

Rich Nonverbal Sensing Technology for Automated Social Skills Training

Mohammed (Ehsan) Hoque,
ROC HCI, University of Rochester

Rosalind W. Picard, *MIT Media Lab*

Automated nonverbal sensing and feedback technologies, such as My Automated Conversation coach (MACH), can provide a personalized means to better understand, evaluate, and improve human social interaction—for both practical and therapeutic purposes, and to advance future communications research.

Lewis, a technically gifted college junior applying for an internship, worries about his interview skills. A 15-minute session with a career counselor yields recommendations to maintain more eye contact, vary vocal pitch and volume to express enthusiasm, and conclude the interview with a friendly smile. Lewis would like more opportunity to practice and receive feedback—but the counselor's availability is limited, and he's embarrassed to ask peers for help.

Kal, a teen diagnosed with social anxiety disorder, finds face-to-face communication extremely challenging. His awkwardness with conversational “rules” involving turn-taking, appropriate facial expressions and eye contact, varying intonation, and the like have often been met with taunting, bullying, and social rejection. Consequently, Kal has retreated into a virtual world, comfortable conversing with new friends only online. Still, he craves real-world

contact with others and wishes for a way to learn the nuances of physical social interaction while remaining in the virtual environment where he's currently most comfortable.

Gretchen, a graduate student excited by the idea of sharing work she's proud of at scientific conferences, dreads public speaking. Despite working with her research group to improve her written presentations and incorporating feedback from peers and advisors, at formal meetings she ends up talking too quickly for audiences to comprehend her ideas, and her body language fails to project her enthusiasm. Gretchen would love to improve her persona as a public speaker, but lacks the time or money to take a professional course.

Marti, a new sales representative who videoconferences with clients daily, has failed to achieve her quotas compared to her colleagues, even though they all use the same basic script. She is eager for personalized feedback to help her understand objectively where her persuasive communication falls short, but she fears that direct comparison with her colleagues might jeopardize her job.

In each of these cases, we find someone seeking one-on-one guidance to improve his or her nonverbal communication skills, stymied by logistics, lack of resources, or personal embarrassment. Is it possible that a technology could help individuals like these—an automated system available ubiquitously, wherever and whenever they want?

AUTOMATED, REAL-TIME SOCIAL SKILLS TRAINING

Here, we share our vision of just such an automated sensing technology, one that allows users to practice multiple

social scenarios electronically and obtain immediate, objective feedback—a system for which Figure 1 represents an ideal. Its potential uses include not only improving job interviews, social anxiety, public speaking, and sales presentations, but also new language acquisition, cross-cultural communication, customer support, doctor–patient understanding, even flirting and dating skills. For any situation in which face-to-face human communication occurs, whether in person or online, capturing the richness of real-time interaction can provide a personalized means to better understand and improve how it unfolds.

Specifically, we address three fundamental questions:

- How might computer technology help people practice successful social interaction skills and receive useful feedback?
- How could an automated system sense, interpret, and represent rich conversational multimodal behavioral data in a format that is both intuitive and educational?
- Is it possible for such technology to elicit measurable improvements in users' social skills?

In considering these, we pose a further question: What future applications might we envision for automated nonverbal sensing technology?

CONTEXT, FUNDAMENTALS, AND CHALLENGES

Designing an automated system for social skills training unites several disparate bodies of knowledge, including research in social and clinical psychology, communication theory, affective computing, smart sensing systems, and information visualization.

Nonverbal behavior in effective social interaction

Nonverbal behavior plays an integral role in social interaction, often influencing others' responses in significant ways. For example, strong nonverbal interpersonal skills have been identified as a solid indicator of better performance in school,¹ patients perceive doctors who exhibit effective nonverbal communication as more competent and caring,² skill in nonverbal communication on the part of

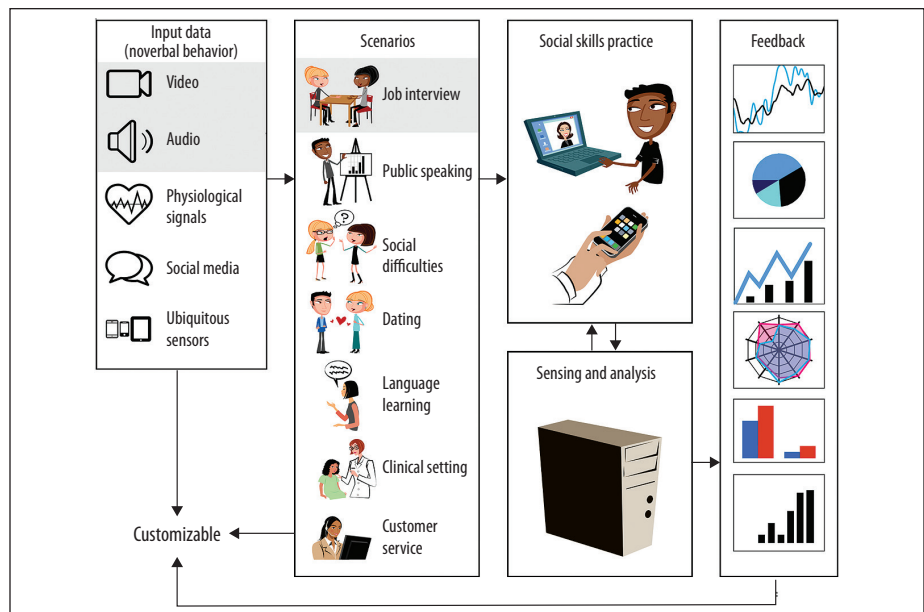


Figure 1. Representation of an overarching, ideal automated sensing and feedback technology that can offer social interaction mentoring for various occasions and purposes, including improving nonverbal communication skills. Shaded boxes indicate areas in which our work has succeeded so far.

customer-facing employees signals overall trustworthiness that boosts a company's business,³ and nonverbal behaviors can predict long-term affinity in intimate relationships.⁴ Clearly, the ability to adjust one's nonverbal behavior appropriately is an important social skill.

This is particularly true in the context of job interviews, during which employers gauge how potential employees will interact with others both inside and outside the company. Numerous studies have found that the most important constructs evaluated in employment interviews are "personality traits and social skills."⁵ Large corporations such as Walmart, Nike, Starbucks, Dunkin' Donuts, and eBay increasingly use automated Web-based video technologies that require candidates to record answers to interview questions for later scrutiny. Using these recordings, employers eliminate unfavorable candidates, often based on simple behavioral rules. For example, reports suggest that Holiday Inn has eliminated candidates who smiled fewer times than a given threshold during interviews.⁶ Such practices underline the growing need for technologies that help people gauge and improve their overall communication skills.

Challenges for automated nonverbal communication recognition

But how difficult is it for computers to recognize human nonverbal behavior? As a way of comparison, consider the extensively studied computational challenge of playing chess. In chess, the first player can open with any of 20 moves, and the second player can do the same; thus, after



the first two moves, there are $20 \times 20 = 400$ outcomes to specify. By extension, according to Avinash K. Dixit and his coauthors⁷ the number of possible moves in chess is on the order of 10^{120} . A computer making a billion calculations a second would take approximately 3×10^{105} years to consider all these possible moves. Because of the complexity involved, chess has become a common platform for computer theorists designing new optimization algorithms and performing complexity analysis.

Now, contrast this with human nonverbal behaviors. Using the 43 muscles of our face, it is possible for humans to produce approximately 10,000 unique facial expressions at any given time (which, for humans, can be counted in milliseconds).⁸ In other words, for two people, the number of possible facial paths after the opening “move” is $10,000 \times 10,000 = 100,000,000$. Adding to the challenge, a chess player might have up to a minute between moves; those engaged in social interactions have barely a moment to register and interpret each facial movement.

People who have trouble deciphering the social “games” others play confront an important challenge: the instantaneity of changing facial expressions. Factor in the complexity imposed by other nonverbal modalities—such as prosody, gesture, and body movement—and the range of possible social–emotional moves in human interaction is truly astronomic, dwarfing that of chess. Rich real-time recognition and understanding of natural nonverbal communication is thus a great challenge for computing, requiring breakthroughs in multiple areas.

Despite the challenges, however, the affective computing community has progressed significantly in developing computational frameworks that can model the interplay, redundancy, and dependency among affective modalities. Further progress in the field of automated nonverbal communication recognition clearly must take into account numerous multimodal sensory channels, including facial expression, prosody, and speech pattern and dynamics.

Facial expressions. Paul Ekman and Wallace Friesen introduced the notion of six basic emotions revealed through facial movement: surprise, fear, happiness, sadness, anger, and disgust.⁸ Further research has sought to expand understanding of facial expression nuances based on naturalistic data congruent with concurrent relevant tasks. For example, our work reveals that smiles can arise from very frustrating situations, and yet appear similar to the smiles Ekman would label as resulting from “happiness” given a static snapshot. Based on this recognition, we were able to develop computer algorithms to look at smile dynamics and automatically distinguish smiles elicited by frustration, delight, or the act of being polite.⁹

Prosody. Studies suggest that the manner in which we communicate is just as important as the content of our

communication.¹⁰ The phenomenon of varying one’s rhythm and melody in spoken language to add communicative information is called *prosody*.¹¹ Its features include register and pitch, speaking rate and phoneme duration, volume, and vocal quality. Our ability to respond to prosody is innate; even newborns extract and respond to prosody in speech.

Prosody is pervasive and conveys meaning on many different levels. For example, Dwight Bolinger¹² notes that varying prosodic markers can result in changes to iterative focus (“Turn right.”/“Turn *right*.”), syntax and phrasing (“Turn right.”/“Turn. Right.”), pragmatic function (“Turn right.”/“Turn right?”), and affect and emotion (to show annoyance, praise, panic, and so forth).

Speech pattern. Speech and its patterns are important modalities to consider in developing models for human–computer interaction. Automated speech recognizer (ASR) activation provides text only, in essence losing all nonverbal markers of the original: dynamics, structure (implicit punctuation, capitalization, formatting, and so forth), and, most important, tone. Developing a framework that can capture relevant nuances of spoken language in real time remains an active area of research.

Advances in sensing technologies could eventually help in this regard, allowing recognition of physiological signals such as electrodermal activity, heart rate, temperature, and movement through quasi-contact or non-contact means using an accelerometer. Analysis of social media has also enabled new ways of capturing information about affect, revealing users’ interaction patterns (mouse clicks on “like,” “share,” “+1,” and “retweet,” among others) with online content. Continued work in this area might provide insights into online relationship dynamics previously unavailable. In addition, now that more than half of Americans carry a smartphone capable of capturing data—including location specifics, pictures, audio level, and communication logs—even more data is available for analysis (www.pewinternet.org/2013/06/05/smartphone-ownership-2013). While privacy concerns remain significant in the collection of such data, social media and multicapacity communication devices are here to stay. They inevitably represent fertile resources for capturing nonverbal communication nuance to improve understanding of human interaction for purposes of developing rich nonverbal sensing in automated systems.

MACH: MY AUTOMATED CONVERSATION COACH

Our own initial attempt to develop a fully automated nonverbal sensing system is MACH (My Automated Conversation coach), designed to sense, interpret, and provide feedback about an individual’s nonverbal behavior during a simulated job interview.¹³ MACH users interact with a

3D character who plays the role of interviewer, asking questions and acknowledging interviewees' answers; at the conclusion of the interview, MACH provides visual feedback to help users evaluate their nonverbal behavior.

MACH's development occurred in three stages:

- First, we modeled the 3D character at the center of MACH on the behaviors and responses of actual interviewers.
- Next, we constructed an architecture to automatically sense and interpret users' nonverbal behavior using a webcam and a microphone, and to visualize these in an intuitive, generalizable format for purposes of providing feedback.
- Finally, we designed an experiment to test MACH's effectiveness as a tool for improving nonverbal communication skills using a sample of 90 job-seeking students.

Modeling the character

To better understand job interview phenomena and model the behavior of the 3D character, we began by running a pilot study in which we observed four MIT career counselors conduct mock interviews with 29 MIT undergraduates. We analyzed the interviewers' behavior to extrapolate information about matters such as expressiveness (for example, interviewers maintain a neutral expression during interviews), listening behavior (they periodically nod as they listen), acknowledgment patterns (they consistently say things like "That's very interesting," "Thanks for that answer," and "I can understand that" to convey acknowledgment), and timing of responses.

We then developed a real-time framework that allows the system to automatically process facial expressions, interpret prosodic features, and recognize speech acts. The 3D character's functionality enables the image to mirror and acknowledge user responses, creating the illusion that it is aware of users' affective states, and then provide personalized feedback regarding users' nonverbal communication modalities. Figure 2 illustrates the process: from sensing and analyzing users' physical/oral data to generating the character's appropriately modeled virtual behaviors to providing feedback for the user.

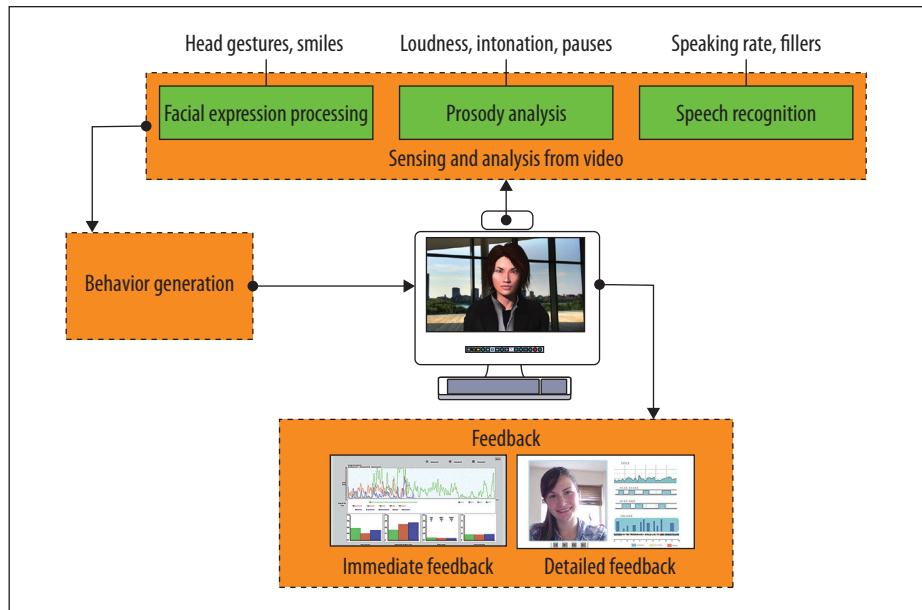


Figure 2. The MACH system works on a regular laptop, processing a user's audio and video inputs in real time. The processed data generate the behavior of the 3D character, which interacts with and subsequently provides feedback to participants.

Constructing the sensing and feedback architecture

During the second development phase, we ran four iterative studies to determine how the system could recognize and provide feedback about nonverbal communication indicators including smiles and head movements; intonation volume, and pauses; and speaking rate and filler word use.

Based on several case studies, we designed user interfaces that could provide both summary feedback for the entire interview in the form of a bar graph (including possible analyses comparing a user's performance over multiple sessions) and specific, focused feedback as users view their recorded interview in real time. Figure 3 shows four sample feedback screens.

Figure 4 illustrates examples of MACH focused feedback, demonstrating the interface's ability to distinguish and point out nonverbal nuances when a person speaks. In each of the three videos analyzed here, the speaker is saying the same sentence—"I am a great candidate; you should definitely hire me"—while varying certain nonverbal speech aspects. In the first video (top of figure), the participant shows no facial expression, exhibits no head movement, and speaks with no variation in pitch or volume. In the second video (middle of figure), the participant increases his volume when he says "strong" and emphasizes the word "definitely" by lengthening its vocal duration. In the third video, the speaker smiles as he starts to say "... you should definitely hire me"; he also nods his head as he says "definitely," aligning the nonverbal action with the intent of the verbal utterance, resulting in more

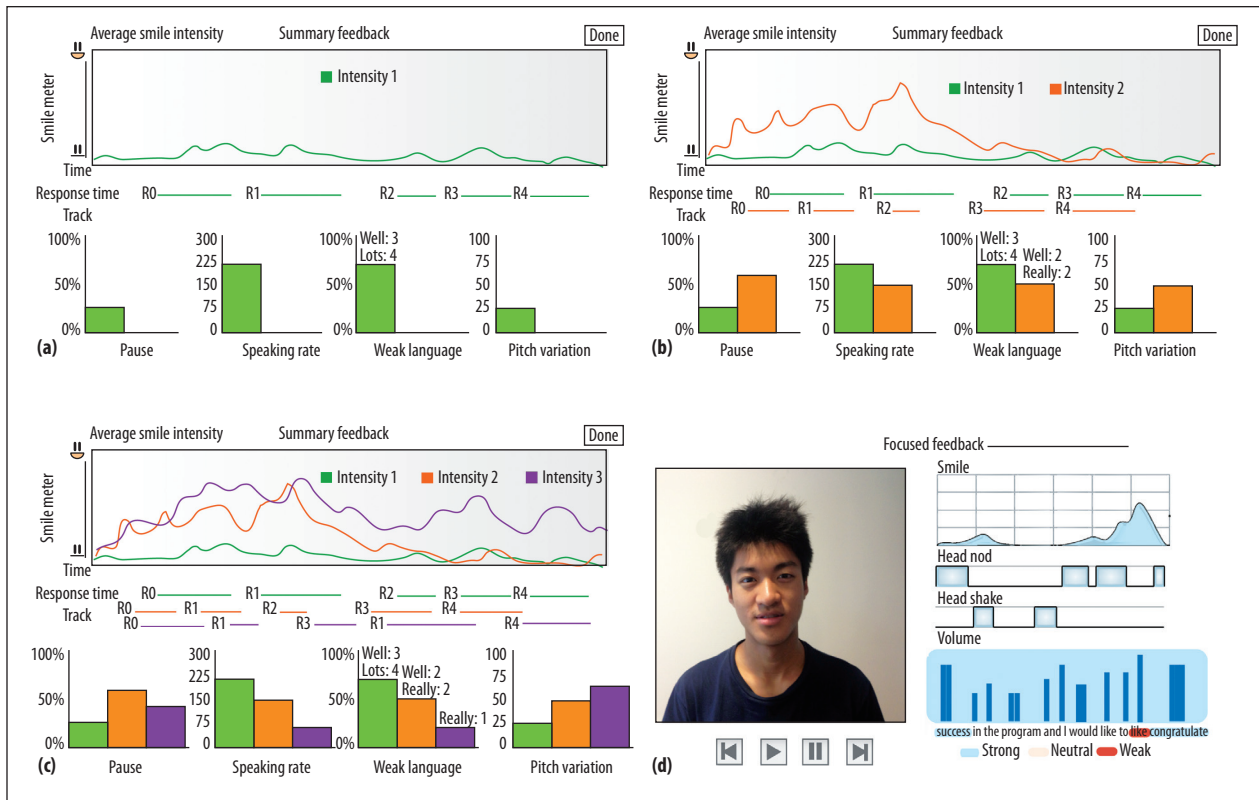


Figure 3. Four examples of MACH user feedback. (a) Summary feedback based on an initial interview session captures the quality of the overall interaction according to various nonverbal communication metrics. (b, c) Users can then compare initial performance quality to their performance in subsequent sessions and, ideally, chart improvement. (d) Focused feedback allows users to view videos of themselves and observe variation in nonverbal behaviors over the course of the interview; charts provide data graphically in real time.

compelling communication. MACH offers graphic tracking of these variations in nonverbal behavior.

Testing

For our initial evaluation of the system, we recruited 90 job-seeking MIT undergraduates and tested them to determine whether the students' interview skills improved after using MACH. We randomly assigned participants to one of three groups. In Group 1, the control group, participants watched commercially available videos demonstrating effective interview techniques recommended by the MIT career services office. Participants in Group 2 practiced interviews with the MACH character and then reviewed their recorded videos only. In Group 3, participants practiced interviews with MACH, reviewed their videos, and received automated feedback on their behaviors, thus interacting with MACH's full functionality.

We began by bringing all participants into the lab for an initial mock job interview with a professional career counselor (who was blind to the study conditions). A few days later, participants in the second and third groups were brought back into the lab for an hour-long MACH session

(those in the control group viewed the educational videos during this period). All participants then returned to the lab one week after the initial interview session and completed a second interview with the same career counselor, this time using a different set of interview questions.

In addition, we recruited two independent counselors from the school's career services office—one male and one female—to observe and rate videos of participants' interview performance in both sessions. We expected these "independent counselor" ratings to be particularly reliable because the observers were blind not only to the study conditions (like the interviewers) but also to the study phase—that is, whether an observed interaction was a baseline or post-intervention interview. Moreover, they had no direct interaction with participants and thus were less likely to be affected by biases, such as rapport with individual participants. Finally, the independent observers could pause and replay the videos, enabling them, if they wished, to consider individual interview performances more thoroughly.

Our analysis of the two MIT counselors' ratings revealed that in general students who used MACH's full functionality (Group 3) demonstrated statistically significant improvement

between their first and second interviews in areas such as “overall performance,” “would love to work with as a colleague,” and “excited about the job”; ratings for Group 1 (the control group) and Group 2 (those who only viewed the video of their first interview without receiving feedback) did not change.

FUTURE POSSIBILITIES FOR AUTOMATED SOCIAL SKILLS TRAINING

In the five months after publishing our initial report on MACH in July 2013,¹⁵ we received responses from more than 2,000 people, all sharing personal stories about problems they had in social interactions. These suggested new ways of thinking about our work and opened up possibilities for designing personalized, respectful automated interfaces to help people improve nonverbal communication skills in multiple social scenarios, well beyond the job interview.

Public speaking

Public speaking skills are essential for advancement in virtually every profession. Yet few people are naturally comfortable speaking in front of an audience; in fact, people often list speaking before a group as their number one fear, higher than the fear of death. Might personalized feedback from automated sensing technologies aid public speaking training?

Most practical public speaking courses focus to a large extent on slowing one’s pace, learning when to pause, varying pitch, using appropriate volume, minimizing use of fillers like “um,” and making eye contact with the audience. Though obvious, these rules are not always easy to apply, so people seek help in professional workshops where they hope a trainer’s personalized feedback can prove valuable and constructive. We attended such a workshop series conducted by Own the Room—a company whose website promises “a proven methodology that helps people learn to be great speakers fast”—where these very rules provided the primary focus. In addition, participants’ presentations were videotaped for class discussion and critique.

While testimonials suggest that workshops like this do benefit many people, there are obviously exceptions. For one thing, professional public speaking training isn’t cheap—the

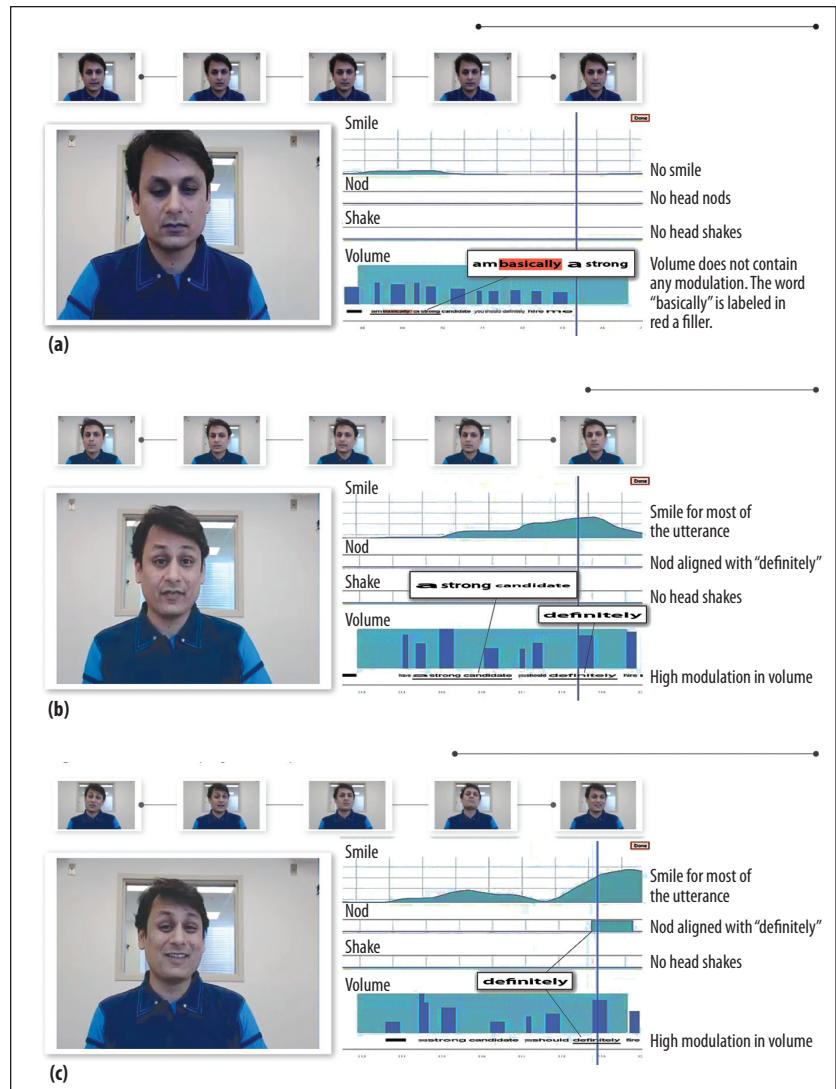


Figure 4. Three examples of MACH focused feedback in which the participant speaks the same sentence (“I am a great candidate; you should definitely hire me”) but varies his nonverbal actions: (a) shows no facial expression, head movement, or vocal variation; (b) increases volume on the word “strong” and lengthens the word “definitely”; (c) smiles while speaking and nods when saying “definitely.”

three-day workshop we attended as guests normally costs \$3,000 per person—and might require participants to take time off work. Second, participants cannot have qualms about being videotaped and critiqued in front of an audience. Finally, whatever feedback participants receive, even from the “expert” trainer, is necessarily subjective.

We feel, then, that sensing technologies like MACH could significantly advance public speaking education, making it possible for people to practice their skills anywhere in the world via a browser, using their computer or mobile device’s webcam and microphone to record their presentations and transmit this data automatically to the cloud for automated analysis and feedback. Participants could even opt for sharing data with friends to critique



via a social network or with anonymous critics via crowdsourcing sites such as Mechanical Turk (www.mturk.com) to receive feedback on a larger scale without fear of personal embarrassment. Public speaking experts might also choose to participate if participants paid a fee.

In fact, an ideal training system might well combine a computer's quantifiable analysis of nonverbal communication skills with more subjective human responses from both expert and lay critics. Whatever the case, the objective data provided by automated sensing systems can motivate users by allowing them to measure concretely the difference that changes in nonverbal behavior can make for effective public speaking.

Customer service and sales

Particularly in service and sales positions, employees' ability to establish rapport with customers not only drives individual job performance (and, in some cases, compensation) but also contributes significantly to a company's overall reputation and success. While technologies such as MACH will probably not go far toward helping workers with the demands and stress involved in direct customer interaction, they can provide other kinds of assistance. For example, repeating the same information hundreds of times each day, a customer service agent might inadvertently start running sentences together and lose dynamic vocal inflection, suggesting lack of energy and enthusiasm. The same is true of sales representatives who work with a basic persuasive script. A variation of MACH could provide periodic feedback for customer-facing employees to evaluate and improve nonverbal elements of their presentations—and do so individually and privately, without any potential embarrassment.

Communication for healthcare professionals

Sensitivity in terms of nonverbal communication is also vital for healthcare professionals—not only because patients generally perceive doctors adept at such skills to be more competent and caring. Medical professionals frequently face the daunting task of revealing disappointing or unexpected information to patients and their families. In addition, doctors and other diagnosticians often have to rely on patients' nonverbal clues to provide them with proper treatment.

It is not surprising, then, that considerable effort has been invested in developing interventions to train medical professionals in social interaction skills, and systems like MACH could supplement this training. For example, programs could be designed with automated characters that present with symptoms of specific mental or physical illnesses and respond in particular ways to different facial expressions, tones of voice, and other nonverbal traits. Trainee users would need to display appropriate body language and also ask the right questions in order to establish rapport with the “patient” and obtain an accurate diagnosis. Follow-up automated feedback on their performance

could provide an objective measure to help improve both the social and diagnostic skills of healthcare professionals.

As one of our respondents wrote to us online,

I am interested in seeing what my personal facial expressions and verbal cues are when speaking to someone. As a mental health therapist, my work is based on communication, but we rarely have the opportunity to view ourselves, much less with this type of feedback.

We believe that many in the profession more broadly share this concern and would welcome the technology we propose here.

Learning a new language

Learning to speak a new language requires extensive conversational practice and feedback. In addition to its semantic and grammatical features, every language has unique prosodic variations and peculiarities that affect the phonetic realization of words and sounds. These interact with meaning and interpersonal context in subtle and intricate ways.

In its simplest form, an intelligent automated sensing interface could allow second-language learners to practice pronunciation with visual feedback. Once words had been broken into syllables and plotted digitally according to a desired prosodic contour, users could compare their pronunciation of a word with its more authentic pronunciation on a prosodic graph quantifying how closely they match the target. More sophisticated interfaces could incorporate a 3D character to monitor and perhaps mirror other nonverbal markers and adapt to users dynamically, even offering encouragement to learners showing signs of frustration or confusion.

MACH might feasibly be adapted for tourists or others wishing to learn nonverbal communication features, such as culturally specific gestures, and basic questions and requests prior to visiting a foreign country—skills acquired most effectively through physical practice. Future technologies have the potential to offer unique, personalized responses for practice in scenarios such as ordering food and asking for directions, and perhaps even more complex interpersonal exchanges with a virtual agent representing a specific language and culture.

Limited learning systems like these are already available for military personnel learning cultural norms prior to foreign deployment.¹⁴ Advances in sensing technologies and cloud computing could soon make it possible for people everywhere to practice the cultural nuances of a new language while at the same time developing linguistic facility.

Helping people with social difficulties

A primary motivation as we developed MACH was our desire to help people whose nonverbal behaviors impede successful and satisfying social interactions.

Unquestionably, the best way to improve interpersonal communication skills is through direct contact with others. However, fearing embarrassment or social rejection and lacking a supportive network, many would benefit from practice using an automated system to provide nonjudgmental, unbiased quantitative feedback.

The technology could be adapted to a wide range of applications, assisting people diagnosed with social phobias and Asperger syndrome or those who simply suffer extreme introversion, like the person who wrote to us,

I have suffered from introversion, shyness, and social awkwardness my entire life. This makes me unwilling to seek and unable to benefit from human coaching. MACH seems like a non-intimidating way for people like me to become more comfortable in social situations.

In this and other ways, potential applications for automated social-emotional technology go well beyond game-playing entertainment and practical coaching, holding great promise for the scientific community. Automated systems can generate social cues and autonomously record continuous fine-grained measurements of social behavior; they offer a repeatable, standardized stimulus-and-recording mechanism, removing subjective bias and increasing both the quantity and quality of data available for study. And because it will ultimately be possible to “turn off” some sensing modules (for example, eye contact and head movement) while making others integral (speech affect or hand-to-face gestures), researchers could selectively probe responses to interaction variables, alone or in combination, that unassisted human observation could not accomplish.

Rich sensing technology has moved into the nonverbal domain and can even now provide personalized feedback for communication skills. As language understanding progresses, allowing comprehension of gestures, affective expressions, and other culturally and contextually sensitive nonverbal messages, these systems will extend their ability to serve multiple human communication needs—both practical and therapeutic—and advance fundamental scientific understanding of successful human interactions. ■

References

1. A.G. Halberstadt and J.A. Hall, “Who’s Getting the Message? Children’s Nonverbal Skill and Their Evaluation by Teachers,” *Developmental Psychology*, vol. 16, no. 6, 1980, pp. 564–573.
2. D.C. Blanch et al., “Is It Good to Express Uncertainty to a Patient? Correlates and Consequences for Medical Students in a Standardized Patient Visit,” *Patient Education and Counseling*, vol. 76, no. 3, 2009, pp. 300–306.
3. J.A. Wood, “NLP Revisited: Nonverbal Communications and Signals of Trustworthiness,” *J. Personal Selling & Sales Management*, vol. 26, no. 2, 2006, pp. 198–204.
4. S. Carrère et al., “Predicting Marital Stability and Divorce in Newlywed Couples,” *J. Family Psychology*, vol. 14, no. 1, 2000, pp. 42–58.
5. A.I. Huffcutt et al., “Identification and Meta-analytic Assessment of Psychological Constructs Measured in Employment Interviews,” *J. Applied Psychology*, vol. 86, no. 5, 2001, pp. 897–913.
6. M. LaFrance, *Lip Service: Smiles in Life, Death, Trust, Lies, Work, Memory, Sex, and Politics*, Norton, 2011.
7. A.K. Dixit, S. Skeath, and D.H. Reiley, *Games of Strategy*, Norton, 2009.
8. P. Ekman and W. Friesen, *Facial Action Coding System: A Technique for the Measurement of Facial Movement*, Consulting Psychologists Press, 1978.
9. M.E. Hoque, D.J. McDuff, and R.W. Picard, “Exploring Temporal Patterns in Classifying Frustrated and Delighted Smiles,” *IEEE Trans. Affective Computing*, vol. 3, no. 3, 2012, pp. 323–334.
10. A. Mehrabian, “Communication without Words,” *Psychology Today*, vol. 2, no. 9, 1969, pp. 52–55.
11. D. Hirst and A. Di Cristo, “A Survey of Intonation Systems,” *Intonation Systems: A Survey of Twenty Languages*, D. Hirst and A. Di Cristo, eds., Cambridge Univ. Press, 1998, pp. 1–44.
12. D. Bolinger, *Intonation and Its Uses: Melody in Grammar and Discourse*, Stanford Univ. Press, 1989.
13. M.E. Hoque et al., “MACH: My Automated Conversation coach,” *Proc. 2013 Int’l Conf. Pervasive and Ubiquitous Computing (UbiComp 2013)*, 2013, pp. 697–706.
14. W.L. Johnson and A. Valente, “Tactical Language and Culture Training Systems: Using AI to Teach Foreign Languages and Cultures,” *AI Magazine*, vol. 30, no. 2, 2009, pp. 72–83.

Mohammed (Ehsan) Hoque is an assistant professor of computer science and codirector of ROC HCI at the University of Rochester. His research interests include making sense of nonverbal data and enabling new human-computer interaction possibilities. Hoque received a PhD from the MIT Media Lab. A member of IEEE, he was awarded an IEEE Gold Humanitarian Fellowship. Contact him at mehoque@cs.rochester.edu.

Rosalind W. Picard is a professor of media arts and sciences at the MIT Media Lab. Her research interests focus on the development of technology to help people comfortably and respectfully measure and communicate affective information. Picard received a ScD in electrical engineering and computer science from MIT. She is an IEEE Fellow. Contact her at picard@media.mit.edu.