System Engineering: Optimizing Creation of Conversational Virtual Human Avatars

Daniel P. Burns
Home Ports Solutions
Los Angeles, California
daniel.p.burns@homeportsolutions.net

Dan M. Davis & Julianne Nordhagen
University of Southern California
Los Angeles, California
dmdavis@acm.org & nordhage@usc.edu

ABSTRACT

Many defense organizations report pressure to replace or augment scarce, expensive and schedule-constrained personnel who are needed to counsel personnel. One method that has demonstrable capability and has found utility is the use of conversational, computer-directed avatars. These can respond smoothly to a predetermined list of germane questions. These avatars can be either animated or live (via a large number of video clips). Early research and development has focused on "hand crafting" question lists, paraphrases, recording scripts, utterances, transcriptions, video editing, validation, testing, and updating. The administrative burden of all of this has begun to overshadow the scientific and technical research effort. As this virtual human capability becomes more acceptable and widely implemented, the skills of system engineering are seen as potential facilitators for optimizing the production process. This paper addresses early experience with enhancing an on-going project via system engineering. These analyses should make the adoption of these avatars more practical and economical. The authors discuss the use of the V-Model and other system engineering tools in a way that should enable other researchers to understand what can be expected of such an approach to computer-generated virtual humans. These management tools and techniques provide a real opportunity in the DoD to make manifest the early successes demonstrated in research settings. Optimizing just the transcription phase would save significant amounts of implementation time. The paper then lays out both early results and suggests future paths for research and use.

ABOUT THE AUTHORS

Daniel P. Burns is a lifelong Systems Engineer, first with the Active Duty Navy, then SAIC, and small business. He served as Naval Chair and Professor of Practice in Systems Engineering at the Naval Postgraduate School (NPS). Captain Burns served as the as the Military Associate Dean and as acting Dean of the Graduate School of Engineering and Applied Sciences at NPS. His research interests center on analyses of both human and resource utilization in defense efforts. Currently he is participating in the creation of a new program for Air Force Officers who seek post-graduate degrees. Captain Burns received a BS degree from the U.S. Naval Academy and an MS from the Naval Postgraduate School. He is currently finishing his dissertation for a Ph.D. from Portland State University.

Dan M. Davis is a consultant for the University of Southern California, focusing on large-scale distributed DoD simulations. Pre-retirement, for a decade he was the Director of USC's JESPP project for JFCOM. As the Assistant Director of the Center for Advanced Computing Research at Caltech, he managed Synthetic Forces Express, bringing HPC to DoD simulations. Prior experience includes serving as a Director at the Maui High Performance Computing Center and as a Software Engineer at the Jet Propulsion Laboratory and Martin Marietta. He has served as 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. He saw duty in Vietnam as a USMC Cryptologist and retired as a Commander, Cryptologic Specialty, U.S.N.R. He received B.A. and J.D. degrees from the University of Colorado in Boulder.

Julianne M. Nordhagen is an Ensign in the U.S. Navy and is attached to the NROTC unit at the University of Southern California, where she is assigned to collaborate on research at the Institute of Creative Technologies (ICT). As an Industrial and Systems Engineering (ISE) undergraduate, Julianne cultivated the analytical and technical skills required for this project, including optimizing the performance of large scale systems and applying computational methods for scheduling and forecasting. She has professional experience using integrated planning software like SAP and NetSuite ERP. She was selected to be one of the mentors for ICT's Mentor PAL project for the Office of Naval Research. Next April, she will begin her Naval Aviator training at the Naval Air Technical Training Center in Pensacola Florida. Ensign Nordhagen earned an ISE B.S. degree from the University of Southern California.

System Engineering: Optimizing Creation of Conversational Virtual Human Avatars

Daniel P. Burns
Home Ports Solutions
Los Angeles, California
daniel.p.burns@homeportsolutions.net

Dan M. Davis & Julianne Nordhagen
University of Southern California
Los Angeles, California
dmdavis@acm.org & nordhage@usc.edu

INTRODUCTION

This paper addresses the issues surrounding the use of System Engineering (SE) techniques, tools and approaches to enhance the conduct of simulation research, specifically Virtual Human implementation. It presents both an overview of SE disciplines and examples of how these have been employed in a program at the Institute for Creative Technologies (ICT) of the University of Southern California which is charged with creating a Virtual Human (VH) mentor series. There are two major objectives: analyzing the utility of SE and providing sufficient direction to enable others to see how to apply similar approaches to their own research projects. There will be sufficient examples and data presented to show application techniques and indicate results that can be anticipated.

The paper is organized in four parts. The first will be a background discussion of the advances made in the capabilities of VH technologies over the last few decades. The second section will cover the Mentor PAL project at ICT, setting forth the project's goals, implementation plans, early experiences, and current results. The third section will present a brief survey of System Engineers, focusing on areas where it will intersect with VH and other simulation initiatives. The fourth section will focus on analysis of the efforts to use SE for optimization in this environment and a discussion as to future implementation opportunities and efforts.

As this project is still part of an active research effort at ICT, new data is being collected daily and new tools and methods are being evaluated. The results and insights from these activities will be presented in the presentation at the conference and the slides from that presentation be made available on-line after the conference. It should be pointed out at the onset that this optimization work was not the directed focus of the project, but is an unfunded contribution of the staff members and of interested U.S. Naval officers, both active duty and reservists. While the authors acknowledge the lack of scientific rigor that would have been possible with a funded project, they believe that valuable insights have been garnered and they think that it is their professional obligation to share the lessons they have learned with their research colleagues.

BACKGROUND

Virtual Humans in the Defense Environment

Declining budgets, increasing operations tempos, and often under-staffed conditions continue to constrain the amount of time that experienced personnel have to train, advise and mentor those who follow them (Gould, 2017) Those same conditions create an increased need for such counseling to maintain readiness, improve morale and enhance retention of the Navy's most valuable asset: the sailors and officers (Payne & Huffman, 2005). The authors have observed that virtually every type of unit the Navy must face these issues. They assert that a method for standardizing and disseminating such advice globally and at any time on any day would make such information transfer more beneficial.

A virtual human is a virtual reality creation in which an avatar is created, often based on a real person, and attempts to recreate the appearance, voice, feel, and interaction that a live human would produce. With the advancement of several new technologies, including but not limited to natural language processing, virtual reality (VR), computer generated imagery (CGI), machine learning, and virtual learning, the uses, as well as the limits of virtual humans are becoming evident. As a research institute with a specialization in virtual reality programming, the Institute for Creative Technologies (ICT) is the home of myriad simulation projects, and it considers virtual reality to be its primary focus. Researchers at ICT have generated SimCoach, New Dimensions in Testimony (NDT), PAL3, MentorPAL, and other generalized programs under learning sciences, medical VR, mixed reality, narrative storytelling, social stimulation, virtual humans, and vision and graphics. The breadth of knowledge available is significant and the use of these resources has allowed advances in the implementing of virtual humans. The presentation to the user can take many forms, as shown in Figures 1-3 (All ICT Photos).

Figure 1 – Video of NDT speaker shown in 3-D holographic display

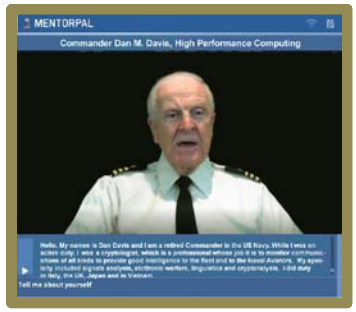

Figure 2 – MentorPAL video clip presented on a 2-D monitor

Figure 3 – Fully Animated CGI in SimCoach

Although a virtual human may seem as simple as remodeling a human using CGI, it turns out that it takes significant study and effort to implement a Virtual Human, and the process can consume considerable computing power to do so effectively. The essential elements that go into the creation of a virtual human with lifelike abilities include natural language processing, machine learning, VR, CGI, and social stimulation. Natural language processing (NLP), will be the main focus of this discussion, though the same argument concerning the limits of virtual humans can be made with several of the other components. Natural language processing composes "an area of research and application that explores how computers can be used to understand and manipulate natural language text or speech to do useful things" (Chowdhury, 2003). Using this definition within the context of virtual environments, NLP tools allow computer technology to recognize voice input, analyze voice tone, provide lifelike conversation, retrieve information, and many other applications in combination with machine learning. Recent developments in NLP have made significant advances, like "a single convolutional neural network architecture that, given a sentence, outputs a host of language processing predictions: part-of-speech tags, chunks, named entity tags, semantic roles, semantically similar words and the likelihood that the sentence makes sense (grammatically and semantically) using a language model" (Collobert & Weston, 2008).

Many applications, through interpretation of language, have been quickly advanced by the researchers focusing their studies on the field. At ICT, progress has been made in training and learning environments (Kenny *et al.*, 2007), multi-party dialogues (Traum & Rickel, 2002), ethics and cooperation (Allwood *et al.*, 2000), health applications (Rizzo *et al.*, 2011), and representation and reasoning (Swartout *et al.*, 2006). Although automated speech recognition (ASR) is far from perfect, prime software makes Virtual Humans practicable, "Google achieved 73.3% of exact recognized phrases with a 15.8% (Word Error Rate)" (Kudryavtsev, 2017), and the technology will continue to improve. It is also important to note that many of the errors in ASR are caused by slurred speech, cultural slang, and context, stemming from a "lack of consistent units of speech that are trainable and relatively insensitive to context" (Lee, 1988). Although this causes problems when comparing these transcriptions with global data, models can be trained locally to a particular person's voice or a low resource language, using software such as CMU Sphinx from Carnegie Mellon University as applied to languages such as Arabic (Satori *et al.*, 2007). Although these customizations and build-your-own languages can be more accurate, they take time to implement and often require increased local processing power.

As it stands now, the primary issues with natural language processing include machine translation, precision, data storage, efficiency, and computation power; meeting these are the foci of this paper. All of these will continue to improve with time, as hardware, software, data storage and ease of access are areas within which new research is emerging and seems directly applicable. Specifically, the speed and application of quantum computing will enable significant advances in NLP and its applications: one of which is the critical area of enhancing virtual humans.

Virtual humans constitute-and should continue to constitute-a major role in virtual learning, virtual storytelling, VR/AR, and constructive simulation. Within the military context, virtual learning environments provide useful mechanisms for initial training as well as lifelong training. An example of this in implementation can be found in PAL3; "the PAL3 system was designed to accompany a learner throughout their career and mentor them to build and maintain skills" (Swartout *et al.*, 2016). Modern learning calls for new methods of information transfer. Online tools like Khan Academy, YouTube, Coursera, Lynda, and other massive open online courses (MOOCs) are growing in popularity, and the trend does not seem to be slowing down. Many of these tools suffer from the lack of the immediacy and motivational qualities of personal interaction with a human.

Storytelling and gaming are also important facets of the virtual field. The New Dimensions in Testimony project "allows people to have an interactive conversation with a human storyteller (a Holocaust survivor) who has recorded a number of dialogue contributions (Figure 4), including many compelling narratives of his experiences and thoughts" (Traum *et al.*, 2015) . The project's mandated mission of preserving the stories of Holocaust survivors can be applied to other important persons and historical sagas of value and interest. Advances in gaming allow for utilization in learning, entertainment, healthcare, and lifelike training for areas like combat, base security, reaction team strategy, medical emergency responses, interfaces with civilian populations, and critical thinking. The SimCoach system was developed to assist Post Traumatic Stress Disorder (PTSD) subjects and to "motivate users to take the first step – to empower themselves to seek advice and information regarding their healthcare" (Rizzo *et al.*, 2011). These virtual systems have been shown to have more success, generating deeper levels of confidence with patients than even live healthcare interactions (Rizzo *et al.*, 2011). The successes of these operations thus far encourage research into potential future applications.

Figure 4– Subject receiving direction in light stage used for generating 3-D imagery for holographic display

The benefits of virtual environments include making the most skilled instructors and resources available and scalable, engaging new learning environments, confidentiality, and global accessibility. Some of the negatives of virtual learning are the difficulties of personalizing learning to fit individual students, handling live question, responding to different learning styles, and similar limits to NLP and other processor and data heavy systems. On the positive side, this approach does allow for using only the best instructors and can avoid the classroom stricture against teaching to the lowest common denominator, having to respond to questions in a generic way. It can be trained to detect and cater to every student's learning style, and other constraints faced by live teachers. Future fields of focus will include streamlining the creation of such virtual environments and interaction. For example, it is possible to conceive a system which takes in a live recorded interview with a U.S. President, or other prominent or knowledgeable figure, and generates an interactive environment in which a user can experience the interview by asking the questions themselves. Or a true virtual assistant that obviates the need for an assistant might provide a radical shift in the way the DoD functions. Such tools are already in existence, and they will only become more prevalent and powerful. However, the need for extensive processing power, efficiency, and data storage and transfer remains a limiting factor for many such developments. The aforementioned applications call for considerable further exploration and research. They may well be amenable to a quantum annealing approach to deep learning enhancement.

MENTOR PAL

One of the needs identified by the DoD is the increasing need for technical personnel, both in uniform and in a wide range of civilian capacities. To meet that need, potential candidates putatively would be well served to have access to thoughtful and experienced technical personnel who can clearly communicate the careers that may be offered to them (Nye, 2017). There are several factors that militate against that availability of effective mentoring. Some of these factors were mentioned by education leaders who have given input into the Mentor PAL project. While anecdotal, they are mentioned here. Many of the counselors noted that a large portion of the population lives in widely dispersed small school districts around the country, many students live in lower Socio-Economic Status (SES) districts in which few technical professionals live, many students are not amenable to face-to-face mentoring by professionals they find intimidating, and many who are available for mentoring have agendas of their own or are hampered by poor mentoring skills.

ICT proposed to the Office of Naval Research that most of these problems could be ameliorated by using their conversational VH technology and providing compelling, engaging and validating mentoring which could be delivered by locally hosted programs or web-delivered client/server service. It was proposed as being able to provide conversation-like mentoring to any student in any locality, in a non-threatening way, and by mentors who have been vetted as being "on-message" and personally engaging and insightful. The program has been funded and initial mentor videotaping and interface design has produced a working prototype. Various versions of this prototype, running on a Microsoft Surface tablet computer, have been assessed by a small number of pre-college students. Input from the students affirmed both the need for the mentorships and the openness to the on-line mentor.

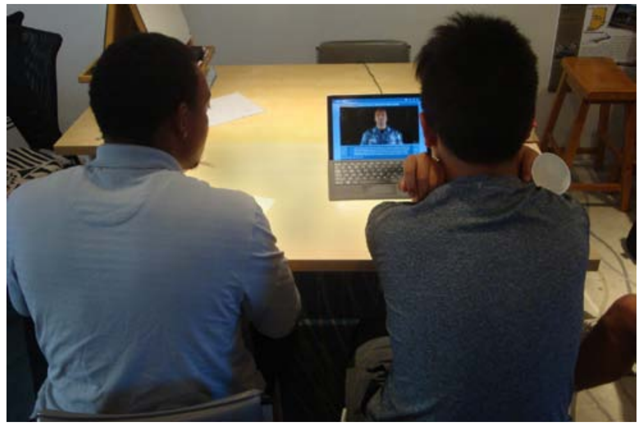

Figure 5 – Two Senior High School Students Evaluate Mentor PAL Advice from a Navy Chief Petty Officer

The project was funded to produce a prototype for analysis, but the early trial tests with students, as represented above, emphasized the range of evaluative comments and the subjective nature of the researchers reaction to those comments. Therefore, before the last student exposure to the program, a way of garnering, documenting, and quantifying the responses was sought. With insufficient time and funding to do a formal survey at this point in the project, one of the researchers created a quick five item Likert survey to capture and record some of the insights that the students had to offer. Below in Tables 1 you will find an image of the four and a quarter inch by five and half inch form that was given all of the students who sat and interrogated the VH mentor. Table 2 has a compilation of the answers given. All 16 participants responded to all five questions. Ten either wrote comments or gave oral feedback of note to the researchers.

Table 1 – Form Given to Students

Responses from Users Solicited

Please Circle the number that best describes how you feel.	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
1. I think or worry about my future life every day.	1	2	3	4	5
2. Having easy access to better advice would be useful.	1	2	3	4	5
3. I liked this on-line mentor and learned from him.	1	2	3	4	5
4. The answers were on-target and useful.	1	2	3	4	5
5. The responses seemed "conversational" and real.	1	2	3	4	5

Table 2 – Record of Responses

Results from Users, N=16

Please Circle the number that best describes how you feel.	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
1. I think or worry about my future life every day.	1	2	2	11	
2. Having easy access to better advice would be useful.			2	10	4
3. I liked this on-line mentor and learned from him.		1	3	6	6
4. The answers were on-target and useful.		1	3	6	6
5. The responses seemed "conversational" and real.		1		7	8

The authors recognize the informal nature of this data collection effort, but it did help record the impressions of the users and also acted as a documentation aide to assist the researchers in remembering the users' impressions. Both these data and the concomitant oral comments seem to support earlier research, cited above, that reflect both the ease of use and the credibility that the users afford to the VH mentors. Using this tool, instead of the time-consuming method of individual mentoring, is seen by the researchers as holding some promise for alleviating the constraints enumerated above. However, there is a cost: the production time of each videotaped mentor.

The procedure for creating such a conversational VH mentor involves a number of steps, in addition to the development and tuning of the software code that makes the program function. Literally volumes have been written about how to system engineer large code programs, running all the way from general, almost philosophical approaches, e.g. Professor Fred Brooks book (Brooks, 1995) to more detailed and didactic tomes (Pressman & Maxim, 2005). However, for the more prosaic support functions, no such tomes are familiar to the simulation computer scientists that would assist in organizing their activities. Project personnel, basing their approach on lessons learned from previous ICT efforts, knew they needed a process for the production of the videos. Therefore, with more courage than confidence, they set up a series of steps to produce the videos: familiarization with the career field to be covered, generation of questions to be posed to the mentor, selection of mentors, setting up the recording facilities, doing the recording, tearing down the facilities, editing the videos, time-stamping the resultant files, correcting voice recognition transcriptions, and validating the results.

These are represented in column one of Table 1 below. The rest of the columns show time estimates, in staff-hours, by two researchers, offering different configurations of the evolutions.

Table 3 – Task List and Estimated Times based on Producing Two Mentor Video Series.

Estimated time to produce a Mentor	Units = Hours			
	N.J. Kaimakis	D.M. Davis	DMD 4X Staffing	DMD 2nd Mentor
Study the new topic including interviewing professionals to help design questions	15	60	60	20
Create a new question list, including paraphrases and rankings of questions	35	80	80	
(Need 60 new questions out of 360 total for all subsequent Mentors)				20
Identify, select, recruit and train a single talking head/mentor	20	40	40	40
Set up/tear down the recording facilities (at USC or remote) first time - 1-2 hours	1	2	2	2
subsequent times 0.25 - 0.5 hour per session	1.25	2.5	2.5	2.5
Do the recording of the subject mentor ~360 - 400 questions (mentor reads own questions)	18	18		18
30 questions per hour, 2- 3 minutes per question, includes getting into standard clothing				
Recording if question read by researcher, with up to two others monitoring (2-4X staff hours)			72	
Edit the resultant video file 1-4 hours of work for 6 hours of video	3	12	12	12
plus 1 hour per 1 hour of video for after effects to render green screen / other edits	18	18	18	18
Timestamp video files for processing 1 hour per 2 hours of recording	9	9	9	9
Fixing transcriptions after voice recognition processing 2 hours per 1 hour of recording	36	36	36	36
Validating questions at various stages of development by presenting to sample users (4 tests)	16	16		16
Validation if sessions are conducted by additional researchers (2-4X staff hours)			64	
Totals	172.25	293.5	395.5	193.5

It was at this point that those actually doing the production began to notice and decry the significant amount of time that was being invested in very tedious and unrewarding assignments. Optimization was done *post hoc* and *ad hoc*, as the researchers had not thought that the small number of iterations justified a lot of time being invested in the process to enhance the production. However, it was realized that implementation by the Navy for a plethora of professions would benefit from such an effort; further, the success of any consideration for other counseling uses would be more promising if such advocacy were to be accompanied by quantification of the price of production. It was at this point that the decision was made to seek out the input from the System Engineering community.

.

SYSTEM ENGINEERING

The Discipline of System Engineering

Definition of Systems

Definition of Engineering

Definition of System Engineering

Goals of a System Engineer

Relationship of System Engineering to Operations Research

Misuse of the Term System Engineer

History of System Engineering

Future of System Engineering

Paucity of System Engineering in Research Project

Extensibility of System Engineering Concepts into Research

Accessibility of System Engineering tools by non-System Engineers

A Short List of System Engineering Tools:

The “V-Model” – grasping the life cycle

Stakeholder Diagram – grouping interested parties into mutually focusing groups

Pugh Matrix – plotting design criteria satisfaction against different design options

Context Diagram – graphing the mutual impacts of external entities and the system

The Analytic Hierarchy Process – quantifying the decision-making process

Articulating Requirements – translating customer desires into technical system requirements

System Requirements Model – putting requirements into a framework

Requirement Relationship Analysis – structuring the relationship of the requirements

Design Structure Matrix – identifying subsystem connections

N^2 Analysis – understanding system interconnection problems

Functional Modeling – establishing functionality and needed interconnections

Implementation Means Analysis – enumerating implementation alternatives

System Requirements Analysis – defining system requirements in a rigorous way

Requirement Means Analysis – studies the relationship of requirements to alternative solutions

Failure Mode and Effects Analysis – listing potential causes and effects

SE Tools Applicable to Mentor PAL

V-Model

Stakeholder Diagram

Context Diagram

Articulating Requirements

System Requirements Model

Implementation Means Analysis

Failure Mode and Effects Analysis

Current Application Efforts

Impact of SE on Mentor PAL to Date

ANALYSIS

When this was all observed and considered, it seemed that there were good reasons to consider the application of some of these techniques and approaches in other research projects. Further, larger projects with even more detailed and onerous procedures, it may warrant the recruitment of a System Engineer as a member of the researc

CONCLUSIONS

The authors assert that the data presented above justify the careful consideration of the adoption of some of the tools described.

ACKNOWLEDGEMENTS

The authors want to acknowledge the leadership and support from the two co-Principal Investigators, William Swartout, PhD and Benjamin Nye, PhD. Their counsel for us and patience with us were invaluable. They contributed ideas and data that were included in the data listed above. Much of the work described above was conducted in response to Office of Naval Research contract: Mentor PAL: Growing STEM Pipelines with Personalized Dialogs with Virtual STEM Professionals, N00014-16-R-FO03. The opinions expressed herein are the authors' own and do not necessarily reflect those of the Department of the Navy or the U.S. Government.

REFERENCES

Allwood, J., Traum, D., & Jokinen, K. (2000). "Cooperation, Dialogue and Ethics." *International Journal of Human-Computer Studie*s, 53(6), 871-914.

Brooks Jr, Frederick P. (1995), "*The Mythical Man-Month*." Anniversary Edition, Addison-Wesley Professional, Boston, MA.

Chowdhury, G. (2003). "Natural language processing." *Annual review of information science and technology*, 37(1), 51-89.

Collobert, R., & Weston, J. (2008, July). “A unified architecture for natural language processing: Deep neural networks with multitask learning.” In Proceedings of the *25th international conference on Machine learning* (pp. 160-167). ACM.

Gould, J. (2107), Military chiefs warn of force-readiness struggles, *Defense News*, Retrieved 15 January 2018 from: https://www.defensenews.com/congress/2017/02/07/military-chiefs-warn-of-force-readiness-struggles/

Kenny, P., Hartholt, A., Gratch, J., Swartout, W., Traum, D., Marsella, S., & Piepol, D. (2007, November). “Building Interactive Virtual Humans for Training Environments.” In Proceedings of the *Interservice/Industry Simulation, Training and Education Conference*, Orlando, Florida, 2007 (Vol. 174).

Kudryavtsev, Andrey. "Automatic Speech Recognition Services Comparison." *Grid Designer's Blog*. N.p., 11 Jan. 2016. Retrieved on 20 Apr. 2017 from: http://blog-archive.griddynamics.com/2016/01/automatic-speech-recognition-services.html.

Lee, K. F. (1988). “Automatic speech recognition: the development of the SPHINX system” (Vol. 62). *Springer Science & Business Media*.

Nye, B., Swartout, W., Campbell, J., Krishnamachari, M., Kaimakis, N. and Davis, D. (2017). "MentorPal: Interactive Virtual Mentors Based on Real-Life STEM Professionals." in the Proceedings of the *Interservice/Industry Simulation, Training and Education Conference*, Orlando, Florida, 2017

Payne, S. C., & Huffman, A. H. (2005). A longitudinal Examination of the Influence of Mentoring on Organizational Commitment and Turnover. *Academy of Management Journal*, 48(1), 158-168.

Pressman, R. S. & Maxim, B. (2005). *Software Engineering: a Practitioner's Approach*. Palgrave Macmillan, Basingstoke, UK, .

Rizzo, A., Lange, B., Buckwalter, J. G., Forbell, E., Kim, J., Sagae, K., ... & Parsons, T. (2011). “SimCoach: an Intelligent Virtual Human System for Providing Healthcare Information and Support.” *International Journal on Disability and Human Development*, 10(4), 277-281.

Satori, H, Harti, M., & Chenfour, N. (2007). “Introduction to Arabic Speech Recognition using CMU Sphinx System.” *arXiv preprint* arXiv:0704.2083.

Swartout, W. R., Gratch, J., Hill Jr., R W., Hovy, E., Marsella, S., Rickel, J., & Traum, D. (2006). “Toward Virtual Humans.” *AI Magazine*, 27(2), 96.

Swartout, W. R., Nye, B. D., Hartholt, A., Reilly, A., Graesser, A. C., VanLehn, K., ... & Wang, L. (2016, March). “Designing a Personal Assistant for Life-Long Learning (PAL3).” In *The Twenty-Ninth International Flairs Conference*.

Traum, D., & Rickel, J. (2002, July). “Embodied Agents for Multi-party Dialogue in Immersive Virtual Worlds.” In Proceedings of the first *International Joint Conference on Autonomous Agents and Multi-agent Systems*: part 2 (pp. 766-773). ACM.

Traum, D., Jones, A., Hays, K., Maio, H., Alexander, O., Artstein, R., Debevec, P., Gainer, A., Georgila, K., Haase, K., Jungblut, K., Leuski, A., Smith, S., & Jungblut, K. (2015, November). “New Dimensions in Testimony: Digitally Preserving a Holocaust Survivor’s Interactive Storytelling.” In *International Conference on Interactive Digital Storytelling* (pp. 269-281). “Springer International Publishing.