# **Jill Watson:**

# **A Virtual Teaching Assistant for Online Education**

Ashok K. Goel and Lalith Polepeddi

# School of Interactive Computing, Georgia Institute of Technology Corresponding author: goel@cc.gatech.edu

In Learning Engineering for Online Education: Theoretical Contexts and Design-Based Examples. Chris Dede, John Richards, & Bror Saxberg (editors), Chapter 7, pages 120-143, New York: Routledge, 2018. Jill Watson: A Virtual Teaching Assistant for Online Education

Ashok K. Goel and Lalith Polepeddi

#### **Motivations: Learning Assistance in Online Education**

Massively Open Online Courses (MOOCs) are rapidly proliferating. According to Class Central (Shah, 2016), in 2016 more than 58,000,000 students across the world registered for more than 6,800 MOOCs offered by more than 700 institutions. Further, these numbers continue to grow rapidly. Today, MOOCs cover almost all disciplines and education levels, and their students cut across most demographics groups such as gender, age, class, race, religion, nationality, and so forth.

However, the effectiveness of learning in many MOOCs is questionable as the student retention ratio typically is less than 50% and often less than 10% (Yaun & Powell, 2013). Although there are several reasons for the low student retention, a primary reason is the lack of interactivity in MOOCs (Daniel, 2012). Thus, one of the principle recommendations for improving the effectiveness of learning in MOOCs, and thereby also improving student retention, is to enhance the interaction between the teacher and the students (Hollands & Tirthali, 2014).

As an example, consider Georgia Tech's recently launched online section of CS 1301: Introduction to Computing<sup>1</sup> based on the Python programming language. This online section is in addition to traditional, residential sections of the Introduction to Computing class. The online class itself has two sections. In Spring 2017, the accredited section was available only to 45 selected Georgia Tech students who had access to three teaching assistants (TAs) in addition to

<sup>&</sup>lt;sup>1</sup> For a full description of the class, see their website at http://www.cc.gatech.edu/ academics/degree-programs/bachelors/online-cs1301

course materials provided by the instructor. The three TAs provided several kinds of support to the online students, such as answering questions, tutoring on the course materials, evaluating student progress, and so forth. The open and non-credited section of the online Introduction to Computing class—the MOOC—currently has more than 40,000 registered students. The students in the MOOC have access to all the same course materials as the students in the other online section. However, the 40,000 MOOC students do not have access to any TA (or the instructor, except indirectly through the standard course materials). Given that computer programming is a technical skill that many students find difficult to master on their own, it is unclear what percentage of students in the MOOC section will successfully complete the course. It seems safe to say the percentage of students who successfully complete the MOOC section without any teaching assistance will be significantly lower than the students in the online section with teaching assistants.

Of course, most humans are capable of learning some knowledge and some skills by themselves. However, reliable estimates of autodidacts with the capacity to learn advanced knowledge and complex skills are not readily available. For the purposes of the present discussion, let us posit that a vast majority of learners can benefit from learning assistance: perhaps more than 90% of the 58 million students taking a MOOC worldwide may need or want some learning assistance, and perhaps as many as 99% may significantly benefit from learning assistance. If we assume just one teaching assistant (TA) for 50 students for a typical MOOC, then we need at least 1 million TAs for supporting the 58 million students registered for a MOOC! It is highly doubtful that anyone can organize or afford such a large army of human TAs. The Georgia Tech CS 1301 MOOC itself would need about 800 TAs to support the 40,000 students, more than the number of TAs in all other Georgia Tech classes in computing combined.

This raises a profound problem: how can we provide meaningful learning assistance to the tens of millions of learners taking MOOCs?

In response to this question, MOOC teachers, researchers, and service providers are engineering online learning by building on several technologies for learning assistance, such as E-Learning (e.g., Clark & Mayer, 2003), interactive videos (e.g., Kay, 2012; Koumi, 2006), intelligent books (e.g., Chaudhri et al., 2013), intelligent tutoring systems (e.g., Azevedo & Aleven, 2013; Polson & Richardson, 2013; VanLehn, 2011), peer-to-peer review (e.g., Faltchikov & Goldfinch, 2000; Kulkarni, Berstein, & Klemmer, 2015), and autograding. Of course, many of these technologies were developed prior to the start of the modern MOOC movement with Stanford University's MOOC on artificial intelligence in 2011 (Leckart, 2012; Raith, 2011). Nevertheless, MOOCs are extensively developing and deploying these technologies to assist online education.

Another strategy for engineering online learning is to design and develop virtual teaching assistants to augment and amplify interaction with human teachers. These virtual teaching assistants may help with many of the tasks human teaching assistants do, for example, cognitive tutoring, question answering, question asking, autograding, formative assessment, and metacognitive tutoring.

In this chapter, we describe a virtual teaching assistant called Jill Watson for the Georgia Tech OMSCS 7637 class on knowledge-based artificial intelligence (KBAI). Jill Watson (JW) has been operating on the online discussion forums of different offerings of the KBAI class since Spring 2016. At the time of writing this chapter in June 2017, some 750 students and some 25 (human) TAs had interacted with different versions of JW. In the Spring 2017 offering of the KBAI class, JW autonomously responded to student introductions, posted weekly

announcements, and answered routine, frequently asked questions. Thus, JW is a partially automated, partially interactive technology for providing online assistance for learning at scale. In this discussion of JW, we describe the motivation, background, and evolution of the virtualquestion-answering teaching assistant, focusing on what JW does rather than how she does it.

### **Background: An Online Course on Artificial Intelligence**

In January 2014, Georgia Tech launched its online Masters of Science in Computer Science program (OMSCS). OMSCS is a fully accredited, highly selective Georgia Tech graduate degree offered to select students from across the world. The online courses are developed by Georgia Tech faculty in cooperation with Udacity staff, offered through the Udacity platform,<sup>2</sup> and supported by a grant from AT&T. The goal of the OMSCS program is to offer the same courses and programs online that are offered through the on-campus Masters program while maintaining equivalent depth and rigor (Joyner, Goel, & Isbell, 2016). In Spring 2017, the OMSCS program enrolled an order of magnitude more students (approximately 4,500) than the equivalent residential program (approximately 350) that cost far less (approximately \$7,000) than the residential program (approximately \$30,000) (Carey, 2016; Goodman, Melkers, & Pallais, 2016). By now a few hundred students have successfully completed the OMSCS program, and the diploma awarded to them does not mention the word *online*.

As part of the OMSCS program, in 2014, we developed a new online course called CS7637: Knowledge-Based Artificial Intelligence: Cognitive Systems<sup>3</sup> (KBAI). The first author

<sup>&</sup>lt;sup>2</sup> For the selection of classes, see https://www.udacity.com/courses/georgia-tech-mastersin-cs

<sup>&</sup>lt;sup>3</sup> For the course description, see https://www.omscs.gatech.edu/cs-7637-knowledgebased-artificial-intelligence-cognitive-systems

of this article (Goel) had been teaching an earlier face-to-face KBAI course on the Georgia Tech campus for more than a decade. While the online KBAI course builds on the contents of the earlier on-campus KBAI course, we rethought the course for the new medium and developed many of the course materials from scratch (Goel & Joyner, 2016). The second author (Polepeddi) took the online KBAI course in Summer 2015 and was a TA for the course in Spring 2016.

The online, semester-long KBAI course consists of 26 video lessons developed from scratch that help teach the course material (Ou, Goel, Joyner, & Haynes, 2016), a digital forum (Piazza) where students ask questions and participate in discussions as illustrated in figure 7.1, a learning management system through which students submit assignments and receive grades (Sakai), a proprietary peer-feedback tool developed at Georgia Tech where students read and submit feedback on each other's assignments, and a proprietary autograder tool developed by Udacity that helps grade the source code of programming projects. The course is administered by the instructor (typically Goel), who is assisted by a small team of TAs. The TAs typically answer questions and facilitate discussions on the digital forum, and they grade assignments, projects, and examinations.

#### [Insert FIGURE 7.1 HERE]

Since Fall 2014, we have offered the OMSCS KBAI course each fall, spring, and summer term. Enrollment in the class has ranged from about 200 to 400 students each term, so that at the time of writing, about 2,000 online students have enrolled in the course. For the most part, student surveys of the online KBAI course have been very positive (Goel & Joyner, 2016; Ou et al., 2016). In addition, in the fall terms of 2014, 2015, and 2016, we have offered the same KBAI course to residential students at both graduate and undergraduate levels. The performance of the online students on the same set of assessments using blind grading has been comparable to that

of the residential students (Goel & Joyner, 2016, 2017). The retention ratio in the online section has been 75–80%, only slightly lower than the 80–85% in the residential sections.

The OMSCS KBAI course has provided us with a research laboratory for conducting experiments in pedagogy for online education. For example, we have experimented with programming projects based on real artificial intelligence research to promote authentic scientific practices (Goel, Kunda, Joyner, & Vattam, 2013) as well as the use of peers as reviewers and TAs as meta-reviewers (Joyner et al., 2016). We also developed and deployed about a hundred *nanotutors* for teaching domain concepts and methods (Goel & Joyner, 2017). A nanotutor is a small, focused AI agent that models students' reasoning on a particular problem engaging a domain concept or method to be learned. Given a student's answer to the problem, a nanotutor first classifies the answer as correct or incorrect and then provides an explanation on why.

# A Challenge in Scaling Online Education: Responding to Student Questions

Teaching the OMSCS KBAI class in the Fall 2014 and Spring 2015 terms revealed a new challenge for the teaching staff: the discussion forum for the online class was very active and thus took a large amount of staff time to monitor and respond. Table 7.1 provides the data from the discussion forums for the online and residential sections from Fall 2016. As table 7.1 indicates, the discussion forum for the online section had more than 12,000 contributions compared to less than 2,000 for the residential class. One obvious reason for this six-fold increase is that the online class had three times as many students as the residential class. Another, perhaps less obvious reason is that the discussion forum acts as the virtual classroom for the online class (Joyner, Goel, & Isbell, 2016). It is on the discussion forum that the online students ask questions, get and give answers, discuss the course materials, learn from one another, and construct new knowledge.

ριαzza cs- 7637-001	42 🗸	Q & A Resources Statistics Manage Class All Companies 🤁 🚔 💻 👹 I	Lalith Polepeddi 🛛 🗘	
logistics 1 other 2 discussion 2 announcer	ments	assignment1 🚺 assignment2 project1 2 project2 🚺 project3 🚺 midtermexam feedback 🚺 finalexam office-hours		
<ul> <li>Unread Updated Unresolved Following</li> </ul>	Q-	Question History:		
New Post Q' Search or add a post		? question *	185 views	
+ PINNED	A			
■ Instr Start of Week 3 Announceme 1/ Welcome to the start of week 3! Here's what's going on this week. Lessons. The lessons for this week are Means-	25/16	Length of first assignment Okay so when first looking at this assignment I said to myself how on earth am I going to come anywhere close to 1000 words. Now that I am deep into it I am going WAY over 1000 words. I feel that by completion I might be close to 2000 words. Is this going to be an issue and yes I was expecting to have the opposite problem. Thanks, assignment		
Instr Peer Feedback : Complete re Hello everyone, You all should have signed up for peer feedback by now. If any of you still have issues then please pos	22/16			
Instr Announcing Office Hours 1/ We've looked at the survey results from the start-of-course survey, and we're ready to schedule office hours. Th	21/16			
Instr [Discussion] What is intellige 1/ This week's topic What is intelligence? Context Defining intelligence can be helpful to identify the inaredients • An instructor thinks this is a good note	20/16	edit · good question 2		
Instr. Project 1 Questions Thread     Copied/pasted from the other similar topic: if     you have small, simple questions on Project 1,     please post them in this     1 Unresolved Followup	/18/16	the instructors' answer, where instructors collectively construct a single answer     Click to start off the wilki answer		
Instr Assignment 1 Questions Thr 1/ If you have small, simple questions on Assignment 1, please post them in this thread. The purpose of this thread is to r	18/16	followup discussions for lingering questions and comments           Start a new followup discussion		
Instr Academic dishonesty, plagia	1/14/16 5	Compose a new followup discussion		
Dear KBAI Friends, Today we need to discuss another important topic: academic dishonesty.		Average Response Time: Special Mentions: Onli	ne Now   This Week:	
I have been teaching this <ul> <li>An instructor thinks this is a good note</li> </ul>		1.8 hr Anonymous answered Assignment 1 didn't submit in 3 min. 2 hours ago	39 330	
Instr Class Projects	/13/16	Copyright © 2013 Piazza Technologies, Inc. All Rights Reserved. Privacy Policy Copyright Policy Terms of Use Blog Report Bug!		

**Figure 1.** While the video lessons in the OMSCS KBAI course are like a textbook, the class forum is like a virtual classroom where students ask questions, discuss ideas, and give feedback. Here, a student asks a question about whether there is a word limit on an assignment.

While the abundant participation on the discussion forum of the online class likely is an indication of student motivation, engagement, and learning (and thus is very welcome), the higher levels of participation create a challenge for the teaching staff in providing timely, individualized, and high quality feedback. On one hand, the quality and timeliness of TAs' responses to students' questions and discussions are an important element of providing learning assistance and thus play a part in the success of student learning and performance. On the other hand, given the high rate of student participation on the discussion forum, the TAs may not have time to respond to each message with a high quality answer in a timely manner.

# A Potential Answer: Virtual Teaching Assistants

In reading through the students' questions on the online discussion forums of the OMSCS KBAI class in Fall 2014 and Spring 2015, we recognized (as many teachers have done in past),

that students often ask the same questions from one term to another and sometimes even from one week to another within a term. For example, questions about length and formatting of the assignments, allowed software libraries for the class projects, and class policies on sharing and collaborating have been asked in different ways every semester since January 2014. Perhaps more important than that, from the online discussion forums of the Fall 2014 and Spring 2015 OMSCS KBAI classes, we had access to a dataset of questions students had generated and the answers TAs had given.

Thus, in Summer 2015, we wondered if we could construct a virtual teaching assistant that could use the available dataset to automatically answer routine, frequently asked questions on the online discussion forum. We posited that if we could create a virtual TA that could answer even a small subset of students' questions, then it would free the human TAs to give more timely, more individualized, and higher quality feedback to other questions. Also, the human TAs may have more time to engage in deeper discussions with the students.

Our thinking about the virtual teaching assistant was also inspired by IBM's Watson system (Ferruci, 2012; Ferruci et al., 2010). Independent of the OMSCS KBAI class, in Fall 2014, IBM had given us access to its Watson Engagement Manager<sup>4</sup> for potential use in support of teaching and learning. We successfully used the Watson Engagement Manager for teaching and learning about computational creativity in a residential class in Spring 2015 (Goel et al., 2016). Building on this educational experience with the Watson Engagement Manager, in Fall 2015, IBM gave us access to its newer Bluemix<sup>5</sup> toolkit in the cloud. Thus, we were familiar

<sup>&</sup>lt;sup>4</sup> For a detailed description of the IBM Watson Engagement Manager, see http://m.ibm.com/http/www-03.ibm.com/innovation/us/watson/watson for engagement.shtml

<sup>&</sup>lt;sup>5</sup> For a detailed description of the Bluemix toolkit, visit https://www.ibm.com/cloud-

with both the paradigm of question answering and some of the Watson tools.

## **Jill Watson and Family**

At the time of writing (June 2017), we have developed three generations of virtual teaching assistants. We have deployed these virtual teaching assistants in the discussion forums of the online KBAI classes in Spring 2016, Fall 2016, and Spring 2017, as well as in the residential class in Fall 2016. All actual experiments with the virtual teaching assistants have been in compliance with an institutional review board (IRB) protocol to safeguard students' rights and to follow professional and ethical norms and standards.

We call our family of virtual teaching assistants Jill Watson because we developed the first virtual teaching assistant using IBM's Watson application programming interfaces (APIs). However, the names and tasks of specific virtual teaching assistants have evolved from generation to generation as described below. More important, starting with the second generation, we have used our own proprietary software and open-source libraries available in the public domain instead of IBM's Watson APIs (or any other external tool). We made this shift to cover a larger set of questions as well as a larger set of tasks.

## Jill Watson 1.0

#### Design

In January 2016, we deployed the first version of Jill Watson, Jill Watson 1.0 (JW1), to the Spring 2016 offering of the OMSCS KBAI class. Although we included JW1 in the listing of the teaching staff, initially we did not inform the online students that JW1 was an AI agent. As noted above, we built JW1 using IBM's Watson APIs. JW1 is essentially a memory of questionanswer pairs from previous semesters organized into categories of questions. Given a new

computing/bluemix/

question, JW1 classifies the question into a category, retrieves an associated answer, and returns the answer if the classification confidence value is greater than 97%.

Initially, we deployed JW1 on the discussion forum with a human in the loop; if JW1 was able to answer a newly-asked question, then we would manually check that her answer was correct before letting her post that answer to the class forum in reply to the question. In March 2016, we removed the human in the loop and let JW1 post answers autonomously.

Every 15 minutes between 9 a.m. and 11 p.m., JW1 checked the discussion forum for newly-asked student questions. We chose this time interval to mimic the working hours for most human TAs as well as to monitor to JW1's performance throughout the day. If there was a question that JW1 could answer and that another human TA had not already answered, she would post an answer.

## Performance

Figures 7.2 and 7.3 illustrate some of JW1's interactions with the online students on the discussion forum of the OMSCS KBAI class in Spring 2015. (Note that we have blackened some portions of the exchanges to maintain student confidentiality.)



# [Insert FIGURE 7.2 HERE]

## [Insert FIGURE 7.3 HERE]

We found that while JW1 answered only a small percentage of questions, the answers she gave were almost always correct or almost correct. We wanted to both increase the range of questions covered by JW1 as well as the task she addresses. The latter goal led us to develop the next generation of Jill Watson.

### **Jill Watson 2.0**

## Design

In the first week of the KBAI class, we ask students to introduce themselves on the discussion forum by posting a message with their name, their location, why they are taking KBAI this semester, other OMS classes they have taken, activities outside of school, and one interesting fact about them. Human TAs then reply to each student, welcoming him/her to the class. However, it is time consuming to respond individually to 200–400 students within one week. Thus, we built the second generation of Jill Watson, Jill Watson 2.0 (JW2), to autonomously respond to student introductions.

Unlike JW1 that was built using IBM's Watson APIs, we developed the software for JW2 in our laboratory from scratch, using only open-source, external libraries available in the public domain. Further, unlike JW1 that used only an episodic memory of question-answer pairs from previous semesters, JW2 used semantic processing based on conceptual representations. Given a student's introduction, JW2 first mapped the introduction into relevant concepts and then used the concepts as an index to retrieve an appropriate precompiled response.

In August 2016, we deployed two separate, virtual TAs to the discussion forums of the Fall 2016 offerings of the KBAI class that included both an online section and a residential section. We redeployed JW1 to answer routine, frequently asked questions as a TA named Ian Braun and we deployed JW2 to respond to student introductions as a TA named Stacy Sisko.

Just like Ian Braun, every 15 minutes between 9 a.m. and 11 p.m., Stacy checked for newly posted student introductions. Just as with routine questions, if there was a student introduction that Stacy could reply to and that another TA had not already replied to, she would autonomously post a welcome message.

Once again, while we listed both Ian Braun and Stacy Sisko among the teaching staff, we did not inform the students that they were AI agents. To prevent students from identifying the human TAs among the teaching staff through Internet searches, all human TAs operated on the discussion forum under pseudonyms.

## Performance

Stacy Sisko autonomously replied to more than 40% of student introductions. Figure 7.4 illustrates Stacy's responses to student introductions.

# [Insert FIGURE 7.4 HERE]

## [Insert FIGURE 7.5 HERE]

Figure 7.5 illustrates Ian Braun's interactions with students on the online discussion forum. We found that although Ian Braun was a redeployment of JW1, he performed better in the Fall 2016 KBAI class than JW1 did in the Spring 2016 class both in the coverage of routine, frequently asked questions and in the proportion of correct answers. This improvement likely was because by Fall 2016 we had a larger dataset of question-answer pairs since the class had been offered a few more times by then.

## **Jill Watson 3.0**

## Design

Given the success of Stacy Sisko in using semantic processing to reply to student introductions, we created a third generation of Jill Watson, Jill Watson 3.0 (JW3), that uses semantic processing for answering questions. Unlike JW1, JW3 does not use IBM's Watson APIs. Instead JW3 relies solely on an episodic memory. Given a student's question, JW3 first maps the question into relevant concepts and then uses the concepts as an index to retrieve an associated answer from the episodic memory of questions organized into categories.

In January 2017, we deployed two separate, virtual TAs to the Spring 2017 offering of the OMSCS KBAI class. We redeployed version JW2 (or Stacy Sisko) to respond to student introductions as a new virtual TA named Liz Duncan, and we deployed version JW3 to answer routine questions as a virtual TA named Cassidy Kimball. Once again, while we listed both Liz Duncan and Cassidy Kimball among the teaching staff, we did not inform the students that they were AI agents. To prevent students from identifying the human TAs among the teaching staff through Internet searches, all human TAs operated on the discussion forum under pseudonyms. We also increased the time interval during which Cassidy checked for newly-asked questions to 6 a.m. and 11:59 p.m., based on our observations of the activity on the discussion forum.

## Performance

Liz Duncan replied to 60% of all student introductions, a performance superior to that of Stacy Sisko in the earlier generation. Figures 7.6 and 7.7 illustrate Liz's interactions with the online students.

## [Insert FIGURE 7.6 HERE]

#### [Insert FIGURE 7.7 HERE]

We found that Cassidy Kimball performed much better than JW1 and Ian Braun. For example, of the questions that students asked about KBAI's three class assignments, Cassidy autonomously answered 34%, and of all the answers Cassidy gave, 91% were correct. Figures 7.8 through 7.11 illustrate Cassidy's interactions on the online discussion forum.

[Insert FIGURE 7.8 HERE]

[Insert FIGURE 7.9 HERE]

## [Insert FIGURE 7.10 HERE]

## [Insert FIGURE 7.11 HERE]

## **Student Reaction**

In the KBAI classes in Spring 2016, Fall 2016, and Spring 2017, we shared the true identities of the virtual AI agents towards the end of the term. Student reactions to our use of virtual teaching assistants in online discussion forums have been uniformly and overwhelmingly positive. Figure 7.12 illustrates a small sample of student reactions from the KBAI class in Spring 2016 after the students learned about the true identity of Jill Watson towards the end of April.

## [Insert FIGURE 7.12 HERE]

## Discussion

There are several questions about the virtual teaching assistants that we have not fully answered in this chapter. The first question is how does Jill Watson work? As we briefly indicated above, Jill Watson 1.0 uses an episodic memory of questions and their answers from previous episodes. We developed JW1 using the IBM Bluemix toolsuite. In the second generation of Jill Watson, Ian Braun was a redeployment of JW1 for answering questions. However, Stacy Sisko used semantic-information-processing-technology developed in our laboratory to reply to student introductions. In the third generation of Jill Watson, Cassidy Kimball too uses semantic-information-processing-technology developed in our laboratory for answering questions as does Liz Duncan for replying to student answers.

Second, is the Jill Watson technology transferrable to other classes with different student demographics and using different educational infrastructures? To answer this question, we are presently building a new version of Jill Watson for a new Georgia Tech class, CS 1301 Introduction to Computing MOOC, that at present has 40,000 students but no TA support

whatsoever.

Third, is the Jill Watson technology effective in lowering the demands on the teaching staff? While it is too early to determine the answer to this question for the task of question answering, anecdotally there is some evidence to suggest that Jill Watson did reduce the load on the teaching staff for responding to student introductions and for posting messages to the class.

Fourth, is the Jill Watson technology effective in enhancing student performance and improving student retention? We are presently conducting studies and collecting data to answer this question about student engagement, learning, and performance; it is too early to have insights into the issue of student retention.

Fifth, what ethical issues arise in using Jill Watson as educational technology in an online classroom? As we mentioned above, we obtained IRB approval in advance of the Jill Watson experiments. Nevertheless, these experiments have raised several additional, ethical issues. For example, when is it appropriate to use AI agents without telling human subjects about them? Does the use of a feminine name for an AI agent implicitly promote gender stereotypes? Might the use of AI agents as virtual teaching assistants eventually result in reduced employment opportunities for human teachers? These are serious questions that require investigation.

## Conclusion

We may view the Jill Watson experiments from several perspectives on learning engineering. First, we may view Jill Watson as an educational technology for supporting learning at scale. In fact, this was our primary, initial motivation for developing Jill Watson, and this is also how we motivated the discussion in this chapter. As indicated above, Jill Watson uses AI technology for supporting learning at scale by automatically answering a variety of routine, frequently asked questions and automatically replying to student introductions.

Second, we may view Jill Watson as an experiment in developing AI agents so that for highly-focused technical domains, highly-selected subject demographics, and a highly-targeted context of human-computer interaction, it is difficult for humans to distinguish between the responses of AI and human experts. We found that in order to improve coverage, the design of Jill Watson gradually moved from using an episodic memory of previous question-answer pairs to using semantic processing based on conceptual representations.

Third, we may view Jill Watson as an experiment in human-AI collaboration. The KBAI class has become a microsociety in which humans and AI agents collaborate extensively and intensively, living and working together for long durations of time.

## Acknowledgements

We thank IBM for its generous support of this work both by providing us access to IBM Bluemix and through multiple IBM faculty awards. We are grateful to Parul Awasthy for her contributions to the development of Jill Watson 1.0 in Spring 2016. We thank members of the Design & Intelligence Laboratory for many discussions on the Jill Watson project. We also thank the students and (human) teaching assistants of the KBAI course in Spring 2016, Fall 2016, and Spring 2017. This research has been supported by multiple Georgia Tech seed grants for research and education. We especially thank Georgia Tech's Office of the Provost, College of Computing, C21U, and the OMSCS program.

## References

- Azevedo, R., & Aleven, V. (Eds.). (2013). International Handbook of Metacognition and Learning Technologies. New York, NY: Springer.
- Carey, K. (2016, September 28). An online education breakthrough? A Master's degree for a mere \$7,000. *The New York Times*. Retrieved from http://www.nytimes.com
- Chaudhri, V., Cheng, B., Overholtzer, A., Roschelle, J., Spaulding, A., Clark, P., Greaves, M., . .
  Gunning, D. (2013). Inquire biology: A textbook that answers questions. *AI Magazine, 34*(3).
- Clark, R., & Mayer, R. (2003). *E-learning and the science of instruction*. San Francisco, CA: Pfeiffer.
- Daniel, J. (2012). Making sense of MOOCs: Musings in a maze of myth, paradox and possibility. *Journal of Interactive Media in Education*, 2012(3).
- Falchikov, N., & Goldfinch, J. (2000). Student peer assessment in higher education: A metaanalysis comparing peer and teacher marks. *Review of Educational Research*, 70(3), 287– 322.
- Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A., . . . Welty, C.(2010). Building Watson: An overview of the DeepQA project. *AI Magazine*, *31*: 59–79.
- Ferrucci, D. (2012). Introduction to "This is Watson." *IBM Journal of Research and Development 56*(3): 1–15.
- Goel, A., Anderson, T., Belknap, J., Creeden, B., Hancock, W., Kumble, M., . . . Wiltgen, B.
  (2016). Using Watson for constructing cognitive assistants. *Proceedings of the Fourth Annual Conference on Cognitive Systems*, Chicago, IL.

Goel, A., & Joyner, D. A. (2016). An experiment in teaching cognitive systems online. In

D. Haynes (Ed.), International Journal for Scholarship of Technology-Enhanced Learning 1(1).

Goel, A., & Joyner, D. (2017). Using AI to teach AI. AI Magazine, 38(2).

- Goel, A., Kunda, M., Joyner, D., & Vattam, S. (2013). Learning about representational modality: Design and programming projects for knowledge-based AI. *Proceedings of the Fourth Symposium on Educational Advances in Artificial Intelligence*, Bellevue, WA.
- Goodman, J., Melkers, J., & Pallais, A. (2016, October). Can online delivery increase access to education? Harvard University Kennedy School Faculty Research Working Paper Series. RWP16-035.
- Hollands, F., & Tirthali, D. (2014). MOOCs: Expectations and reality: Full report. Center for Benefit-Cost Studies of Education, Teachers College, Columbia University, NY.
- Joyner, D., Ashby, W., Irish, L., Lam, Y., Langston, J., Lupiani, I., . . . Goel, A. (2016, April). Graders as meta-reviewers: Simultaneously scaling and improving expert evaluation for large online classrooms. *Proceedings of the Third ACM Conference on Learning at Scale*, Edinburg, Scotland, (399–408),
- Joyner, D., Goel, A., & Isbell, C. (2016). The unexpected pedagogical benefits of making higher education accessible. *Proceedings of the Third Annual ACM Conference on Learning at Scale,* Edinburgh, Scotland.
- Kay, R. (2012). Exploring the use of video podcasts in education: A comprehensive review of the literature. *Computers in Human Behavior*, 28(3), 820–831.
- Koumi, J. (2006). Designing video and multimedia for open and flexible learning. Oxford, UK: Routledge Falmer.

Kulkarni, C., Bernstein, M., & Klemmer, S. (2015). PeerStudio: Rapid peer feedback

emphasizes revision and improves performance. *Proceedings from The Second ACM Conference on Learning at Scale*, Vancouver, Canada, 75–84.

- Leckart, S. (2012, March 20) The Stanford education experiment could change higher learning forever. *Wired*.
- Ou, C., Goel, A., Joyner, D., & Haynes, D. (2016). Designing videos with pedagogical strategies: Online students' perceptions of their effectiveness. *Proceeding of the Third* ACM Conference on Learning at Scale, Edinburg, Scotland, 141–144.
- Polson, M., & Richardson, J. (2013). *Foundations of intelligent tutoring systems*. Hillside, NJ: Psychology Press.
- Raith, A. (2011). Stanford for everyone: More than 120,000 enroll in free classes [Website]. *MindShift*. Retrieved from http://ww2.kqed.org/mindshift/2011/08/23/stanford-foreveryone-more-than-120000-enroll-in-free-classes/
- Shah, D. (2016, December 25). By the numbers: MOOCS in 2016. Class Central. Retrieved from https://www.class-central.com/report/mooc-stats-2016/
- Van Lehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197–221.
- Yuan, L., & Powell, S. (2013). MOOCs and open education: Implications for higher education [White Paper]. Center for Educational Technology and Interoperability Standards.
   Retrieved from http://publications.cetis.org.uk/wp-content/uploads/2013/03/MOOCsand-Open-Education.pdf