

## Hardware and Software Enhancements: Conversational Computer Agents with Initiative

**Dan M. Davis**  
ICT, Univ. of Southern California  
Los Angeles, California  
dmdavis@acm.org

**Evan Jaksha & Skander Guizani**  
Visiting Researchers, ICT, USC  
Playa Vista, California  
{evan.jaksha, skander.guizani.tn}@westpoint.edu

### ABSTRACT

This paper presents several innovative approaches to creating a more human-like computer agent that is capable of not only responding to live-user questions, but can also formulate and convey appropriate questions and comments. The paper begins with recounting the need for such agents and the current paucity of both utilities for and research into this challenging field. Data and analysis from the authors' work in virtual human agents are adduced to further identify and characterize the current state of the art. Further, anecdotal examples are given as to the impact of having a computer agent capable of only responding to direct and germane questions. This is contrasted with a live human who would not only respond to questions, but would conceptualize potential impacts, comment on inputs, and then request information or suggest future courses of discourse. To implement a virtual equivalent, several alternative algorithms and approaches are reviewed, with the varying features of each being evaluated. The use of Artificial Intelligence (A/I) is considered, as are the potential hardware suites and long-haul digital communications necessary for the success of the system. Potential misfeasance by the A/I functions is considered and an Ethical A/I approach is recommended. An outline is given and justified for the appropriate metrics for such an effort and the best way to measure and evaluate the data against goals of the project. Additional research and probable impacts are considered and resolved. The implementation and exploitation of emerging hardware capabilities, *e.g.* quantum computing and software advances, *e.g.* Deep Learning in big data sets are considered. Future research directions are identified and justified. Analysis of the needs for these emerging capabilities and the opportunities for extensibility into other initiatives are set forth and implementation implications are discussed to benefit the readers who may want to pursue these approaches.

### ABOUT THE AUTHORS

**Dan M. Davis** is active as a consultant at the Institute for Creative Technologies (ICT), University of Southern California (USC), focusing on large-scale DoD simulations and avatar uses. Prior to retirement, he was the Director of the JESPP project at USC for a decade. As the Assistant Director of Advanced Computing Research at Caltech, he ran Synthetic Forces Express, bringing HPC to DoD simulations. He also served as a Director at the Maui High Performance Computing Center and in computer research roles at the Jet Propulsion Laboratory and Martin Marietta. He was the Chairman of the Coalition of Academic Supercomputing Centers and has taught at the undergraduate and graduate levels. As early as 1971, Dan was writing programs in FORTRAN on one of Seymour Cray's CDC 6500's. While in the Marine Corps, he saw duty in Vietnam as a Cryptologist and retired in 2002 as a Commander, U.S.N. He received B.A. and J.D. degrees from the University of Colorado in Boulder.

**Evan Jaksha** is studying Computer Science and currently focusing in the use of virtual humans as an effective interface to address a range of critical issues in the military. He comes from a military family with both his father and two siblings currently serving in the US Armed forces. His research interests include cyber warfare and national security. Growing up in San Diego, he is a long-time resident in Southern California. He is scheduled to graduate in (2022 with a degree in Computer Science, with a Minor in Cyber Security from the United States Military Academy in West Point, New York.

**Skander Guizani** is studying Electrical Engineering and currently focuses his research efforts on hardware, infrastructure, and operations of long range secure communications systems. Skander is from Tunisia and is a foreign exchange student at the United States Military Academy at West Point, New York. Back in Tunisia, his father is a Colonel and commands a unit in that country. While at the Institute for Creative Technologies, Skander will be working with the MentorPal team to enhance the virtual human conversations for advising high school seniors. He is intent on pursuing a doctoral degree and I focusing his research on that goal. He anticipates receiving a degree in Electrical Engineering in the spring of (2022 from the United States Military Academy, West Point.

# Hardware and Software Enhancements: Conversational Computer Agents with Initiative

**Dan M. Davis**  
ICT, Univ. of Southern California  
Los Angeles, California  
dmdavis@acm.org

**Evan Jaksha & Skander Guizani**  
Visiting Researchers, ICT, USC  
Playa Vista, California  
{evan.jaksha, skander.guizani.tn}@westpoint.edu

## INTRODUCTION

The issues to be addressed are the advances needed and capabilities emerging that will allow the human avatar computer agents to show more initiative in their virtual conversations with human users. While great strides have been made in Artificial Intelligence (A/I) and Natural Language Processing (NLP) that now enable digital computers to respond to users in a conversational way that emulates a human dialogue, their ability to initiate conversations is much more severely limited. The authors will present insights from their design, implementation and evaluation of systems designed to be a surrogate for human-to-human conversations. Then, the limitations of these response-oriented systems will be described using results from human testing, with an emphasis on why a computer-enabled initiative may be important. Further they will discuss current software, hardware and system constraints that make solving these problems difficult, if not virtually impossible. Then a few emerging techniques and technologies will be discussed and the authors will offer their conjectures as to how these new capabilities may enable of more human-like interaction with the users. To implement a virtual equivalent, several alternative algorithms and approaches are reviewed, with the varying features of each being evaluated. The use of Artificial Intelligence (A/I) is considered, as are the potential hardware suites and long-haul digital communications necessary for the success of the system. Potential misfeasance by the A/I functions is considered and an Ethical A/I approach is recommended. An outline is given and justified for the appropriate metrics for such an effort and the best way to measure and evaluate the data against goals of the project. Additional research and probable impacts are considered and resolved. The implementation and exploitation of emerging hardware capabilities, *e.g.* quantum computing and software advances, *e.g.* Deep Learning in big data sets are considered. Future research directions are identified and justified. Analysis of the needs for these emerging capabilities and the opportunities for extensibility into other initiatives are set forth and implementation implications are discussed to benefit the readers who may want to pursue these approaches.

## Background

The sophistication of the human-computer interface has come a long way since the advent of common accessibility of computers beginning in the 1970's. One of the authors, (Davis) can regale the younger computer scientists with tales of sitting at an IBM punch card machine for hours typing out one line per card, then only to have to hand them over to an operator and wait for an hour or so, then pick up the printed output typed by line printers on sprocket paper. Today, virtually every house has several devices that respond to the human voice: "Hey Siri, what's the temperature in New York City?". When 2001: A Space Odyssey was released, the populace was at first enthralled by, then threatened by "HAL" who engaged "Dave" in very real conversations. Some of that technology was hard to implement, but it can be done today. "Bixby, where is the closest Italian restaurant?". That is not on a space ship with a huge room-sized computer, but on one's cell phone. However, the state of the art seems not too close to creating a malevolent presence, but it is still missing a human-like initiative.

## Military Context

The paper adheres to the position that there is a significant interest within the military concerning communication issues, and they are not the same as industry, non-governmental organizations and academia. These issues may be articulated alone, but they are synergistic in their impact on the issues of national importance. As in all human behavior, observations may be insightful, but should usually not be taken as being either universal or dispositive.

One issue is the formalization and visible manifestation of the need for discipline. Another issue is the criticality of the missions assigned. The military is distinguished by a very atypical understanding of and mutual appreciation of job boundaries. The team duty roster is much more ingrained and accepted than in civilian life. Another item upon which the authors agree is that the military personnel come with a much more deeply ingrained sense of oneness. This may be most clearly typified in the Marine Corps insistence on not having hyphenated Marines: One is not a Reserve-Marine, a Woman-Marine, an Asian-Marine, a Cryptologic-Marine, or an “Anything-Else”-Marine: one is a “Marine” or not. Even unit patches are not common on Air Force, Navy and Marine Corps uniforms. The last issue is that of the lives at stake. As opposed to some civilian industries, DoD personnel are there to fight and evolution may be significantly more stressful in the military context.

Another issue with which most DoD personnel would be very familiar is the recognition of and adherence to the military job training and assignment principles. The authors have all seen how personnel trained and certified in one occupational specialty, but manifesting a much greater skill in another needed capability, will be informally and exclusively assigned to duties for which they are neither trained nor certified. That tradition even finds its analog in the command structure. There are times in the military when, while retaining the position of the titular head of a unit, a wise commander will designate much of the day-to-day operation of the unit to a more compelling and effective leader. Again, the widespread regard for mutual respect for position and hierarchy is invariably reflected in the junior’s strict adherence to deference to the senior, albeit the exercise of the command functions delegated to him from that senior. All of these issues keep the Tuckman insights in play, even in the military context.

In addition to that, in the military, when a new unit is “stood up,” it is very likely that all of the personnel will not have ever met the others in the unit. In the civilian experience, it is often the case that new teams are made up of a majority of members who have been on similar teams with the members of the new team. Again referring to the motion picture industry, a producer usually has a reasonably free rein on the professionals he employs and some experience with the personalities of those candidates, *e.g.* it was well-known in the Southern California film industry that hiring David Niven would entice other actors to join the cast, as he was famous for his civil congeniality that made every film of his a happy production experience (Lord, 2019). Military commanders have virtually none of that freedom in selecting their own crews.

One major issue that the officers who are also co-authors herein see is the geographical separation of prospective team members. The first several stages of the Tuckman analysis require personal contact and face-to-face experience. The authors maintain that this is also best accomplished with a mediating senior being present. Falling once again on the motion picture industry, there are often protracted face-to-face negotiations as to team composition and duties prior to the first day of team formation.

The most salient difference could be the differences in goals. A slow or mal-forming team in civilian life may lead to loss of money or competitive market advantage. In the military context, the same delay or misbegotten effort in team formation may very well result in, not only equipment and monetary losses, but in mission failure or loss of life.

The authors assert that many of the above issues can be successfully, and perhaps only, addressed by the effective application of computational science and global network communications. Some of these computer assets may necessarily be other than digital computers. Research into both Quantum computing and a resurgence of interest into analog and hybrid analog/digital computers may be necessary to provide the desired and required computation power. An operational quantum annealer is now operational



**Figure 1 – A D-Wave Quantum Annealer**

## **VIRTUAL CONVERSATIONS RESEARCH**

One area of rapid growth in the computational sciences is the area of Virtual Humans and Natural Language Processing. As the world has recognized, early predictions of virtual human simulations and other advances in this arena have not matched the predictions. Some areas have sprung ahead in ways not envisioned. Dick Tracey's wrist

watch phone has been put to shame by current models of smart watches that not only provide instant communications anywhere in the world, but also sense an unexpected fall and summon medical care on their own (Technokogizer, (2019). On the other hand The "HAL 9000" fully conversational, sentient and self-aware computer from the movie Space Odyssey is still a significant ways off (Wikipedia, 2019). However, great strides have been made at USC's Institute for Creative Technologies and elsewhere.

### Advances at The Institute for Creative Technologies




Figure 2 – ICT's MentorPal Display

An early version is shown in Figure 2 above. The prototype for this interactive mentor targeted the Microsoft Surface® line of tablet computers. The limited power, RAM and secondary memory of these small portable devices mandates optimal sizing of all processes and data storage and transfer (Nye *et al.*, 2017).

The authors assert that the insights gained and the techniques employed apply to model reductions in practically any specific domain, rather than only to chatbots. Part of this assertion is recognizing the need to prepare for operations under conditions of limited computational power and data capacity, while at the same time designing code to make it amenable to scaling if a less constricted environment is available. This would allow computer assistance for team formation, no matter where the team members were stationed prior to the teams converging. These constrained environment issues are known to be common in much of the work currently being conducted at ICT (ICT, 2018).

The basic objective of the Mentor PAL project was to provide a proof-of-concept version of career mentoring to students who may otherwise have no, or severely limited, access to advice or mentoring as they face the daunting task of selecting a career. The ultimate goal is to increase the number of technically trained personnel available. They might be used in the uniformed services or as civilian researchers for the DoD. To accomplish this, prospective career candidates may make other choices if they do not have someone to whom to speak who is experienced in and feels positively toward the desired careers. The authors envision a similar mentoring process in the team formation scenarios of the future. Users, who do not otherwise, have access to a mentor, may benefit from a virtual mentor. The reason for such lack of access may be due to a number of reasons, the two most prevalent of which the authors hold to be: geographic remoteness and poor network connectivity (Crisp *et al.*, 2009).

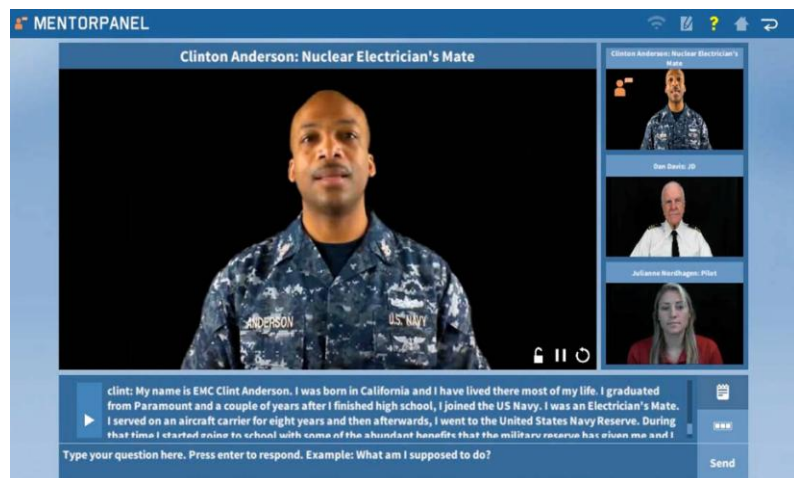


Figure 3 – Mentor PAL for Advising High School Seniors about STEM Careers in the Navy

The solution advanced by ICT was to produce a computer-generated mentor to be available via any system with the appropriate software installed. The user would be presented with a computer screen interface via which he could engage a mentor in a conversation-like exchange about issues of concern. Input from the user would be by text entry or audio speech recognition software. This mentor was envisioned as being both compelling and engaging. ICT advanced the position that the mentors would be most compelling and engaging if they were real people (Traum *et al.*, 2915) recorded live and exhibiting a demonstrated screen presence.

### Making a Conversational Agent with Initiative

Initially, a list of germane questions would be drafted. Depending on the topic, this might total a thousand or so items. The creative process requires a broad knowledge of both the questions that are a part of the mentor's experience and represent the questions that the target users are likely to ask. In the previous projects, the team considered input from several members who had professional and Navy experience in order to generate a thorough list of questions. A previous paper (Nye *et al.*, 2017) discussed this process at greater length.

The next step is to have the mentors record a varying number of video clips relating to their experience in the service and their duty in their own profession. These clips are then stored in a standardized database. When the user poses a question the program selects the most appropriate clip to play. This must be done very rapidly to sustain the desired "conversational" effect of the exchange. Based on the team's experience with briefing users, the clips were designed to be on the order of 90 seconds or less. This technology rests on previous research into the use of Virtual Humans acting as program interfaces in counseling (Morbin, *et al.*, 2012) and in history capture (Artstein *et al.*, 2014). The team found that these videotaped mentors are difficult to classify easily: sometimes being referred to as Virtual Humans, but they are actually intelligent agents that sequence and present recordings of real people. This has been termed "time-offset interaction" in some related work. There still needs to be a substantial corpus of carefully collected questions and answers available to the system that can be answered within the latency limits that circumscribe the feel of conversational dialogue.

In Figure 3, the flow begins at the top left and proceeds clockwise through the process. The icons are notional only and should not deceive the reader as to their importance, weight or size, *e.g.* the storage icons approximately the same image-on-page size, whereas, in physical fact they differ by three orders of magnitude in data size. The important steps are the data flows from the storage locations into and out of the computational functions.

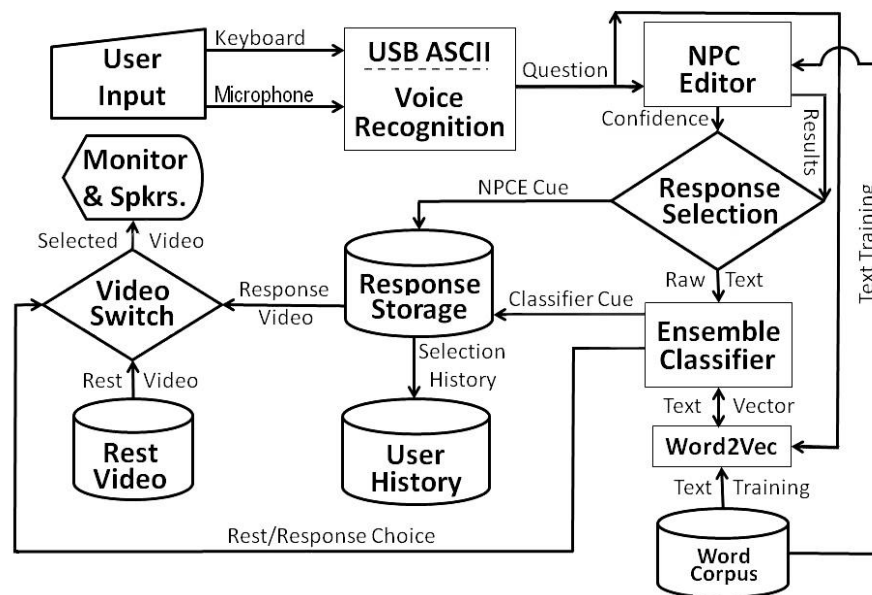


Figure 4- MentorPal Notional Flow Chart  
ICT's Virtual Mentor Implementation

## OPPORTUNITY FOR IMPLEMENTATION

Based on what the authors have observed in their professional careers and on what has developed with his research and several decades of further observation, the authors now turn their attention to problem areas that may be amenable to computer-enabled intervention. This paper takes the position that there are basically four areas in want of additional assistance:

- Providing guidance, counsel, and stress relief
- Articulating acceptable "field expedients" featuring a more free-form skills assessment and assignment;
- Helping to bridge the communications gap for personnel assigned around the globe.
- Reiterating the imperatives mandated by the national goals and the lives that are in the balance.

### 24 Hour Mentor Access

Using the technologies described above, it is possible to design, implement and field a set of mentors who are focused on assisting the prospective team members get through the roughest stages of new team formation. Some enhancements to current capabilities would be useful, but just implementing what is already been proven will have a markedly beneficial impact.

### Implementing a Virtual Conversation with Initiative

It is now possible to create a Virtual Conversation with initiate. A number of choices need to be made, but the decision is just a search of optimum effect and a wrong choice would impact efficacy, but would not be fatal to the concept. An example of such a choice would be an animated literal virtual human, speaking with synthesized voice driven by text. This is what was done in the ICT project called SimCoach. The other option, which was discussed above, is the utilization of a thousand or so video clips of live humans. This technique is often also called a "virtual human," but the purest point out the subject is not "virtual" at all, but a real, live human. This terminology difference will be addressed later.

There are two known hurdles to this vision. First: studying, evaluating, and collating the appropriate question data set. The authors have submitted a paper on this topic to another conference on how such a data set was created for the current project, called PAL3, A/I. A read-ahead copy should be available by the time of the SIW conference in February. Please contact Dan Davis, if you would like to read it (Davis *et al.*, 2020). Second, it has been the authors' observation that much of good mentoring is asking the salient questions. The second is that NLP is very good at answering direct inquiries, but it would be a new investigating effort to identify, draft, implement and validate a series of germane questions or probes to get the users to discuss and analyze their issues on their own, as well as a bank of considered suggestion to resolve common issues that are "bound" to come up.

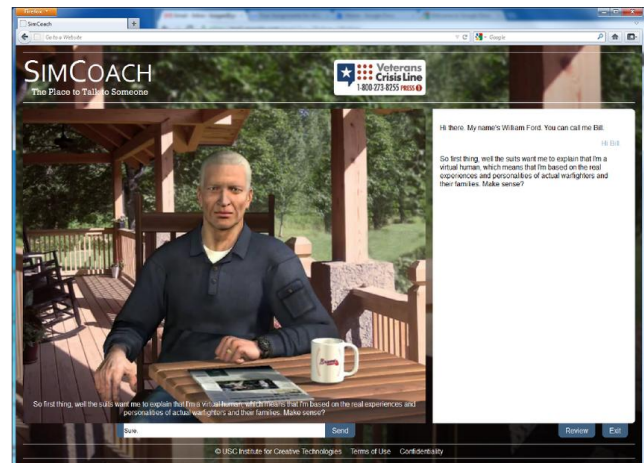


Figure 5 - SimCoach for Advising VA Patients

It is envisioned that part of the solution will be to urge the participants to set up and schedule regular video conferencing meetings. There are now a range of video conferencing programs available to the public at no direct cost. Most of them now support multi-user sessions and various forms of document sharing. The one issue that I think has yet to be resolved would be the participation of those relying in satellite connections: remote sites in areas not served by cable and ships at sea. The latencies imposed by the round trip to geostationary satellites can result in 900 to 1500 milliseconds, which makes conversation problematic, and just the speed of light limits adds at least 550 milliseconds to the response time (V-Sat Systems, 2019) and that is just enough delay to cause some minor

It is envisioned that part of the solution will be to urge the participants to set up and schedule regular video conferencing meetings. There are now a range of video conferencing programs available to the public at no direct cost. Most of them now support multi-user sessions and various forms of document sharing. The one issue that I think has yet to be resolved would be the participation of those relying in satellite connections: remote sites in areas not served by cable and ships at sea. The latencies imposed by the round trip to geostationary satellites can result in 900 to 1500 milliseconds, which makes conversation problematic, and just the speed of light limits adds at least 550 milliseconds to the response time (V-Sat Systems, 2019) and that is just enough delay to cause some minor

disruptive interference to the normal flow of conversation. Even so, it could be a very effective tool for resolving some of the counter-productive impacts of the Tuckman “Storming” stage.

### Standardizing Terminology

In an emerging technology such as Virtual Humans, one of the issues that continually hampers the achievement of consensus as to progress and the discovery of the optimal path for future research is the lack of consistency in both the terms used to describe the various techniques and the absence of carefully considered quantification of performance. As mentioned above, the term Virtual Humans is used in the technical vernacular to refer to any human avatar, be the image a true animation object portraying a Virtual Human or a really just a computerized delivery of a real live human. Comparative studies are needed to show the relative costs, time constraints, performance, and impacts of both implementations.

The call for more studies on performance brings to light the need to better quantify many factors now left to subjective assessments like face validity and anecdotal reports. The human behavior disciplines, e.g. social psychology, systems engineering and operation research rely on quantified measures of merit in order to produce the analyses that would facilitate future analyses and enable more optimal pursuit of goals that would result in actual benefit to the Warfighter.

The simulation community may be able to make good use of the experience of the standards discipline in managing the standards for various physical and process operation. Just this professional expertise could help the simulators in creating and promulgating and new *lingua franca* for Virtual Human research.

### NEW TECHNOLOGIES

Nearing the end of the Moore’s Law growth of digital computing, many simulation professionals are concerned with solving the grand challenges set forth above. One of the alternatives frequently mentioned is Quantum computing. It has hopefully been considered an extension of computational capability since the Nobel Laureate Richard Feynman presented the seminal paper in 1982. In that paper he held that: “... with a suitable class of quantum machines you could imitate any quantum system, including the physical world.” (Feynman, 1982). The authors have assiduously followed the development of such a machine and those devices still are, as near as can be ascertained, almost entirely at the test-bench phase. There seem to be no such “general purpose” quantum computer that is even nearing operation. There is one operational design in the quantum world: while not a general purpose quantum computer, it relies on very cold temperatures (15 milliKelvin) to create a useable quantum effect. (Lucas, 2013)

This adiabatic quantum annealing device has been conceived, designed, produced, and delivered to the University of Southern California. It has been in operation since 2012. In its current configuration, D-Wave computers have a design providing approximately 2,000 qubits. The “approximate” figure is required as some delivered machines have some fraction of these qubits turned off and a small number of the qubits (~ 1%) are not stable after the processor reaches the target 15 mK. Figure 5 shows the D-Wave Two, as installed in the USC-Lockheed Martin Quantum Computing Center (QCC) at the Information Sciences Institute (ISI) in Marina del Rey. There is another in the San Francisco Bay area in a joint Google and NASA project, and one has putatively been ordered by the DoD’s High Performance Computing Modernization Program. Others are in varying stages of procurement and delivery.



Figure 6 - D-Wave Quantum Annealer at USC


In recent years, other authors have touted quantum computing's ability to produce more power, using terms like "magic" to stir the imagination and whet the appetites of the user community. They point out that the capability of quantum computers arises from the different way they encode information. Digital computers represent information with transistor-based switches having a state of 0 or 1, labeled as a bit. In contrast, the basic unit of quantum computer operation, the quantum bit or qubit, can exist simultaneously as 0 and 1. A quantum bit, called a qubit, might be represented by an atom in one of two different states, which can also be denoted as 0 or 1. Two qubits, like two classical bits, can attain four different well-defined states (0 and 0, 0 and 1, 1 and 0, or 1 and 1). But unlike classical bits, qubits can exist simultaneously as 0 and 1, with the probability for each state given by a numerical coefficient. Data in a two-qubit quantum computer thus requires four coefficients. In general,  $n$  qubits demand  $2^n$  numbers, which rapidly becomes a sizable set for larger values of  $n$ . For example, if  $n$  equals 50, about  $10^{15}$  numbers are required to describe all the probabilities for all the possible states inside the quantum machine. That large number exceeds the capacity of the largest digital computer. A quantum computer should demonstrate incredible computational power because it can be in multiple states at once, a condition called "superposition." Also, perhaps more importantly, it can act on all its possible states simultaneously. The quantum computer can evaluate a series of optima that would be beyond the power of the largest digital cluster (Gershenhorn, 1998). Evaluating a large number of parameters to "learn" patterns and behaviors is at the heart of Virtual Human enhancement.

The authors have witnessed and participated in the development of high performance computing for several decades and have developed a significant body of experience with newly introduced technologies. They were engaged in the very early introduction of parallel computing and aware of its rivalry with sequential computing and with vector computing. They heard the detractors of parallel computing argue the limits of parallelism and the proponents who argued that it could be used more universally. While acknowledging there are many problems that have remained outside of the easily parallelized arena, it is evident that the majority of all large-scale computational problems are now run in parallel. This is due to the application of new techniques to decompose both data and computation in effective ways. Such technology has proven very useful to the simulation community, which has many issues identical to the test and evaluation environment. By using super-cold processors, the D-Wave has been able to demonstrate accepted quantum computing, see Figure 7 to the right. Even if the projected speed-ups are not realized on this design, it is a workable and verifiable quantum computing device.

The D-Wave machine in its developmental adolescence; it is no longer helpless, but its potential remains to be seen and characterized. This phenomenon was described by the late Professor Clayton Christensen in his series of books on disruptive technologies (Christiansen, 2011), *i.e.* a period in which the new technology is operational, but not yet competitive with technologies that have had decades of development and optimization. The criticisms and skeptical comments are familiar to all who have witnessed the displacement of older technologies. However, this community has also seen the claims of new capabilities that were then quickly set aside when old technologies improved faster than the new technology could prove itself. These issues are all part of this community's history and provide insights for analyzing the value of new technologies.

Several things need to happen before that D-Wave becomes a tool rather than a proof of concept device. The hard-wired connectivity paths need to be increased to represent a viable representation of the experimental environment sought to be simulated. A huge code base, both theoretical underpinnings and practicable software approaches, must be conceived and implemented. Perhaps decades of developmental advances will be needed to bring even this limited capability to its promise of revolutionary change. But this community will also remember the naysayers of the early days of personal computing and the skeptics of the first half of the 90's when the internet was first becoming available to the public at large.

Current plans at D-Wave include: more connectivity pathways, more qubits and, perhaps most critically, improved software in both the applications and the digital-to-quantum interface. The programmers at USC have used all of the D-Wave programs and become familiar with their Application Program Interface (API). D-Wave currently supports C/C++, Python, and MATLAB, three very common research languages. The company provides both on-line and live tutorial sessions to bring new users on-board. Much of the USC work focused on the treatment of large databases,



**Figure 7 - D-Wave Cooling Mechanism (D-Wave photo)**

known today as “big data.”.

## Deep Learning

Deep learning is another term that is used to describe an emerging technique in machine learning and artificial intelligence.

The authors take it to be an extension of earlier work (DeMuth, 2014) in the areas of neural networks, evolutionary computing, and data mining. But in deep learning, the refinement is done by several layers of convolution processing, in which some evaluation process sends along only such information that the layer values as beneficial. These successions of processing steps are called hidden layers and are the distinguishing feature of deep learning. Deep learning advocates note that this hierarchical approach to sifting large amounts of data is more in keeping with real world issues, albeit it's need for more compute power.

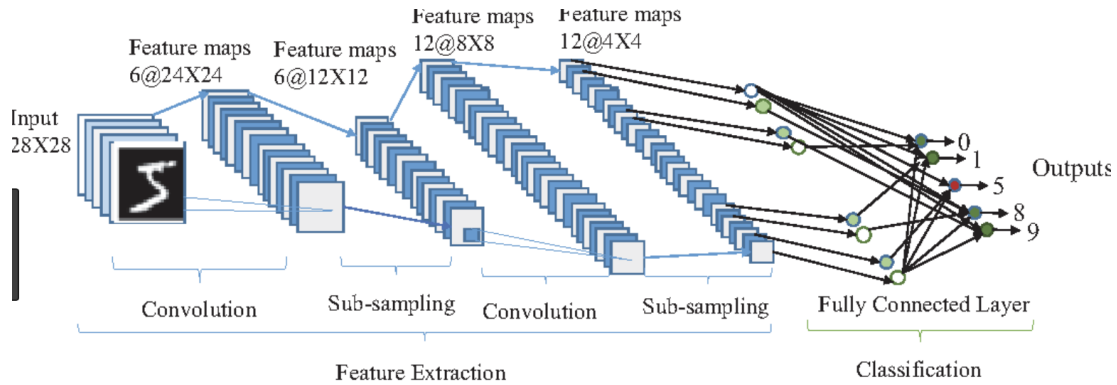
Deep learning has achieved state-of-art in performance in a wide range of problems, including hand writing recognition, object recognition, classification in images, speech recognition, understanding natural language text, and adversarial games, (LeCun, *et al.*, 2015). A key characteristic of being able to tackle these problems is the need to find complex structures in high-dimensional data. Deep learning with its multiple layered structure and distributed representation turns out to be well suited for these problems as compared with other machine learning approaches and other manual knowledge engineering approaches.

The current interest in deep learning can be viewed as the third wave of neural network development. In the late 1950's the Perceptron ® captured the public imagination as a new type of electronic brain (Bernstein, *et al.*, 1981). But, the interest in Perceptron waned when Minsky and Papert pointed out the limitations of a single Perceptron, which can only capture simple linearly-separable structures. In the 1980's, interest renewed as researchers constructed much more capable neural network models with multiple layers with multiple neurons per layer. But, again interest waned, because multiple layered networks are difficult to train and the computation power was not sufficient at that time. The current interest in neural network started in 2006 as Hinton *et al.* introduced a novel method to pre-train the network evaluation weights for deep belief networks using Restricted Boltzmann Machines (RBM) (Hinton, *et al.*, 2002). Moreover, the advent of graphical processing units (GPUs) with their Single Instruction Multiple Data (SIMD) capabilities reduced the computational bottleneck, often reducing weeks of computation to a few hours. GPUs are increasing in their power, but also face physical limits to their computational growth.

Of the three major deep learning approaches (feed-forward neural networks, recurrent neural networks, and Boltzmann machines), the Boltzmann machine approach is the most amenable to quantum annealing. Boltzmann machines are probabilistic neural networks that implicitly define probability distribution over the activation states of the neurons in the deep learning network. Training a Boltzmann machine requires being able to repeatedly sample from the distribution of activation states. However, sampling for Boltzmann network with loops can be computationally expensive.

The reason restricted Boltzmann machines can be trained efficiently is because they avoid edges within a layer, which enables layer-based approximate sampling of the network. With quantum annealing, there is the potential of efficiently sampling networks with loops. This enables the creation of a broader range of networks, with more complex topology. The current generation of quantum annealers does support complete graphs; they also impose limits on the topology of the intra-layer edges. This is the definition of limited Boltzmann machines (LBM), which are strict supersets of RBMs, and demonstrate the effectiveness of the additional edges.

The rise of deep learning is related to the rise of big data. Deep learning models require very large datasets to properly train the neural network weights. The activation of a neuron depends on the activation of its neuron neighbors mediated by the weights of the connection, see Figure 8. The winner of the 2012 ImageNet Large-Scale Visual Recognition Challenge AlexNet (Krizhevsky *et al.*, 2010) has over 60 million weights with 5 convolutional layers, max pool layers, and 3 fully connected layers. The network was trained on over 15 million images with over 22 thousand labeled categories. In 2014, Simonyan and Zisserman developed VGG net, which has 19 convolutional layers, and Microsoft ResNet (Jaderberg *et al.*, 2015) has 152 layers.



**Figure 8 - Diagram Showing Successive Layers of Deep Learning Analysis (Medium, 2019)**

Deep learning does have limitations and there are recent research efforts that are addressing these limitations. One limitation is brittleness. Deep learning models have demonstrated remarkable ability to perform well in noisy data, for example the ability to recognize objects in photos with busy backgrounds. But, paradoxically it is often quite fragile. Given new classes of objects to recognize, the model must be extensively retrained with many instances of the new objects as well as the previous known objects. In the games area, a deep learning-based Alpha Go program can beat top ranked human players, but the number of handicap stones is hardwired. Changing from the default handicap of 7.5 stones requires extensive retraining. An approach to address this limitation is to develop less monolithic deep networks, such as developing visual attention models to incrementally decode images and combining deep networks with artificial reasoning capabilities. Another limitation is the need for large labeled training dataset. Large datasets are often easy to find, but labeling is often a labor intensive process, *cf.* discussion of “sarcasm” recognition below. Researchers are developing unsupervised learning techniques, such as using Boltzmann machine-based auto-encoders; and incorporating artificial intelligence techniques (like reinforcement learning in the AlphaGo) to self-generate labeled data.

## ETHICAL ISSUES OF CONCERN

In their work on developing on-line mentors, the authors have recorded two instances of sensitivities not otherwise anticipated, despite the ICT teams having had one prior service officer observing all sessions. One was not particularly serious, but involved a Naval Academy graduate who refused to say one of the “repeat after me” words that were used as segue terms if an answer were not to be found. Such terms as “I am not sure I can answer that.” and you’ll have to ask me something else.” The term at which he balked: “Maybe.” He protested that commissioned officers never say maybe and was ostensibly offended the team would suggest that he do so.

Another was junior officer who responded, when asked to differentiate the differences between officer and enlisted ranks, opined that there was no difference. When this response was reported to several flag and general officers, the emotionality of the response indicated that this too was a very sensitive area. All those contacted responded with vociferous and emphatic analyses explicating why and how there was a difference and justifying their responses with voluminous historical examples.

A computer, not recognizing what had happened, might pursue a sensitivity to the point of severely injuring the user. There are examples of human insensitivity resulting in suicides by those offended and humiliated. The ethics of both humans and computer programs must be carefully examined and appropriate limits articulated and validated.

As one might expect, both of those sensitive issues were quickly recognized by the humans who were in the loop, but it is unclear how they may have been recognized by a computer. The authors propose two, not mutually exclusive approaches:

1. Using a Quantum Computing-enabled Deep Learning program, the computer could be given recorded clips of offensive statements from the entertainment industry (cinema and television

video files) and the Deep Learning programs set to isolating those indicia of an offended person, e.g. certain words or tones, body posture, physical attack,

2. Emerging sensor technology designed to assess user emotional states in body language, facial expression, and voice dynamics could be employed to help the computer both “learn” and to recognize angst in the user

Closely related to these issues are the issues of confidentiality. To be effective, the computers need to keep a database of approaches and efficacy. But the users need to have solid assurance that their confidences are as safe or safer with the computer than they would be with a human therapist.

As a way to assess the issues at hand, the authors suggest using a process not unlike the well-established requirements of an Internal Review Board and concomitantly set up an external ethics panel made up of ethicists, theologians, therapists and clinical psychologists. The goal would be to assure that any user could feel a peace with the steps taken to avoid the danger cited above, along with the more sensational dangers portrayed in movies such as 2001: A Space Odyssey and The Terminator.

## CONCLUSIONS

While there is yet much to be accomplished, the early progress in Virtual Human implementations has shown both the promise and the problems in this simulation discipline. This paper has attempted to lay out these by using the work of Professor Tuckman as an exemplar of the issues before the researchers.

### **Team Creation is one of the Issues that Require Virtual Human Conversation with Initiative**

There are several areas of potential use of virtual human conversation that require the computer agent initiate and guide the conversation and not just be responsive to the user. Some of these are team creation coaching, PTSD patient therapy, career counseling, and academic tutoring. The experience of the authors supports the proposition that team formation in large organizations is essential to the achievement of organizational goals. The fusion of disparate personalities, often initially located in site remote from each other, is critical to success. There appears to be a significant lag in team efficacy and this lag may be significantly different in civilian and in military environments.

### **Emerging Technologies and Capabilities in Simulation and Virtual Humans may be Effective**

In virtually every human endeavor, there is a need for mentoring and mediation in order to optimize both the timely achievement of goals and the optimal efficiency in that achievement. This is especially true in the military where goals have national survival implications and costs are human lives. Operations schedules and funding constraints may make the use of computer-generated human avatars more feasible, especially in the light of geographic dispersion of personnel that will form the envisioned team.

Advances Currently **are Hampered by a lack of Community Awareness**

For future research to be effective and rapid, a set of community standards needs to be defined and adopted. Much of the work centers on human behavior and a new area “Virtual Human Behavior.” These are areas in which there will be a need to compare capabilities and progress. This situation bespeaks a need for both community-accepted terminology and rigorously-defined quantifications. The existing standards professionals could be of invaluable assistance in this process.

### **Emerging Technologies have Potentials as Enablers of Virtual Humans with Initiative**

Virtual Conversation with Initiative is physically possible with current technology, but not practicable due to latencies imposed by digital computing limitations. Advances in Deep Learning and Quantum Computing hold promise for breaking the digital barriers and enabling a truly human-like conversationalist. Many benefits from this

implementation would accrue to the warfighters who are dispersed and time constrained in their interpersonal relationships with professionals for whom they have a need.

### **There are Ethical Issues that arise from a Computer that has Initiative**

Several ethical issue that will likely arise with the advent of a virtual human with initiative One of the more obvious and potentially problematic is the issue of sensitivity. Humans spend their lives generating carefully drawn sensitivities and a quick interview with a person rapidly gives a good sense of that person's awareness of topics that can be destructive to other individuals. There may be A/I approaches a not yet implemented that would address these issues, but the authors in no way represent this is a closed issue. Even the use of a quantum computer and Deep Learning techniques may take a long time to develop a trustworthy program that would sense the apprehension in a user's face and steer clear of objectionable comments. Naturally, many readers will point to similarly insensitive humans and it is not outside of the realm of possibility that humans will be less offended by a machine mistake that is putatively not malicious and a human mistake that may have been intended to injure.

The problems are real but the promise of new hardware and more sophisticated A/I techniques create a hopeful environment for the 21<sup>st</sup> Century.

### **ACKNOWLEDGEMENTS**

Virtually all of the work involving Virtual Humans reported in this paper came from a series of research projects of which the Principal Investigator was Dr. Benjamin D. Nye, Director of Learning Sciences at ICT. Without his technical vision and unflagging support, none of this would have taken place. The various projects mentioned were supported by a range of DoD organizations and, while this paper reports insights gained during that work that was not the focus of the research. Nevertheless, the positions taken in the paper are the authors' own and do not represent in any way the views of the Department of Defense or the US Government.

### **REFERENCES**

- Artstein, R., Traum, D., Alexander, O., Leuski, A., Jones, A., Georgila...& Smith, S. (2014). Time-offset Interaction with a Holocaust Survivor. In *Proceedings of the 19th international conference on Intelligent User Interfaces* 163-168ACMFebruary,
- Bernstein, Jeremy, (1981). A. I.. *New Yorker Magazine*, New York, New York,
- Christensen, Clayton. (2013). *The Innovator's Dilemma: when new Technologies Cause Great Firms to Fail.* Harvard Business Review Press, 2013.
- Crisp, G., Nora, A., & Taggart, A (2009). Student characteristics, pre-college, college, and environmental factors as predictors of majoring in and earning a STEM degree: An analysis of students attending a Hispanic serving institution. *American Education Research Journal*, pp 924-942.
- Davis, M.C., Nye, B.D& Davis, D.M. (2020). Populating a Virtual Conversation Database: Community Data Collection and Analysis In the *Proceedings of the ModSim World Conference*, Norfolk, Virginia, (Pending).
- Demuth, Howard B., (2014). *Neural Network Design*. Martin Hagan,.
- Feynman, R., (1982), Simulating Physics with Computers. *International Journal of Theoretical Physics* 21 (6-7): 467-488.
- Gershenfeld, Neil, and Isaac L. Chuang. (1998) Quantum computing with molecules. *Scientific American* 278.6 (1998): 66-71
- Hinton, Geoffrey E. (2002). Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation* 14.8: 1771-1800.

- Hornsby, JS. (2018). Critical Elements of Team Formation to Enhance Organizational Innovation In *The Challenges of Corporate Entrepreneurship in the Disruptive Age* (pp123-140). Emerald Publishing Limited.
- ICT (2018). *Research Project One-Sheets* Institute for Creative Technologies, University of Southern California. Retrieved on 04 (2019 from: <http://ict.usc.edu/research/pdf-overviews/> ..
- Jaderberg, Max, Karen Simonyan, and Andrew Zisserman. (2015). Spatial transformer networks. *Advances in Neural Information Processing Systems*.
- Krizhevsky, Alex, and Geoffrey E. Hinton. (2010). Factored 3-way Restricted Boltzmann Machines for Modeling Natural Images. *International conference on artificial intelligence and Statistics*.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015). Deep learning. *Nature* 521.7553: 436-444.
- Lord, G. (2019). In *Thespian Praise of: David Niven*, Paulburgin.blogspot.com25 January 2006Retrieved from the Internet on 02 Dec..
- Lucas, Robert F., Gene Wagenbreth, John J. Tran, David R. Pratt, and Dan M. Davis. (2013) Practical Adiabatic Quantum Computing: Implications for the Simulation Community. in the *Proceedings of the Interservice/Industry Simulation, Training and Education Conference*, Orlando, Florida, 2013
- Medium, (2019). *Deep Learning*. Downloaded from the Internet in November of 2019 from URL: <https://medium.com/tag/deep-learning/latest>
- Morbini, F., Forbell, E., DeVault, D., Sagae, K., Traum, DR., & Rizzo, AA. (2012). A Mixed-initiative Conversational Dialogue System for Healthcare. In Proceedings of the *13th Annual Meeting of the Special Interest Group on Discourse and Dialogue* pp 137-139.
- Norris, PE., O'Rourke, M., Mayer, AS., & Halvorsen, KE. (2016). Managing the Wicked Problem of Transdisciplinary Team Formation in Socio-Ecological systems. *Landscape and Urban Planning*, Vol 154, pp 115-122.
- Nye, B.D., Swartout, W., Campbell, J., Krishnamachari, M., Kaimakis, N and Davis, D.M. (2017). MentorPal: Interactive Virtual Mentors Based on Real -Life STEM Professionals. In the *Proceedings of the Interservice/Industry Simulation, Training and Education Conference*.
- Pelton, JN., & Jacqué, B. (2016). Distributed Internet-Optimized Services via Satellite Constellations. *Handbook of Satellite Applications*, pp 1-21.
- Takai, S & Esterman, M. (2017). Towards a Better Design Team Formation: A Review of Team Effectiveness Models and Possible Measurements of Design-Team Inputs, Processes, and Outputs In *ASME (2017 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*.
- Technogizer, (2019). ; Time.com *Dick Tracy's Watch: The Most Indestructible Meme in Tech Journalism* <http://techland.time.com/2013/02/11/dick-tracys-watch-the-most-indestructible-meme-in-tech-journalism/> .Retrieved from the Internet on 03 December 2019.
- Traum, D., Jones, A., Hays, K., Maio, H., Alexander, O., Artstein, R., ..& Jungblut, K., (2015). New Dimensions in Testimony: Digitally Preserving a Holocaust Survivor's Interactive Storytelling. In *International Conference on Interactive Digital Storytelling*, pp269-281. Cham, Switzerland: Springer.
- Tsai, J., Bowring, E., Marsella, S., Wood, W., & Tambe, M. (2012). A Study of Emotional Contagion with Virtual Characters. In *International Conference on Intelligent Virtual Agents* (pp81-88). Springer, Berlin, Heidelberg. September,.
- V-Sat Systems, (2019). *Latency- why is it a big deal for Satellite Internet?*. Retrieved from the Internet from: <https://www.vsat-systems.com/satellite-internet-explained/latency.html> on 02 December 2019
- Wikipedia, (2019), *HAL 9000*. Retrieved from URL: [https://en.wikipedia.org/wiki/HAL\\_9000](https://en.wikipedia.org/wiki/HAL_9000) on 02 December 2019b

