
Artificial Intelligence and Technology in Teaching Negotiation

Samuel “Mooly” Dinnar , Chris Dede ,
Emmanuel Johnson, Carrie Straub,
and Kristjan Korjus*

Artificial intelligence (AI), machine learning (ML), affective computing, and big-data techniques are improving the ways that humans negotiate and learn to negotiate. These technologies, long deployed in industry and academic research, are now being adopted for educational use. We describe several systems that help human negotiators evaluate and learn from role-play simulations as well as applications that help human instructors teach negotiators at the individual, team, and organizational levels. AI can enable the personalization of negotiation instruction, taking into consideration factors such as culture and bias. These tools will enable improvements not only in the teaching of negotiation, but also in teaching humans how to program and collaborate with technology-based negotiation

*Corresponding author: Samuel “Mooly” Dinnar, 95 Simonds Road, Lexington, MA 02420, USA.

Samuel “Mooly” Dinnar is an instructor in the Program on Negotiation at Harvard Law School and a lecturer in the Program for Graduate Engineering Leadership at the Massachusetts Institute of Technology (MIT). His e-mail address is sdinnar@meedance.com.

Chris Dede is the Timothy E. Wirth Professor in Learning Technologies at Harvard’s Graduate School of Education and an advisor to Mursion, Inc. His fields of scholarship include emerging technologies, policy, and leadership. His e-mail address is chris_dede@gse.harvard.edu.

Emmanuel Johnson is a PhD student in computer science at the University of Southern California and is advised by Jonathan Gratch. His research explores the benefits of AI simulations that provide personalized feedback for negotiation training. His e-mail address is ejohnson@ict.usc.edu.

Carrie Straub is the Executive Director of Education Programs and Research for Mursion, Inc. Previously, she was the research director for TeachLivE, the project at the University of Central Florida that originally developed and tested the core technology utilized by Mursion. Her e-mail address is carrie.straub@mursion.com.

Kristjan Korjus is a cofounder and CTO of Pactum AI, Inc. He holds a Ph.D. in computer science from the University of Tartu and a First-Class Honours masters degree in mathematics from the University of Manchester. His e-mail address is kristjan@pactum.com.

systems, including avatars and computer-controlled negotiation agents. These advances will provide theoretical and practical insights, require serious consideration of ethical issues, and revolutionize the way we practice and teach negotiation.

Keywords: negotiation, artificial intelligence, AI, machine learning, avatars, computer-controlled negotiation agents, ethics, affective computing, teaching

Introduction

Practical technology systems that are deployed in industry and academic research can shed light on how we may teach negotiation in the future. Today's technology enables an abundance of computing power to capture data and to log files from various negotiations. Researchers and technologists can use machine learning (ML) and artificial intelligence (AI) to improve the ways humans negotiate, as well as train machines how to negotiate. In teaching and learning negotiation, the use of technology to review actual negotiations (be these real-world negotiations or controlled simulations) can help improve future results. The challenges are in choosing what details to monitor and how to turn the data into helpful guidance that will be received and implemented. Will negotiators be open to receiving feedback and coaching from a computer? How different is it to negotiate with a person than to negotiate with a computer-generated avatar controlled by a person? How does one negotiate with a computer-controlled agent (sometimes called a bot)? And can we use what we learn to teach computers how to negotiate instead of (or as agents of) human beings? In this article, we explore several of the many current technologies that are being developed worldwide to teach and practice interpersonal skill-building and organizational learning. We conclude with some reflections on future technological developments.

Various technologies such as video recordings have been used in past decades for capturing and collecting data about negotiation interactions. In recent years, technological capabilities have accelerated significantly to enable the inclusion of more sensors that collect large amounts of multidimensional data. Current technology can also be used to represent one or more parties in a negotiation. This can be done using a computer-generated image, called an avatar, that may look like any type of human, but acts according to manipulations from a particular human operator. In some situations, the computer

itself can be used to make the negotiation decisions, which can be presented as either recommendations or instructions. The recommendations can be presented to a human trainee-negotiator as though they are coming from a coach avatar. The instructions can be translated into moves by a negotiation counterpart avatar (or agent) to which the human trainee-negotiator must respond. Trainees may or may not realize that they are negotiating with an AI agent instead of a person.

We divide our discussion of artificial intelligence and technology in teaching negotiation into five categories, before looking at future convergence possibilities:

- using digital images to teach negotiation skills
- technology for coaching humans to negotiate
- technology for extending negotiation best practices to organizations
- negotiation with AI-automated computer agents
- negotiation with avatars (human driven)

Negotiation teaching and research can both be enhanced through the use of AI and other technologies. While our focus is on teaching we touch on some areas where technology might improve research as well.

Using Digital Images to Teach Negotiation Skills

The practice of negotiation requires many interpersonal skills including problem solving, communication, and interactive decision making. Such skills often are taught through role-play negotiation simulation exercises. In these exercises, students negotiate with each other, with a role-playing teacher, or with an actor (who usually follows a script). At the close of the exercise, students evaluate their negotiating skills through self-reflection and/or receive feedback from their instructor or classmates. Negotiation students learn best when they are able to reflect on their own performance, receive personalized feedback, assess themselves, and integrate their reflections into their learning.

For many years, audio-video recordings of real and simulated conflicts have been used to teach students negotiation strategies and best practices (see, e.g., Williams, Farmer, and Manwaring 2008). For negotiations conducted in writing, such as letters or e-mail, the written text is analyzed; for negotiations conducted over the telephone or similar media, voice recordings are used. When negotiations are in person or over video conference, facial expressions and body language may be included in the analysis.

For decades, technology has been used to aggregate and analyze results from pairs or teams that have negotiated the same simulation. Recent advances in technology have opened up many more possibilities for using digital media to teach negotiators and students how to negotiate. In particular, dedicated platforms for teaching negotiation online such as iDecisionGames (see <https://idecisiongames.com>) allow one to assign roles, collect results, and capture all the video, voice, and text in ways that can be used for both learning and research.

Today's technology allows computers to generate nearly human-looking images and voices that can be used for training. Such images can be used as characters in a movie clip, as virtual players in a game scenario, or to create an avatar that represents the actions and voice of a human. In such virtual settings, there are two types of digital representations of humans: agents (controlled by computer software) and avatars (controlled by humans in real time) (Bailenson et al. 2006).

Simulations involving people or virtual characters have also long been used in fields other than negotiation. For example, early video systems for teaching medical students used branching video-based simulations in which actors portrayed customers or patients, and trainees were provided with a limited menu of responses with preprogrammed consequences. As technology progressed, digital simulations enabled graphic representations such as computer-generated human images and voices to interact with trainees. Systems such as these are useful for demonstrating certain behaviors but rely on a computer application's limited menus and branched responses.

Ebner and Beecher's (2016) "Star Wars: A Trisolan Map" is a branched video-based system that teaches negotiation by enabling students to play characters in a fictional universe. The system uses virtual characters with limited menus that branch off in different directions based on the user's negotiation moves and provides tailored feedback on the user's choices.

Another system that uses limited menus and branching is Trenario's Negotiation Scenarios, developed by Samuel Dinnar. A mobile phone app module that can also run on a computer (see <https://www.trenario.com/negotiation>), Trenario's Negotiation Scenarios teaches basic negotiation terms and concepts. The advantage of such a mobile app system is its ability to provide asynchronous learning that can be done on a personalized schedule in short sessions that are repeated as needed without the presence of classmates or a human instructor, and at the student's own pace. Moreover, the computer-generated images may be modified to assume different genders, ages, voices, languages, and other characteristics. Because the images represent the diversity of the human spectrum, the system can support research into bias as well as

teach users how to negotiate with people from diverse cultures and backgrounds.

We believe that the explosive growth of computing power—similar to the advances taking place in the online Internet gaming industry—will allow for a greater number of choices, branches, and precanned scenarios. Advances in understanding natural language and visual inputs will allow such systems to receive free language (text, voice, and gestures) and give the user an experience that is more realistic than a fixed-menu choice can provide. Responses can be expanded to give users the freedom and experience of more nuanced and layered messages.

Technology for Coaching Humans to Negotiate

Technology that processes a negotiator's past performance can provide individuals with personalized coaching advice during preparation for an upcoming negotiation or while debriefing a session. In some cases, the technology can provide advice to people in real time, during the negotiation itself.

As we have noted, negotiation students learn best when they evaluate and reflect on their own performance, receive personalized feedback, and integrate the self-reflections and feedback into their learning. This reflection, feedback, and integration of data is made more efficient by the technological advances discussed above. Two examples of teaching tools that are made more robust by the use of technology are journaling and postnegotiation assessment.

Students are often asked to reflect on a negotiation after it is complete and to keep a journal of their reflections. This is especially useful over many negotiations (such as a semester series of simulations) and when a mentor reviews the journal and provides feedback. Such reflections can be combined with postnegotiation assessments—often self-assessments based on students' answers to a series of multiple-choice questions. (One popular tool is the Thomas-Kilmann Conflict Mode Instrument that assesses an individual's behavior in conflict.) Today's technology allows online delivery of these assessments or delivery through a mobile app.

One such mobile app is Negotiation 360. Developed by Michael Wheeler, Negotiation 360 (see <https://negotiation-360.com/>) combines assessment functions with short mini-lessons such as video clips and text segments about recommended best practices. Such an app allows an individual to create a habit of preparing for a negotiation and debriefing afterward, as a way of accumulating self-knowledge and self-awareness. By allowing the teacher (mentor) to review and comment on the student's notes, a coaching relationship can develop. The personal device version of the app enables this "microlearning" after each session. The

classroom version of the app allows students to exchange feedback with each other after each exercise and enables both students and teachers to monitor and evaluate progress over many cases, thus enabling the “macrolearning” that occurs over a longer term such as a college semester.

Developments in natural language processing will enable the convergence of assessments and journaling by allowing students’ answers (or journal entries based on prompts) to be spoken or written in free text, thereby facilitating the scoring or categorizing of answers in order to provide feedback. A coaching platform’s integration of video clips, quotes, and statistics from past negotiations, when combined with future developments in computing, will allow for greater personalization of postnegotiation reporting and prenegotiation instruction, based on the negotiator’s accumulated assessments of his or her style, strengths, areas of improvements, and other measures.

Some applications (such as Trenario’s Negotiation Scenarios and the Mursion avatar platform, discussed below) use a computer-generated image to provide feedback and mentoring to the user. Despite initial concerns that students might be reluctant to receive feedback from an avatar (as opposed to a human), users of the platforms have expressed appreciation for this feature, and have commended the empathy with which the feedback was provided. This is a possible area for further research and improvement.

When platforms are enabled to track real-time developments, they provide real-time metrics and facilitate real-time coaching. One technology platform that delivers feedback and guidance to negotiators in real time is Cogito (see <https://www.cogitocorp.com/>). Cogito provides real-time on-screen feedback to human call center representatives while they are handling customer support calls. Such calls often involve negotiations over price, features, or service. Cogito is one of several technologies that came out of Sandy Pentland’s MIT lab. The platform uses AI to perform in-call voice analysis and deliver real-time guidance to human representatives who are handling the calls. Simultaneously, the platform provides detailed monitoring and aggregate insights to their managers. Intuitive alerts (on the representative’s screen) provide instant feedback on speaking behaviors and customer perceptions. For example, human representatives may be guided to speak with more empathy, confidence, professionalism, or efficiency. The early signs of customer frustration or intent to purchase that Cogito provides can help improve service and close more deals. Cogito is a proven early commercial implementation of AI technology that captures emotions and influences how people converse.

Another area of research is the incorporation of social agents into instructional settings. Haring et al. (2019) evaluated the utility of a

virtual agent that provides support during a military strategic mission planning scenario involving a human team leader (who makes the final decisions) and three supporting roles—two humans and the artificial agent. The mission outcome was experimentally designed to fail and to introduce a conflict between the two supporting humans and the leader, a conflict that was mediated by the artificial agent during the debriefing process. Findings indicated that socially intelligent agents show potential for conflict mediation but careful design and implementation are needed to improve team processes and collaboration.

The past two years have seen much growth in the implementation of, and research and developments in, affective computing technology across various domains. Opinions differ as to the effectiveness of this technology and whether its implementation is socially and ethically justified, and while important questions have been raised, public awareness and acceptance have increased through published articles and websites (see, e.g., Greene 2020). Affective computing (a term coined by Picard (1997) and used here to include variants such as emotion recognition, affect recognition, and emotion AI) attempts to use data to recognize emotional states, personal intent, or outward emotional expressions. It also tries to measure, replicate, simulate, or influence such states. As with other implementations of AI and ML, discoveries and capabilities are derived from the computer's ability to process large sets of variables and inputs, including some that may not be intuitively related or easy to track otherwise. Affective computing makes available a broad range of data types, especially combinations of signals such as:

- facial expressions (state of facial muscles, combinations of edges, intersections, and patterns from still or video facial images);
- voice (speaking speed, rhythm, tone variations, gaps, pitch contours, volume);
- body language (eye movement, pupil dilation, gaze direction, head position, gait, posture, gestures, path of travel);
- location (proximity, speed, path, movement);
- physiological measurements (blood flow, heart rate, electrodermal activity, breathing, temperature);
- words (explicit statements of emotion, tone from choice of words, choice of emojis); and the
- use of medium (number and times of messages sent, keyboard force and use patterns, amount of time spent scrolling through and scanning content, social media “likes” and reactions).

One question needing further exploration is how effectively AI can recognize, influence, and simulate human emotion and affect across different cultures. The potential for AI to perpetuate biases related to national origin, religion, race, age, disability, sexual orientation, and gender also requires attention.

The increasing use of AI raises urgent concerns around issues of privacy and control related to the monitoring, recording, and analyzing of collected data. Most research at universities and other organizations is subject to guidelines set by institutional review boards (IRBs), which review the potential harm of a data release and enforce high levels of security and privacy for sensitive information. However, teachers who collect data to aid student learning are not subject to such oversight, and both teachers and students may collect and store data in ways that are insecure or that reveal private data. (For a discussion of these issues, see National Academy of Education 2017.) The adoption of clear institutional policies regarding data ownership, transmission, storage, and confidentiality in both instruction and research is increasingly important.

As we have seen, technology can be used to enable better tracking, processing, and accessing of past experiences and lessons learned. But how can technology build upon these capabilities to improve team negotiation learning, as well as to aid organizations to negotiate better?

Technology for Extending Negotiation Best Practices to Organizations

Artificial intelligence in negotiation can provide a continuous learning system that documents and promotes more effective (best-practice) individual skills, team capabilities, and organizational practices across locations and cultures.

Successful individual negotiators are self-aware and have learned effective practices and common mistakes. To do so, they practice the habits of a negotiation learning loop that includes three efficacious steps: (a) preparing for a negotiation; (b) detecting and responding to events at the negotiation table; and (c) reflecting after the negotiation using a postaction review that captures what worked well and what they might do differently next time. These lessons, learned over time, become the foundation of their individual “personal negotiation theory of practice.” By sharing these lessons learned with their colleagues, negotiators can become a learning team. Each team member benefits from all members’ successes and shortcomings. They learn as a team how they can support various individuals in the most effective way. Extending this process from the team level to the larger organizational level can help drive a culture of accountability and transparency.

In today's technologically connected and mobile world, lessons can be easily shared. By sharing information and practices, a team may develop its own "team theory of practice" and knowledge base of best practices. A mentor, who may be far removed from the negotiation "front lines," is able to distill and contribute his or her insights by providing feedback and guidance. Machine learning and artificial intelligence use context and personalization to facilitate negotiators' postaction debrief.

Negotiation is both an individual skill and a team sport. In many companies, there are sales teams and procurement teams, each with their own unique negotiation culture. In larger corporations, there are teams of teams. These corporations may treat negotiation as a corporate capability (Movius and Susskind 2009). Organizations develop their negotiation learning culture in the same way as individuals—either by default with no learning, or by leveraging the individual learning of each of their negotiators on a daily basis. The biggest barrier to creating an effective learning culture is finding ways to habituate members of organizations to prepare for and debrief every negotiation, practices that may be forgotten in the busy work environment. An organization will not build a useful negotiation knowledge base if reflection and reporting are not integrated into operations. Moreover, lessons learned are useful only if employees can easily find the information they need as they prepare to negotiate. Today the latest developments in AI technology can help negotiators find relevant tips for upcoming negotiations simply and intuitively. Dedicated mobile phone applications with artificial intelligence can deliver systems that are efficient and user friendly.

Shamaym is one such application. The system, built by pilots experienced with postaction debriefings, enables businesses to turn individual mistakes into shared opportunities for ongoing improvement and greater success (see <https://www.shamaym.com/>). First used in the military and adapted to commercial and civic environments, Shamaym is based on the open and honest sharing of performance data and the reflections that follow every "sortie." The process allows entire teams to learn from each member's past performance and to avoid repeating mistakes by taking responsibility for them. When the methodology is supported by an innovative AI-based technology platform for managing the various lessons learned, the system can produce fast results, leading to immediate and constant improvements in an organization. The platform is useable in a variety of sectors, such as high tech, pharma, health care, government, and the military. Its implementation requires some training initially to help with cultural adoption, self-learning demonstrations, workshops, team programs, and effective lessons that illustrate the importance of actionable thinking.

Shamaym and other AI systems can customize best practices for individuals or roles, teams, and organizations. For example, where a CFO belongs to a division's sales and procurement teams as well as the corporation's global "CFO learning forum," lessons that the CFO learns can be shared at the team level and the teams-of-teams level. By applying artificial intelligence technology that connects the individual and the organizational levels, negotiators can promote a continuous learning culture of excellence.

So far, we have focused on the teaching and learning of human negotiators and how technology can help present teaching material, collect data for learning, and assist in coaching. We now turn our attention to technologies that enable the computer to be the one that is negotiating.

Negotiation with Computer-Controlled Agents (Bots)

Humans can learn about negotiation by programming and guiding computers to negotiate, by observing how computers negotiate with humans or other computers, and by negotiating "against" a computer—from one with limited capabilities (text messages only) to computers of the future that will look and feel human.

Negotiation educators have long considered experiential learning an essential component of negotiation pedagogy and role-play simulations with feedback an essential classroom teaching method. The pedagogical value that is derived from simulations and other exercises depends on many factors including efficacy, context, and culture (see, e.g., Druckman and Ebner 2013). Negotiating against a computerized platform—a bot—is one type of experiential learning that can analyze a user's performance and debrief a negotiation. IAGO is a platform that is used for research and Pactum is a platform that is used for commercial negotiations.

IAGO is a virtual agent platform that allows students to practice negotiating and to receive personalized feedback based on well-established pedagogy (see <https://myiago.com/>). It is also used for research into the teaching of negotiation. In IAGO, the bot models a set of basic cognitive functions that support negotiation (such as theory-of-mind reasoning), as well as basic tactics that expert negotiators use to create and claim value (such as anchoring or logrolling). Negotiators can exchange offers or information (i.e., Do you like option A more than B?) and send other messages such as threats and emotional expressions (using emojis). The platform provides tools to customize the agent's behaviors and it can incorporate common biases shown by negotiators (such as the fixed-pie bias). The platform also has an automated feedback component that analyzes a log of the student's interactions

with the agent (offers made, questions asked, and preferences shared) and provides customized feedback on the student's ability to create and claim value. This allows the system to critique a student's understanding of an opponent's interests and a student's ability to share information without being exploited, and improves the efficiency of the final agreement.

IAGO has been used by a community of researchers to build human-like negotiating agents based on a range of principles that (based on prior research) predict good negotiation outcomes. Studies have shown that virtual agents are a good means to teach some aspects of negotiation. Using a developed set of metrics that can measure the quality of negotiation outcomes, researchers have determined that certain principles used to confirm positive negotiation outcomes in a human-human negotiation also indicate success in a human-agent negotiation. Using these metrics and comparing the impact of feedback framing on negotiation outcomes, Johnson et al. (2019) showed that personalized feedback based on the user's recent negotiation session is most effective at helping a student claim more value. The researchers were not as successful in teaching users how to *create* more value.

Research suggests that web-based virtual agents have great potential to be effective tools for teaching negotiation, using the personalization of a negotiation role-playing exercise and feedback (see, e.g., Gratch et al. 2016). Systems like IAGO focus on the cognitive aspects of negotiation such as tactics and problem solving but do not include the complexities of spoken language (relying on menu-based interaction). Further, their cognitive abilities are currently limited, making them at present most appropriate for teaching introductory negotiation concepts. With advances in algorithms and access to a large data set, these weaknesses can be remedied, and the broader use of virtual agents as supplemental tools for helping students learn is possible.

Pactum is an AI-based system that helps global companies offer personalized, commercial negotiations on a large scale (see <https://pactum.com/>). The platform assists corporations with negotiating their low-value, high-volume deals. To automate the process, a human team of analysts maps a "value function" for a given set of negotiations, then feeds these functions into a negotiation chatbot that is capable of conducting negotiations autonomously based on algorithms developed by scientists. The full negotiation is limited in complexity and is done via text messages, so the autonomous bot is able to negotiate at large scale, bargaining with many vendors in parallel and responding quickly to their messages. In some initial implementations of the bot, the negotiation was framed not as a zero-sum game but as a text message discourse where the intention was to improve the deal for both sides. Afterward,

a majority of these selected vendors indicated that they would prefer to negotiate again with a bot rather than a human, citing the clear, fast, and polite communication as one of the main reasons. Perhaps even a text-only interaction with a bot (that resulted in added value) was a better experience than the alternative of having no negotiation, or one with little added value, due to lack of human availability.

Given the large number of real-world vendor negotiations, additional insights can be gleaned by testing modifications to the chatbot's algorithmic settings (called a/b testing), determining those that create more value. The human scientists can adjust the settings for the chatbot accordingly and update their best practice assumptions. Over time, with further machine learning, some of the algorithmic adjustments will be proposed by the computerized platform itself.

With this commercial bot implementation, it is the humans who are “coaching” the computer-controlled agent how to negotiate. The human scientists are then able to use the automatically collected data to further improve the computer-controlled agent's performance, as measured by certain predetermined metrics. It is the humans who set the goals that the computerized agent will try to maximize. As with other technologies where the role of humans has been elevated from the “doer” to the “designer” or “programmer,” it will be our role in the future, as negotiation educators, to teach humans not only to manage other negotiators, but also to direct bots that negotiate autonomously. This new role raises serious questions and should be undertaken with caution. Since research has shown that humans allow themselves to be less ethical when negotiating through human agents (see, e.g., Rees, Tenbrunsel, and Bazerman 2019), might it be wise to train human designers (or operators) to view the computer-controlled agent's adherence to ethical values as one measure of success? What happens when automatic negotiation systems become too competitive, such as when they are not design-limited to situations that can be framed as improving the deal for both (or all) sides of the negotiation? Will systems developers be able to agree on a “code of ethics” that does not allow using deception, or other unethical methods, in the negotiation? These are only some of the questions raised by this technology.

Now that we have explored technology that helps humans learn to negotiate better and technology that teaches humans how to coach bots to negotiate, let us turn our attention to technology that aids human-like coaching of human negotiators through the use of digital humans called “avatars.”

Negotiation with Avatars

As flight simulators are used to train pilots, so can complex simulators use mixed-reality digital systems to complement traditional role-play simulations with particular scripts, data recordings, playback, and coaching.

When humans negotiate with other humans, there are infinite combinations of behaviors from which to choose. But when prescribing behaviors for a human-looking computerized agent, there is a finite number of scripted scenarios with forced-choice and limited branches. One way to overcome these limitations has been to use avatars (human-driven, computer-generated images) rather than agents. In the world of Internet games, authoring shells were used to produce virtual worlds like *Second Life*, in which (as an illustration) the digital avatar of a human-simulated patient could interact with the digital avatar of a medical resident. While this enabled interactive textual conversations without forced choices or branching, the limitations of the gaming interface reduced communication to text messages only. As has been demonstrated in the education field, the lack of verbal and nonverbal interactions undercuts the authenticity and transferability of the experience and the learning (McGarr 2020).

In the real world, people must emote using a rich array of verbal and nonverbal signals. Developing these skills through practice and coaching is central to learning sophisticated forms of interaction such as negotiation (Bondie, Macenido, and Dede in press). Clicking on a limited menu of options or typing what one wants to say is far from adequate. To overcome such shortcomings, Mursion (see <https://mursion.com/>) employs a technology called “digital puppeteering,” which is based on three elements:

1. a virtual environment that presents a detailed digital context;
2. personalized digital avatars controlled by a human being (a simulation specialist) that can assume many possible physical forms (e.g., a small girl, a grown man) and enable the simulation specialist sequentially to control the avatars’ interactions with realistic verbal and nonverbal capabilities (e.g., voice volume, tone, and prosody; nonverbal facial expressions, postures, and gestures); and
3. freeform responses from the trainee (sitting in a training simulator), who has the full range of verbal and nonverbal communication signals, which are sensed by the system and simulation specialist to drive what happens next.

The system in which a simulation specialist leverages AI for algorithmic assistance may be referred to as an Avatar-Mediated Interactive Training and Individualized Experience System (AMITIES) (Nagendran et al. 2014). To understand how an AMITIES works, an analogy to the larger-than-life wizard in the Wizard of Oz, who is actually a little man behind a curtain, may be helpful. The polymorphism of the avatars (their ability to take on many forms) is a strength in studying implicit bias and in preparing trainees to work effectively across a wide variety of developmental levels and demographic characteristics. In the past five years, such systems have become practical and affordable. They have been used to prepare people for one-on-one negotiation through the depiction of one avatar and for multiple-party negotiation through the use of multiple avatars. The systems may be used, for example, to train teachers how to control a group of children or police officers how to manage a small crowd.

The Mursion platform is powered by a blend of artificial intelligence and live human interaction, and provides experiential learning to help people improve their knowledge, skills, and dispositions. These mixed-reality simulations are built on a technical framework that supports avatar behaviors (those initiated, influenced, or directly controlled by a human simulation specialist), while also allowing for some platform agent-based behaviors (those controlled strictly by computer software, while the human is focused elsewhere). The software provides algorithmic assistance for a simulation specialist to inhabit a 3D (virtual reality) or 2D (mixed reality) digital environment and drive up to five digital characters at one time during a simulation. The artificial intelligence module blends the input from the simulation specialist—including head motion, voice morphing, and lip-syncing—with the avatars' unique body language and facial expressions to create the apparent notion that a learner is interacting with several avatars. This produces a strong suspension of disbelief in the learner.

By using trained professionals to orchestrate the interactions between learners and avatar-based characters, Mursion simulations are sufficiently realistic to measurably improve skills in many areas, including negotiation and leadership. For example, Michael Wheeler—a leading teacher and researcher of negotiation—worked with Mursion to develop a scenario-based interaction involving an employee who wants to negotiate a raise with two members of the organization's leadership, at a time when the organization is losing money and laying off personnel. Wheeler worked with the simulation specialist to ensure the quality of the experience and to test the system for its efficacy in developing the interpersonal skills required in negotiation. Several students and professors participated in Mursion training using this scenario. They each

individually participated in a simulation facilitated by a host avatar who set the learning objectives and strategies for the interaction. After the learner's first attempt within the simulation (negotiating with the two executive avatars), they each engaged in reflective debrief facilitated by the host avatar. Some of the learners completed a second attempt at the simulation, resulting in the opportunity for active experimentation and corrective-action implementation. All found the simulation engaging and were impressed by the suspension of disbelief that quickly made the experience seem real and valuable. Interestingly, some students reported finding it more realistic to negotiate with the avatar than to negotiate in-person with a fellow student playing the role of his or her boss.

Experiences with virtual- and mixed-reality digital puppeteering systems like Mursion can be used to complement traditional role plays. In particular, scripts can be developed that focus on each student's interpersonal weaknesses such as an inability to say no to a likable person, to speak out against microaggressions, or to challenge implicit biases. This is analogous to the role flight simulators play in training pilots—students are put in difficult situations and learn to practice the skills needed to deal with them.

Future iterations of the software will include features that provide predetermined and optional input mechanisms for simulation specialists to tag impromptu events of interest during the simulation, resulting in a rating process of learning performance against objectives. Additionally, the audio data from the learner will be processed using piecewise-temporal analysis to extract features such as pitch, formants, jitter, shimmer, silence, and other vocal characteristics. These data will be processed to yield metrics that are significant and valuable in the simulation, including active speaking time and listening time. Finally, software capabilities will allow the reporting of complex constructs such as engagement (measured, for example, by the energy in the head motion and vocal activity) during the analysis. All ratings from the simulation specialist will be time-stamped to the millisecond, allowing for synchronization of the rating data with the raw-data streams (verbals and nonverbals of both learner and simulation specialist), which will be used to produce customized metrics. Each session will be recorded and learners will receive an after-action review for each simulation.

A current obstacle to controlling cost and increasing scale is the need for skilled simulation specialists, since simultaneously manipulating the verbal and nonverbal attributes for a sequential series of avatars is complex. However, with further advances in AI and multi-model analytics, people with lower level qualifications will be able to realistically portray avatars and the software will become more user friendly, reducing the need for specialized, expensive personnel and enabling

teaching assistants or students to play the operator role in negotiation instruction.

Teaching with complex simulators that use mixed-reality digital systems and are powered by human-driven avatars and/or AI computer-driven agents enables teachers to create diverse, personalized, replicable experiences (ranging across different cultures, locations, group sizes, etc.), while the system collects data for continued improvements in learning, teaching, and research. The future of this and the other technologies described above holds great potential, as we discuss in the next section.

Conclusion

In this article, we have discussed technological advances that have significantly changed the practice and teaching of negotiation, focusing on using digital images to teach negotiation skills, using technology to coach humans to negotiate, extending negotiation best practices to organizations, negotiation with AI-automated computer agents, and negotiation with human-driven avatars.

Our exploration has led us to the following reflections and conclusions:

- Improvements in the processing of audio data (voice) and physical engagement (facial expressions, body language) will yield valuable metrics and enable the reporting of complex constructs, which will be used to produce customized smart-metrics. AI computing will allow further personalization of postnegotiation reporting and prenegotiation preparations for human negotiators. Technology that tracks and processes data from past negotiations can help trainees to learn from their own experience and create a personalized negotiation theory of practice; it will also assist teams and organizations to create their own theories of practice.
- Future training platforms that use avatars will be used to tag and collect more data that can be used to research the effects of biases and culture on negotiation. Complex constructs and metrics will be used for offline research as well as real-time tracking during negotiations. As avatar technology improves, the qualifications needed to become a simulation specialist will be simplified and the use of avatars will be expanded. Eventually students will learn to become simulation specialists to deepen their own experiential learning with increasing complexities (such as the number of negotiators).
- Computer-controlled agents show great potential as effective tools for teaching negotiation, with students tasked to negotiate with

a counterpart, sometimes not knowing whether the entity on the other side is a real human or a computerized virtual entity. The computer-controlled agent will be adjustable to address the customized pedagogy objectives for each student, negotiation role-playing exercise, or feedback session. Future advances in spoken language and cognitive elements will help students learn additional and more nuanced negotiation skills, potentially matching or exceeding what is currently possible using real people or avatars.

- Future wide-scale adoption of computerized agents to conduct real-world negotiations will provide additional opportunities for theoretical and practical insights. Corporate scientists and academic researchers will identify new methods of creating value and will formulate new best practices that will be useful in the classroom. Over time, the computerized AI platform will be able to “coach” itself on how to improve.
- Teaching negotiation will include teaching humans how to use avatars, how to program AI computerized agents, and how to “coach” the complex AI negotiating systems to achieve better results and/or create more value. As in many artificial intelligence implementations, this will need to be done with special care to ethical considerations and the need to prevent the perpetuation of historical biases. At the same time, humans will continue to negotiate with both machines and other humans. Technology can help teach and coach individuals how to negotiate and become better team players, as well as how to negotiate as a team or a large organization.

AI and ML technologies are being deployed at an exponential rate. With them will come significant opportunities and challenges for conducting negotiations and for teaching negotiations. The days of the negotiation “Turing machines” are here, or coming soon, where an individual cannot determine whether they are negotiating with another human or a computer-controlled agent. As researchers and teachers of negotiation, we need to think how best to equip our students (and ourselves) with the ability to negotiate at yet unforeseen levels and scales. And may the best bots win-win. . . .

REFERENCES

- Bailenson, J. N., N. Yee, D. Merget, and R. Schroeder. 2006. The effect of behavioral realism and form realism of real-time avatar faces on verbal disclosure, nonverbal disclosure, emotion recognition, and copresence in dyadic interaction. *Presence: Teleoperators and Virtual Environments* 15(4): 359–372.
- Bondie, R., Z. Macenido, and C. Dede. in press. Interaction principles for digital puppeteering to promote teacher learning. *Journal of Research in Teacher Education*.

-
- Druckman, D., and N. Ebner. 2013. Games, claims, and new frames: Rethinking the use of simulation in negotiation education. *Negotiation Journal* 29(1): 61–92.
- Ebner, N., and M. Beecher. 2016. Star Wars: A Trisolan map. Available from https://s3-us-west-2.amazonaws.com/eportfolioabeecher/Trisolan+Map/story_html5.html.
- Gratch, J., D. DeVault, and G. Luca. 2016. The benefits of virtual humans for teaching negotiation. In *Intelligent virtual agents: Proceedings of the 16th international conference: IVA 2016, Los Angeles, CA, USA*, edited by D. Traum, W. Swartout, P. Khooshabeh, S. Kopp, S. Scherer, and A. Leuski, 283–294. New York: Springer.
- Greene, G. 2020. The ethics of AI and emotional intelligence: Data sources, applications, and questions for evaluating ethics risk. Available from Partnership on AI, <https://www.partnershiponai.org/the-ethics-of-ai-and-emotional-intelligence/>.
- Haring, K. S., J. Tobias, J. Waligora, E. Phillips, N. L. Tenhundfeld, G. Lucas, E. J. de Visser, J. Gratch, and C. Tossel. 2019. Conflict mediation in human–machine teaming: Using a virtual agent to support mission planning and debriefing. Presented at the 28th IEEE International Conference on Robot and Human Interactive Communication, New Delhi, India. Available from <https://ieeexplore.ieee.org/document/8956414>.
- Johnson, E., G. Lucas, P. Kim, and J. Gratch. 2019. Intelligent tutoring system for negotiation skills training. In *Artificial intelligence in education. AIED 2019. Lecture notes in computer science*, edited by S. Isotani, E. Millán, A. Ogan, P. Hastings, B. McLaren, and R. Luckin, 122–127. New York: Springer.
- McGarr, O. 2020. The use of virtual simulations in teacher education to develop pre-service teacher behavior and classroom management skills: Implications for reflective practice. *Journal of Education for Teaching* 46(2): 159–169. Available from <https://doi.org/10.1080/02607476.2020.1724654>.
- Movius, H., and L. Susskind. 2009. *Built to win: Creating a world-class negotiating organization*. Boston, MA: Harvard Business Review Press.
- Nagendran, A., R. Pillai, A. Kavanaugh, G. Welch, and C. Hughes. 2014. A unified framework for individualized avatar-based interactions. *Presence: Teleoperators and Virtual Environments* 23(2): 109–132.
- National Academy of Education. 2017. Big data in education: Balancing the benefits of educational research and student privacy. A workshop summary. Washington, DC: National Academy of Education.
- Picard, R. W. 1997. *Affective computing*. Cambridge, MA: The MIT Press.
- Rees, M. R., A. E. Tenbrunsel, and M. H. Bazerman. 2019. Bounded ethicality and ethical fading in negotiations: Understanding unintended unethical behavior. *Academy of Management Perspectives* 33(1). Available from <https://doi.org/10.5465/amp.2017.0055>.
- Williams, G. R., L. C. Farmer, and M. Manwaring. 2008. New technology meets an old teaching challenge: Using digital video recordings, annotation software, and deliberate practice techniques to improve student negotiation skills. *Negotiation Journal* 24(1): 71–87.