in communication with Eve, who wishes to exploit assumptions about the protocol (and implementation) to mislead Alice. In a three-party scenario, or man-in-the-middle attack, legitimate participants Alice and Bob may likewise be malignly influenced by Eve such that conclusions they draw from the protocol are misplaced. A robot telepresence or video teleconference attacker is able to manipulate the streams of communication and control to mislead legitimate protocol participants about the

***Common primitives for constructing protocols include digital signatures, cryptographic hashes, and encryption/decryption using shared secrets or public/private key.***

conclusions they can draw from the protocol. With realistic manipulation or even substitution of realistic human communicative cues over a compromised channel, this characterization proves similarly useful. We consider three scenarios to explore the realm of possibility, each of which uses the manipulation of the communications channel to reinforce a social engineering attack.

### Angry Alice

In our first example (Figure 3), we consider a man-in-the-middle attack in which either Alice or Bob's video conferencing system has been compromised, allowing (subject to real-time constraints) the video and audio streams to be modified. Alice works for a large software company and is evaluating network hardware from several competing vendors; Bob is a sales representative from a vendor. Eve is an employee of a competing firm, who wishes to maliciously interfere with the completion of a protocol (negotiation) between Alice's and Bob's companies.

Bob and Alice have similar goals: explore each other's expectations of cost but more generally build the trust required to complete the business negotiation. In person, by video phone, or via telepresence, various cues allow this negotiation to proceed—sales representatives and buyers alike are experienced in avoiding missteps that may make reaching an agreement more difficult (or impossible). In itself, the use of computer-mediated communication makes negotiation more challenging because of the reduction in nonverbal cues [28], [29]. However, if Eve is able to manipulate the channel to further reduce Alice or Bob's effectiveness at reading the cues of the other, then an agreement may not be reached.

For example, Eve may manipulate the video and/or audio stream to make Alice appear angry (for instance by frequently furrowing her eyebrows or raising the pitch of her voice [30]). This in turn may make Bob feel less comfortable and thus less inclined toward reaching an agreement with her.

This attack could be carried out on a larger scale by compromising the supply chain. Perhaps, Eve's employer is a competing network vendor to Bob and also happens to supply his company's video conference software. Eve's employer could ensure favorable outcomes by modifying all communications between Alice's and Bob's companies, causing them to lose trust in one another, thus propelling Alice's company to purchase hardware from Eve's employer instead.

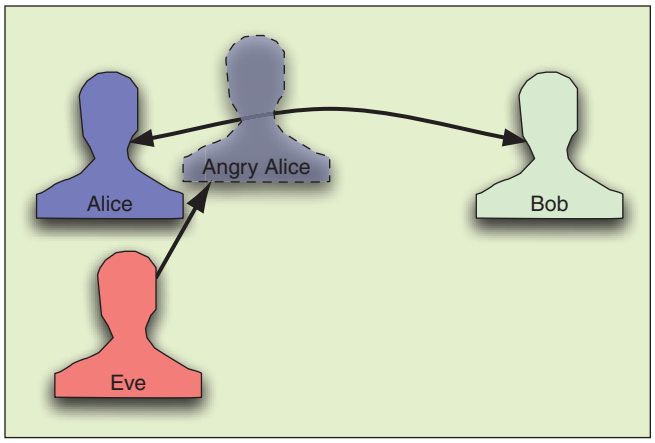


***Figure 3.*** *Eve modifies the videoconference stream between Alice and Bob by injecting new facial expressions leading Bob to conclude Alice is displeased.*

### Alice's Physical Avatar

Physical avatars, or teleoperated robots that serve as physical manifestations of a remote user [31], present several unusual opportunities of attack beyond those of static video teleconferencing systems. In our second example (Figure 4), we consider how Eve might co-opt Alice's physical avatar, while she is remotely delivering a courtroom appearance.

The courtroom environment is of particular concern: telepresence and recorded testimony have raised significant questions in jurisprudence [32] but are becoming accepted practice. The ability to interfere with, and worse, and to inject credible modifications (or even complete substitution) of court testimony threatens both the integrity and availability of legal processes.

Alice is testifying in court via a mobile LCD screen (such as via a Texas Robot [3]). We assume that this robot conveys a realistic representation of her face and voice. Furthermore, Alice can teleoperate the robot and its camera with a joystick to convey gaze and attentional posture cues.

After allowing Alice to establish credibility as a witness and answer a series of question, Eve may hijack Alice's control of the robot and then use a realistic reproduction of Alice's voice and appearance to inject malicious content. If done well, the attack would leave neither Alice nor the jury any wiser—they would have been misled as to the testimony (attack on integrity). Even if the attack is detected, the presentation of compromised evidence to a jury might well lead to mistrial (attack on availability). This attack is reminiscent of attacks against message authentication codes (MACs) such as naive use of

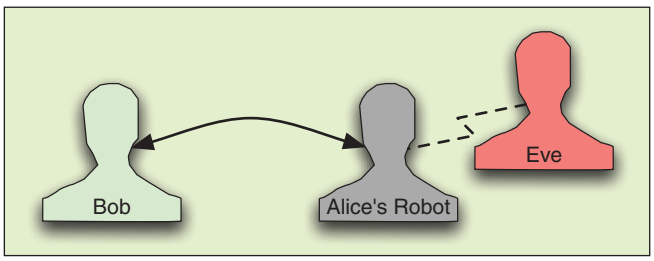


***Figure 4.*** *Bob is unaware that Eve has compromised Alice's robot and therefore gives Eve all the rights and services he would normally provide to Alice.*

keyed MD5; an attacker is able to append new data to the end of messages and have them accepted by the receiving end [33].

### *Alice's Neighbor*

In our third example (Figure 5), we consider a two-party protocol, in which Eve has built an accurate video and audio models of Bob such that a short conversation can be held in real time with reasonable verisimilitude. The malicious Eve uses Skype to place a call from her hotel room to Bob's neighbor, Alice. Using Bob's face and voice, she tells Alice that she (apparently Bob) is away but needs Alice to allow a plumber access to the house via the spare house key. Alice, who knows Bob albeit perhaps not all that well, may reasonably accept these instructions and allow the supposed plumber into the house.

Of particular interest in this scenario is that Eve does not even need to steal keying material or compromise an end-host: participants in many human-to-human protocols rely implicitly on their ability to identify those with whom they are communicating. If the other participant looks, sounds, and acts like a given person, they are that person for all intents and purposes, even if talking on someone else's phone.

## Mitigation

While the techniques we describe in the article remain on the edge of the possible today, improvements in associated technologies are rapid and concerning; what is currently in the realm of high-end cinema studios will rapidly become the domain of end-user PCs. Proper application of traditional platform and channel security approaches to video teleconferencing and robot telepresence systems is clearly critical in preventing some of these attacks. However, end users relying on these technologies must be aware that, in the presence of an exploited vulnerability, a new suite of tools is available to an attacker. When combined with social engineering, the results could be both surprising and devastating.

The security research and engineering communities continue to explore a range of approaches for dealing with security problems, not least defense in depth, in which the assumption of vulnerability motivates the adoption of a multilayered approach and the extensive use of mitigation technologies. Similarly, we wonder whether there aren't mitigation approaches to be found specifically for the new problems we describe. Intrusion detection technology has seen mixed reviews because of its problems with false positives and negatives, as well as the inherent difficulties with anomaly detection [34], but it is easy to imagine similar approaches applied here. Manipulation of communication streams may leave identifiable traces, both at low levels (perhaps identifiable via video recompression artifacts) and at higher levels (biometric techniques could be applied to identify discontinuities or inconsistencies with a model).

Certainly, these are reasonable metrics with which to consider the effectiveness of attacks, and they bear further exploration—especially as biometric techniques are, themselves, frequently used for authentication and are likewise potentially vulnerable. Of course, a significant challenge in the arena of detection is legitimate modification that may also take place in the future; the makeup of the future may remove undesired facial expressions or ticks, even out voice tone, and more, using the same technologies we describe as a potential threat.

***Increasingly, seeing should not imply believing in the world of telepresent communication.***

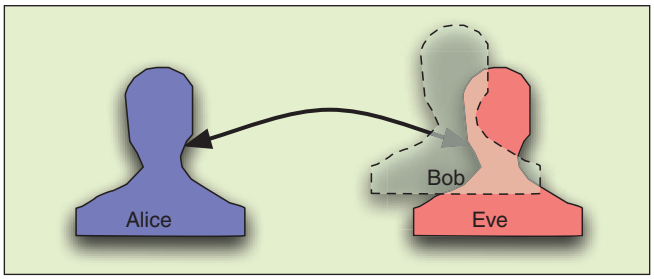


***Figure 5.*** *Eve impersonates Bob via videoconference to convince Alice, Bob's neighbor, to let the plumber (Eve) into his house.*

## Conclusion

We have described computer techniques for characterizing and reproducing realistic behaviors and appearances for individuals that are becoming state of the art, with a potentially serious (and largely unconsidered) impact on security. These include recent advances in resynthesized avatars [15], personalized text-to-speech engines [19], and still-photo to animate avatar systems such as SitePal 3-D photoface [12].

Increasingly, seeing should not imply believing in the world of telepresent communication; strong measures must be used to protect the security of the channels, or there is a risk of discrediting the technologies themselves as reliable and safe to use. While we discuss potential mitigation techniques, it seems that the only sure-fire approach is to redouble our efforts in traditional security areas (platform and channel security) and to ensure that these techniques are used wherever they are appropriate.

## Acknowledgments

Riek and Watson thank Jon Anderson, Shazia Afzal, Cynthia Breazeal, Markus Kuhn, Steven Murdoch, Peter Robinson, Ross Anderson, and our anonymous reviewers for their ideas and suggestions. Riek gratefully acknowledges the support of the Qualcomm Research Studentship and Watson the support of Google, Inc.

## Keywords

Gesture, posture, social spaces, facial expressions, security protocols, biometrics, telerobotics, virtual reality, interfaces.

## References

[1] ABIResearch. (2010). Telepresence at an inflection point; Sales projected to reach $2.7 billion by 2015 [Online]. Available: http://www.abiresearch.com/press/1636

[2] C. Gaylord, "Send your robot to work," *Christian Science Monitor*, May 2008.
[3] WillowGarage. (2010). Texai overview [Online]. Available: http://www.willowgarage.com/pages/texai/overview
[4] HeadThere. Headthere: Maker of the Giraffe video conferencing robot [Online]. Available: www.headthere.com/
[5] D. Sakamoto, T. Kanda, T. Ono, H. Ishiguro, and N. Hagita, "Android as a telecommunication medium with a human-like presence," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction (HRI'07)*, New York, NY, ACM, 2007, pp. 193–200.
[6] C. Dillow, "Japanese Geminoid F Bot realistically mimics human facial expressions, speech," *Popular Science*, Apr. 2010.
[7] J. Burgoon and J. Hale, "Nonverbal expectancy violations: Model elaboration and application to immediacy behaviors," *Commun. Monographs*, vol. 55, no. 1, pp. 58–79, 1988.
[8] Z. Zeng, M. Pantic, G. Roisman, and T. Huang, "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 31, no. 1, pp. 39–58, 2009.
[9] A. Raouzaiou, N. Tsapatsoulis, K. Karpouzis, and S. Kollias, "Parameterized facial expression synthesis based on MPEG-4," *EURASIP J. Appl. Signal Process.*, vol. 2002, no. 1, pp. 1021–1038, 2002.
[10] P. Taylor, *Text-to-Speech Synthesis*. Cambridge: U.K.: Cambridge Univ. Press, 2009.
[11] C. Bregler, M. Covell, and M. Slaney, "Video rewrite: Driving visual speech with audio," in *Proc. 24th Annu. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH'97)*, ACM Press/Addison-Wesley, 1997, pp. 353–360.
[12] SitePal. (2009). 3D photoface [Online]. Available: http://www.sitepal.com/photoface_pop
[13] J. Bailenson and J. Blascovich, *Avatars*. Great Barrington, MA: Berkshire Publishing Group, 2004, pp. 64–68.
[14] S. Boker, J. Cohn, B. Theobald, I. Matthews, T. Brick, and J. Spies, "Effects of damping head movement and facial expression in dyadic conversation using real–time facial expression tracking and synthesized avatars," *Philos. Trans. R. Soc. B: Biol. Sci.*, vol. 364, no. 1535, p. 3485, 2009.
[15] S. Boker, J. Cohn, B. Theobald, I. Matthews, M. Mangini, J. Spies, Z. Ambadar, and T. Brick, "Something in the way we move: Motion dynamics, not perceived sex, influence head movements in conversation," *J. Exp. Psychol.: Hum. Perception Perform.*, to be published.
[16] T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," in *Computer Vision—ECCV'98* (Lecture Notes in Computer Science, vol. 1407), H. Burkhardt and B. Neumann, Eds. New York: Springer-Verlag, 1998, p. 484.
[17] A. Asthana, A. Khwaja, and R. Goecke, "Automatic frontal face annotation and AAM building for arbitrary expressions from a single frontal image only," in *Proc. IEEE Int. Conf. Image Processing*, 2009, pp. 2445–2448.
[18] A. Asthana, J. M. Saragih, M. Wagner, and R. Goecke, "Evaluating AAM fitting methods for facial expression recognition," in *Proc. IEEE Int. Conf. Affective Computing and Intelligent Interaction (ACII'09)*, 2009.
[19] CereProc. Voice creation FAQs [Online]. Available: http://www.cereproc.com/support/faqs/voicecreation
[20] R. Ebert, "Hello, this is me speaking," *Chicago Sun Times*, Feb. 2010.
[21] WowWee. (2010). WowWee: Rovio [Online]. Available: http://www.wowwee.com/en/products/tech/telepresence/rovio/rovio
[22] J. Goetz, S. Kiesler, and A. Powers, "Matching robot appearance and behavior to tasks to improve human-robot cooperation," in *Proc. IEEE 12th IEEE Workshop Robot and Human Interactive Communication (RO-MAN)*, 2003, pp. 55–60.
[23] iRobi. (2010). Yujinrobot: iRobi [Online]. Available: http://www.yujinrobot.com/english/product/irobi.php
[24] H. Ishiguro, T. Ono, M. Imai, T. Maeda, T. Kanda, and R. Nakatsu, "Robovie: An interactive humanoid robot," *Ind. Robot: An Int. J.*, vol. 28, no. 6, pp. 498–504, 2001.
[25] R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*. New York: Wiley, 2008.
[26] N. Ferguson, B. Schneier, and T. Kohno, *Cryptography Engineering: Design Principles and Practical Applications*. New York: Wiley, 2010.
[27] K. Mitnick and W. Simon, *The Art of Deception: Controlling the Human Element of Security*. New York: Wiley, 2003.
[28] L. Thompson and J. Nadler, "Negotiating via information technology: Theory and application," *J. Social Issues*, vol. 58, no. 1, pp. 109–124, 2002.
[29] J. C. Tang, "Approaching and leave-taking: Negotiating contact in computer-mediated communication," *ACM Trans. Comput.-Hum. Interact.*, vol. 14, no. 1, p. 5, 2007.
[30] G. Clore and J. Palmer, "Affective guidance of intelligent agents: How emotion controls cognition," *Cogn. Syst. Res.*, vol. 10, no. 1, pp. 21–30, 2009.
[31] L. D. Riek, "Realizing Hinokio: Candidate requirements for physical avatar systems," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction (HRI'07)*, New York, NY, ACM, 2007, pp. 303–308.
[32] Z. Hillman, "Pleading guilty and video teleconference: Is a defendant constitutionally 'present' when pleading guilty by video teleconference?," *J. High Technol. Law*, vol. 7, pp. 41–226, Jan. 2007.
[33] M. Bellare, R. Canetti, and H. Krawczyk, "Keying hash functions for message authentication," in *Proc. Advances in Cryptology (CRYPTO'96)*, Springer, 1996, pp. 1–15.
[34] V. Paxson, "Considerations and pitfalls for conducting intrusion detection research," Int. Comput. Sci. Inst., Lawrence Berkeley Nat. Lab., Berkeley, CA, Tech. Rep. 2007.

**Laurel D. Riek** received her B.Sc. degree in logic and computation from Carnegie Mellon University in 2000. She is a Ph.D. candidate in the Computer Laboratory at the University of Cambridge. She researches natural human–robot interaction, in particular, facilitating nonverbal communication with robots. Her research explores expression synthesis on android and humanoid robots using naturally evoked human data. She also explores sustainable interaction with robots by applying social signal processing techniques to the analysis of dyadic human conversations. Prior to starting her Ph.D., she worked for eight years as a senior artificial intelligence engineer and roboticist at MITRE, on projects involving search and rescue robots, unmanned vehicles, and human language technology.

**Robert N.M. Watson** received his B.Sc. degree in logic and computation from Carnegie Mellon University in 1999. He is a Ph.D. candidate in the Computer Laboratory at the University of Cambridge. His research is in the area of concurrency and security. Prior to joining the Computer Laboratory, he was a senior principal scientist at McAfee Research, now SPARTA Information Systems Security Operation (ISSO), where he directed commercial and government research and development projects in computer security, including the TrustedBSD MAC Framework used for access control in FreeBSD, Mac OS X, and iPhoneOS. His research interests include operating system security, network stack structure and performance, and windowing system structure. He is also a member of the board of directors for the FreeBSD Foundation, a 501c3 nonprofit supporting development of FreeBSD, a widely used open source operating system.

***Address for Correspondence:*** Laurel D. Riek, Computer Laboratory, University of Cambridge, Cambridge, United Kingdom. E-mail: Laurel.Riek@cl.cam.ac.uk.