

A final project presented to the faculty of the
Instructional Design Masters Degree Program
University of Massachusetts at Boston

Measuring the Effectiveness of Animated vs. Static Learning Agents

Submitted by
Judy Unrein

in partial fulfillment for the requirement of the degree

MASTER OF EDUCATION

May 2011

Approved by Dianne M. Nerboso

Faculty

For my boys... all of them.

Judy Unrein

M.Ed. Instructional Design Graduate Program

Capstone: Measuring the Effectiveness of Animated vs. Static Learning Agents

May 2011

Abstract

This study aims to contribute to a better understanding of the effects of animated versus static learning agents when used in elearning courses. It measures effects on both learner engagement and ability to recall course information.

Two versions of a course were developed that were identical except for the learning agent, an onscreen “character” who presents the information and simulates some interaction with the learner. In Version 1 of the course, the learning agent was fully animated, including lip movement and head and hand gestures. In Version 2 of the course, screen shots of the agent were used in place of the animation. Both versions of the course featured human narration.

The course tracked recall using a post-test immediately following the course; engagement was tracked using a short survey following the post-test as well as a tracking system built into the course that recorded each page learners navigated to and each external link or interactive element they accessed.

Results from the 54 study participants were very consistent throughout; overall, neither version of the course showed a significant advantage in terms of engagement or recall. Given that this study reflected only one of many ways learning agents could be designed and developed, further research is encouraged to shed more light on this topic.

Table of Contents

Introduction	5
Design Considerations	6
Feedback Methodology	9
Objective Results	10
Subjective Results	15
Conclusions	16
Future Extensions of this Research	17
Bibliography	19
Vita	21

Introduction

Elearning offers significant advantages to corporate training departments that are faced with economic pressure to roll out training as cheaply as possible and organizational pressure to deliver consistent training to global audiences. However, a learning experience entirely facilitated by computer does not provide one of the major advantages of instructor-led training: interaction with other humans, particularly an instructor. Instructional designers try to compensate for this in a variety of ways; one of them is the inclusion of an instructional character — a learning agent — in elearning courses.

The use of a learning agent to present content and simulate interaction with the learner has been demonstrated to increase recall and engagement in some contexts and with some audiences. For example, Roxana Moreno and Richard Clark's research, "The Case for Social Agency in Computer-Based Teaching: Do Students Learn More Deeply When They Interact With Animated Pedagogical Agents?", reports that students learn more when an onscreen agent is used than when one is not used. Byron Reeves reports in "The Benefits of Interactive Online Characters" that these characters are perceived as real and are responded to as such by human computer users in both educational and commercial interactions. (Please note that while learning agents are often called *animated pedagogical agents* in academic research, the term *learning agent* will be used in this report to prevent confusion, as both animated/moving and static/unmoving agents are discussed.) However, learning agents can be used in a wide range of circumstances, for a wide range of purposes, and can be created using a wide range of multimedia. Research on the effectiveness of different uses, purposes, and development methods — not to even mention the potential influence of different audiences and different types of content — is far from complete.

One of the areas of development methodology in which research is somewhat sparse and at times conflicting is in the emotional and cognitive effects of fully animating the learning agent versus representing him or her as a static image on screen. Tools designed to create animated characters, complete with expressive hand and head gestures and lip syncing, are often marketed to elearning designers as easy ways to improve learner engagement and retention of content; however, they can carry hefty price tags and/or dramatically increased production times over using static images. Without clear research either supporting or calling into question the use of an animated agent for the purpose intended within a specific course, instructional designers and the organizations that hire them have very little basis upon which to make an educated decision to adopt this method.

This project examined the effects on the learner's engagement and content recall when the learning agent is fully animated versus static. Though this (and all) research has limitations, it is hoped that this study provides contributes toward a more solid basis of informed decision-making in selecting design and development approaches for learning agents in elearning.

Design Considerations

Study Design

The study utilized two versions of an elearning course: one with a learning agent that is fully animated and one with a learning agent that is a static image. The versions were otherwise identical to eliminate any outside factors influencing the results of the research. Each version of the course included the following sections:

- Demographic survey
- Course content
- Post-test
- Engagement survey

Using variables, the course delivered either Version 1 or Version 2 to each learner randomly. The course was delivered in SCORM-compliant format inside of a Learning Content Management System (LCMS) to facilitate registration, course delivery, consistency of course version even if the learner moved from computer to computer, and tracking and reporting of learner activity.

Course Design and Development

For the purposes of the study, it was desirable to generate enough interest in the course to attract volunteer participants who would register and take the full course, thereby delivering as much data as possible. However, it was also necessary for the success of the study to measure varying levels of recall and engagement. Therefore, it was determined that a design that maximized recall and engagement might not serve the study as well as one in which demotivating factors were included in order to ensure that not all participants attained the highest levels of recall and engagement. Another reason for including demotivating factors was to more closely mirror elearning taken in a corporate environment, which may include content that is dry, technical, or required but not relevant to the learner's job on a daily basis.

The course was designed to balance these competing needs by reflecting strategic design decision-making in the following areas:

Course topic and content selection. The course topic, Using Social Media for Your Job Search, was chosen to attract volunteer participants who were motivated enough to spend the time required to complete it. Originally this project was a partnership with the non-profit organization Hiring for Hope to remake some of the content that they offer to job searchers in other forms; in addition, independent research and expert interviews were undertaken to flesh out the course content. However, in the end the depth of content was kept fairly shallow throughout the course to limit engagement — in essence, to limit the rewards of exploration.

Visual design. Good visual design can increase engagement, as well as learner perception of a course and its designers. Because engaging visuals were not needed to promote participation, the interface was kept fairly bland to limit engagement.

Navigation. Unclear navigation can be a huge demotivating factor in elearning; however, the navigation is kept simple and fairly clear so as to prevent the learner from giving up before course completion. Also, since the course tracks the use of pages and links as a measure of engagement, it is undesirable for the learner to navigate around the course for any reason other than to further explore the content presented.

Narration and onscreen text. The relationship between narration and onscreen text has been studied in some depth, and it is now generally considered an ineffective practice to have the narrator “read” the onscreen text verbatim. The course includes some of the “reading from the screen” approach toward the beginning in an attempt to limit engagement. However, a human narrator rather than text-to-speech is used to provide a more human-like experience — a factor that should promote engagement.

Assessments and surveys. These are kept to a minimum to encourage full participation, as they comprise a large part of the study feedback mechanisms. Because the topic is non-technical and possibly already familiar to the audience, the assessment mainly tests recall of specific information from the course, to limit the applicability of prior knowledge.

The creation of the learning agent required audio recording of a narration script, which was then imported to an animation program to create the animated learning agent. Lip syncing was created automatically by the program; physical gestures were added by the program and tweaked and changed by the course developer.

For the static version of the learning agent, the same audio was used on the same screens, and screen shots were taken from the animated learning agent to ensure that the static agent only reflected gestures and poses that were also used in the animated version. The screen shots for the static version were placed at the same onscreen locations as the videos for the animated version.

Feedback Methodology

The goals of this study were to measure both recall and engagement. To measure recall, a five-question post-test covering specific information in the course was employed. While this is the most basic level of retention, it was all that was practical to cover in the study (and often, this level is the most that is done in corporate learning environments, as well).

To measure engagement, both of the following were employed: 1) a three-question engagement survey following the course, and 2) a tracking system built into the course to record how much the learner explored, accessing optional pages in addition to the main navigation, links external to the course, and interactive elements within the course. Every time the learner clicked on a new page, clicked an external link, or used an interactive element, the course added a specific value to a variable string that would be stored by the LCMS. The values were then tabulated and compared across versions. Thus, we can observe which assets were explored by which learners and, more to the point for this study, *how many* assets were accessed by learners taking each version of the course. In order to generate a meaningful measure of how many assets were used, the course navigation was kept as optional and open as possible. All attempts to take the course were counted, even if the learner did not finish the entire course; it is assumed for both versions that unfinished attempts

reflect the learner's "voting with their feet" and exiting the course because some or all elements of the course failed to engage them.

The engagement survey contained three items with responses based on a Likert scale.

The items were as follows:

- Item 1: This course kept my interest.
- Item 2: I would be interested in taking other courses developed by the same company.
- Item 3: I would recommend this course to others who are looking for work.

Objective Results

Course Versions

Of the 54 participants, 27 received Version 1 of the course (with the animated learning agent) and 27 received Version 2 (with the static learning agent). For seven of the 27 who received the animated version, the LCMS registered no interactions whatsoever, apparently due to technical glitches. Follow-up with some of these study participants has revealed that they took the course at least once and at least past the first page, so some values should have been reported. It is unknown why this error in reporting happened and why it was limited to only the animated version of the course. The higher bandwidth demands of downloading video rather than images and audio do represent one difference in the delivery of the different course versions, but whether the video downloaded should have had no effect on the variables registering with the LCMS, as that is a function of the HTML-based SCORM wrapper.

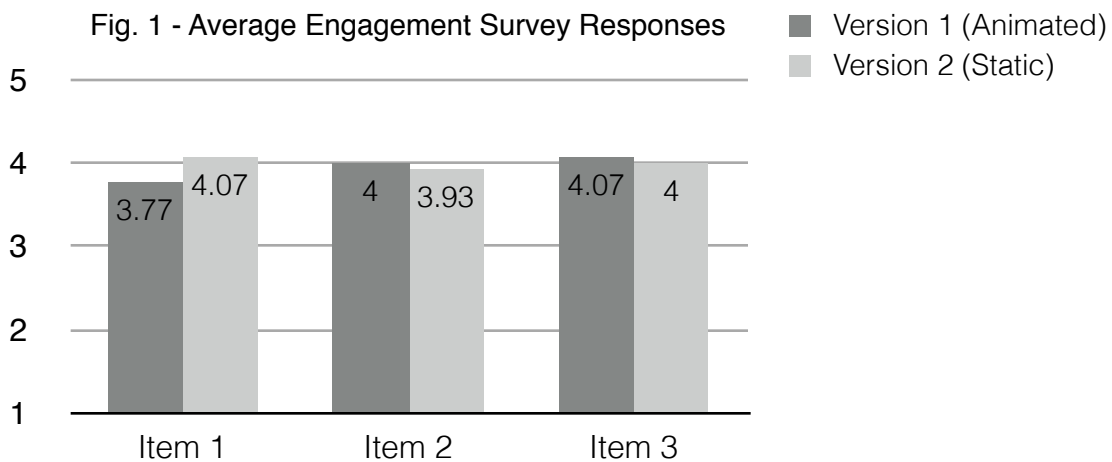
Demographic Breakdown

The results of the demographic survey follow. For any category that does not total 54 participants, the remaining participants declined to specify that information or no tracking information was recorded for that course attempt.

- Gender:
 - 17 were male.
 - 29 were female.
- Age:
 - None were under 18.
 - One was between 18 and 25.
 - 19 were between 26 and 40.
 - 20 were between 41 and 60.
 - Six were over 60.
- Ethnicity:
 - 41 were white.
 - One was black.
 - Two were bi- or multi-ethnic.
- Highest level of education:
 - One had completed some high school.
 - None had completed high school only.
 - Four had completed some college or an Associate's degree.
 - 10 had completed a Bachelor's degree.
 - 21 had completed a Graduate degree.
 - 10 had completed a Postgraduate degree.
- Employment:
 - One was not looking for work, as a student.
 - Four were not looking for work, as retirees.
 - 19 were not looking for work, as they were employed.
 - 16 were looking for work, as they were underemployed or looking for a change.
 - Three were looking for work, as they were unemployed.

Engagement Measures - Survey

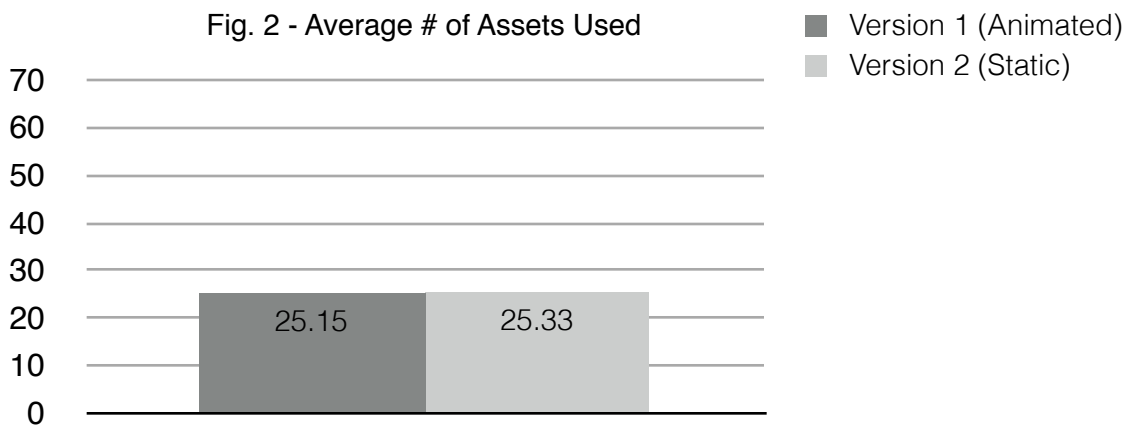
The participants taking each version of the course responded similarly to each other on questions designed to test engagement (Fig. 1).



Version 2 showed somewhat of an advantage on the first survey item; Version 1 showed a negligible advantage on the second and third items. (On the Likert scale, a value of 5 was assigned to Strongly Agree responses and a value of 1 was assigned to Strongly Disagree Responses.)

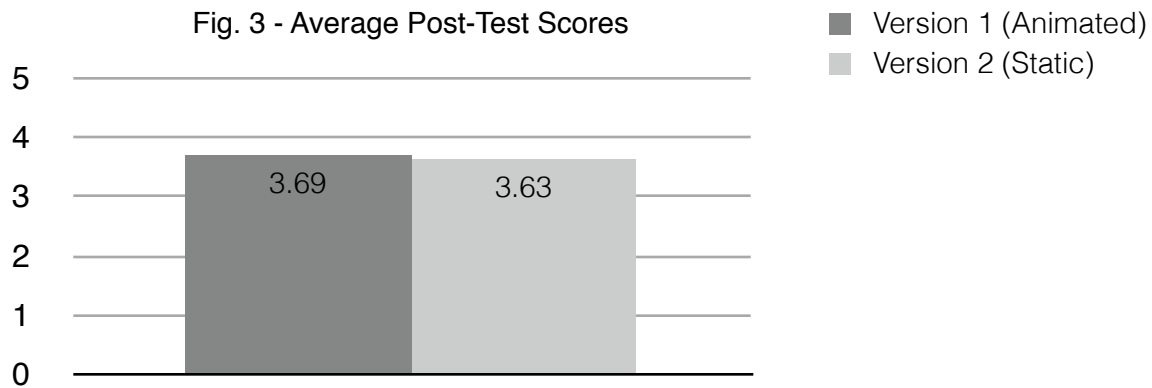
Engagement Measures - Assets Used

The participants also explored their respective versions of the course similarly. Of the 33 course pages and 40 assets that could each be accessed at least once each, participants accessed an average of 25.15 pages and assets combined in Version 1 and 25.33 pages and assets combined in Version 2 (Fig. 2). Again, the difference is negligible.




Recall Measures - Post-Test Score

Participants in both versions scored similarly on the post-test. Those who took Version 1 had an average score of 3.69 correct answers out of 5 and participants who took Version 2 had an average of 3.63 correct answers (Fig. 3).




Factoring In Demographics

Perhaps the most interesting demographic measured in this study was employment status, because it may speak to the motivation of the learners coming into the course (again, the study participants were volunteers and the course content covered tips for finding and getting a job). For the purposes of analysis, the study participants who specified that they were current looking for work were grouped into the Looking category and participants who specified that they were not looking for work were grouped into the Not Looking category.



On the survey, participants in the Looking and Not Looking categories reported similar levels of engagement regardless of which version of the course they took, with Version 2 (static) showing an advantage in the Looking category and Version 1 (animated) showing a less significant advantage in the Not Looking category (Fig. 4). It was somewhat surprising that individuals who reported themselves looking for work showed similar levels of engagement as those who were employed, as it was expected that the former individuals would have a higher level of motivation going into the course.


Study participants did show larger differences in the number of assets used when taking the course when broken down by employment status (Fig. 5). In the Looking category, participants taking Version 2 of the course used significantly more assets than those taking Version 1. The trend reversed itself in the Not Looking category and also showed less of a difference.



However, the expected overall difference in engagement between the Looking category and the Not Looking category did appear when examining number of assets used, possibly suggesting that the number of assets used is a more telling measure of engagement than the survey. Version 1 participants in the Looking category accessed an average of over 7 more assets than those in the Not Looking category, and Version 2 participants in the Looking category accessed over 14 more resources than their counterparts in the Not Looking

category. This could show greater prior motivation on the part of job-seekers, or it could show less familiarity with the subject matter, or both. In any case, the participants in the Looking category who took Version 2 of the course showed a much more significant increase in engagement (as measured by assets used), than did their Version 1 counterparts.

The post-test scores broken down by employment category show the most consistency in any of the data reported: Version 1 of the course showed a slight advantage in both the Looking and Not Looking categories (Fig. 6).



Interestingly, participants in the Not Looking category showed a somewhat higher level of recall than participants in the Looking category, who one might expect to have more motivation to learn the material. However, it is possible that this reflects a higher level of prior knowledge on the part of Not Looking participants.

Subjective Results

Some study participants sent written feedback on the course, though little of it was about the learning agent. Two participants reported liking the animated agent very much. Others reported being unimpressed with either the animated or static agent.

One of the responses on this topic was received via Skype from a 19-year-old research participant followed up a couple of days after she completed the course. After reporting a few of her impressions, she was asked whether she had seen the animated or static version of the learning agent. She was unable to recall which version she had received. (From the LCMS records, she had received the animated version.) It is worth noting that this participant belongs to the age demographic that instructional designers are often told respond well to — or even require — higher-level multimedia in learning experiences.

Conclusions

While I cannot say that a study of this scope definitively proves the effectiveness (or lack thereof) of one approach over another, it seems clear that the use of the animated agent did not contribute significantly to either engagement or recall over the static agent in this study. There are many variations on the use of a learning agent that may yet yield different results; for example, they can be highly character-like, with strong personalities and distinctive clothing (such as uniforms), or they may be used to introduce and conclude content presentations rather than providing instruction throughout. Either of these approaches can be taken with either an animated or static agent; further research is needed to determine whether animation provides an advantage in these cases.

As a basis for making elearning design and development decisions, the individual preferences of instructional designers and stakeholders will often be the final determining factor in which design approaches and development methods to use for any particular course. However, these study results do not seem to support use of an animated agent when the extra time, software, and technical skill required to develop an animated agent are taken into

account. It must also be considered that for some delivery environments, the use of video required by an animated agent, rather than images with audio, may be prohibitive.

Future Extensions of this Research

This study was the result of many hours of research, study design, course design, and course development. I was not able, in the end, to attract a large and diverse enough participant pool to provide specific results to instructional designers working with any audience demographics (or, conversely, establish that demographics had no correlation with the effectiveness of these methods). And while I find this study's results informative, what struck me even more throughout the process is how many variations of this research could be done with most of the same materials, further exploring this topic and others. I firmly believe that the elearning industry could benefit from more research into the effectiveness of our methods and design/development approaches, but the barriers to performing that research are very high because of the time and effort involved in creating elearning courses. (Indeed, some of the researchers already working in this area have taken considerable time and effort to recreate previous courses so that existing research could be extended.) To help lower those barriers, I am offering the following to other researchers under a creative commons license:

- Both versions of the published *Using Social Media for Your Job Search* course
- Native files used to develop the course
- Access to the LCMS implementation (in an observer role, with personally identifying participant data removed)
- Data extracted from the LCMS in the original study (with personally identifying participant data removed)

There are numerous possibilities for altering these materials to create new research studies. Just a few of the possibilities include:

- Creating a third version of the existing course with the static agent and on-screen text instead of narration

- Creating additional versions of the existing course with text-to-speech rather than human narration
- Editing the gestures on the animated learning agent to test the effect of contrary non-verbal communication
- Performing the same research with larger groups or groups with different demographics
- Altering the measures of recall and engagement, or adding delayed assessments to test longer-term recall or behavior change.
- Taking the existing data and further analyzing it to draw additional conclusions

The materials will be hosted at elearningresearch.org. My hope is that others who “remix” and use these materials and data will post their results, and possibly even contribute their own sets of materials to contribute to more and better research-based elearning design.

Bibliography

A. Kamil Mahmood, and Elaine Ferneley, "Embodied Agents in E-Learning Environments: An Exploratory Case Study," *Journal of Interactive Learning Research* 17.2 (2006), *Questia*, Web, 21 Mar. 2010.

A. Silvers, personal communication, April 2011.

Atkinson, Robert K. "Optimizing Learning From Examples Using Animated Pedagogical Agents." *Journal of Educational Psychology* 94.2 (2002): 416. *Academic Search Premier*. EBSCO. Web. 26 Feb. 2010.

Baylor, Amy L. "Running Head: Effects of Image and Animation on Agent Persona." *Department of Educational Psychology and Learning Systems*, Florida State University.

Baylor, Amy L. "Simulating Instructional Roles through Pedagogical Agents." *International Journal of Artificial Intelligence in Education* 15.2 (2005): 95-115. *ACM Digital Library*. Web. 5 Feb. 2011.

Bozarth, Jane. [Social Media for Trainers: Techniques for Enhancing and Extending Learning](#). Pfeiffer. 7 September, 2010.

Clarebout, Geraldine, Jan Elen, W. Lewis Johnson, and Erin Shaw, "Animated Pedagogical Agents: An Opportunity to Be Grasped?," *Journal of Educational Multimedia and Hypermedia* 11.3 (2002), *Questia*, Web, 19 Mar. 2010.

Clark, Ruth Colvin & Chopeta Lyons. "How to Use Visuals to Support Psychological Learning Processes." [Graphics for Learning](#). Pfeiffer: San Francisco, 2004. 47-220.

Clark, Ruth Colvin & Richard E. Mayer. "Applying the Personalization Principle." [E-Learning and the Science of Instruction: Proven Guidelines for Consumers and Designers of Multimedia Learning, 2nd Edition](#). Ed. Rebecca Taff. Pfeiffer: San Francisco, 2008. 168-173.

Craig, Scotty D., Barry Gholson, and David M. Driscoll. "Animated Pedagogical Agents in Multimedia Educational Environments: Effects of Agent Properties, Picture Features, and Redundancy." *Journal of Educational Psychology*. 94.2 (2002): 428-434. Web. 10 Feb. 2011.

D. Elkins, personal communication, February 28, 2010.

Gulz, Agneta. "Benefits of Virtual Characters in Computer Based Learning Environments: Claims and Evidence." *International Journal of Artificial Intelligence in Education* (2004): 416. *ACM Digital Library*. Web. 12 Feb. 2011.

Kathryn Hershey Dirkin, Punya Mishra, and Ellen Altermatt, "All or Nothing: Levels of Sociability of a Pedagogical Software Agent and Its Impact on Student Perceptions and Learning," *Journal of Educational Multimedia and Hypermedia* 14.2 (2005), *Questia*, Web, 19 Mar. 2010.

Moreno, Roxana, et al. "The Case for Social Agency in Computer-Based Teaching: Do Students Learn More Deeply When They Interact With Animated Pedagogical Agents?." *Cognition & Instruction* 19.2 (2001): 177-213. *Academic Search Premier*. EBSCO. Web. 28 Feb. 2010.

M. Lentz, personal communication, April 2011.

R. Smail, personal communication, April 2011.

Reeves, B. (2000). *The Benefits of Interactive Online Characters*. *Center for the Study of Language and Information*, Stanford University.

Reuttimann, Laurie. *Getting the Job You Want: Accelerate Your Job Search with Sales and Marketing Strategies* [video file]. Retrieved from <http://vimeo.com/14370070>.

Reuttimann, Laurie. *Social Media for Job Seekers: A Career Coaches Panel* [video file]. Retrieved from <http://vimeo.com/14684376>.

Vita

Judy Unrein

(503) 512-0140

judyunrein@gmail.com

Education

MEd, Instructional Design, University of Massachusetts Boston. May 2011 (anticipated)

MBA, Baker University. December 2007.

BA, English from University of Texas at Austin. Liberal Arts Honors Program. Recipient – Alcoa Scholarship, Brown Scholarship, Study Abroad Scholarship. December 2000.

Publications - Contributions and Chapters

The Gamification of Learning by Karl Kapp (expected release early 2012)

McGraw-Hill's Catholic High School Entrance Exams, 2ed.

McGraw-Hill's Nursing School Entrance Exams

McGraw-Hill's GMAT

The Princeton Review's Cracking the PSAT

The Princeton Review's WordSmart for the GRE

Publications - Articles

Thinking Outside the Toolbox. Learning Solutions Magazine, 4 April 2011.

Beyond the Hype: Understanding HTML5 and its Potential for e-Learning and mLearning.

Learning Solutions Magazine, 31 May 2010.

Questions to Ask When Choosing an Online Graduate Program. eLearn Magazine, 9 July 2009.

Presentations

HTML5 and Elearning (April 2011). Presentation given via webinar by invitation of the DC-ASTD Technology Special Interest Group.

HTML5: Are We There Yet? (March 2011). Presentation given at Learning Solutions Conference, sponsored by the e-Learning Guild.

How Development Tools Shape Our Design (March 2011). Presentation given at Learning Solutions Conference, sponsored by the e-Learning Guild.

HTML5: Are We There Yet? (February 2011). Presentation given at TechKnowledge 2011, Sponsored by the American Society for Training and Development.

Advanced Camtasia (February 2011). Presentation given at TechKnowledge 2011, Sponsored by the American Society for Training and Development.

HTML5 and Its Potential for E-Learning (December 2010). Presentation given via webinar by invitation of EINet Australia.

Learning Agents (October 2009). Presentation given via webinar by invitation of the DC-ASTD Technology Special Interest Group.

Honors and Awards

Blue Cross and Blue Shield of Kansas City Showcase of Quality

2010 – Runner-up, President’s Award

2010 – Runner-up, Blue Culture Award

2009 – Grand Champion, Strategic Planning and Business Results Award, Small Category

2008 – Grand Champion, Innovation Award, Small Category

2008 – Runner-up, Strategic Planning and Business Results Award, Small Category

2008 KC-ASTD Outstanding Achievement in Workplace Learning & Performance through
Technology and Innovation

Associations

Planning Committee, ASTD TechKnowledge 2012

E-Learning Guild member since 2010

ASTD National member since 2007

ASTD-Cascadia chapter member since 2011

ASTD-Kansas City chapter member 2008-2010

Professional Experience

Learning Designer

October 2010–Present

Artisan E-Learning

Instructional Designer

May 2007–October 2010

BlueCross BlueShield of Kansas City

Kansas City, MO

Instructor/Subject Matter Expert/Author

August 1997–December 2008 (Contract/Part-Time)

Princeton Review, McGraw-Hill Professional

Remote/Various Locations

Product Development Project Manager

December 2004–April 2007

Fred Pryor Seminars/CareerTrack

Overland Park, KS

Production Manager

March 2003–September 2004

Cutter Communications, Inc.

Bend, OR

Software Desktop Applications Trainer

December 2000–November 2002

ExecuTrain, Inc.

Austin, TX