
An Experiment in Teaching Cognitive Systems Online^{1,2}

Ashok K. Goel

GOEL@CC.GATECH.EDU

David A. Joyner

DAVID.JOYNER@GATECH.EDU

Design & Intelligence Laboratory, School of Interactive Computing, Georgia Institute of Technology,
85 5th Street NW, Atlanta, GA 30338

Abstract

In Fall 2014 we offered an online course CS 7637 Knowledge-Based Artificial Intelligence: Cognitive Systems (KBAI) to about 200 students as part of the Georgia Tech Online MS in CS program. We incorporated lessons from learning science into the design of the project-based online KBAI course. We embedded ~150 microexercises and ~100 AI nanotutors into the online videos. As a quasi-experiment, we ran a typical inperson class with 75 students in parallel, with the same course syllabus, structure, assignments, projects and examinations. Based on the feedback of the students in the online KBAI class, and comparison of their performance with the students in the inperson class, the online course appears to have been a success. In this paper, we describe the design, development and delivery of the online KBAI class. We also discuss the evaluation of the course.

1. Background and Motivations

If we want artificial intelligence (AI) to grow as a field of study, then we need to educate new generations of scientists in the discipline. Further, given the early stage of AI research, we want to train the new students in the different paradigms of AI research. However, while the number and variety of AI courses offered by major research universities appears to have grown over the last generation, the number of courses in the cognitive systems school of AI has significantly reduced, and at present there are few online courses focused on the paradigm. Thus, in Fall 2014, we began offering an online course called CS7637: Knowledge-Based Artificial Intelligence: Cognitive Systems (KBAI for short) to around 200 students as part of the new Georgia Tech Online Masters of Science in Computer Science program (OMSCS for short). Although the online course partially builds on an inperson course on KBAI we have taught at Georgia Tech for several years, it nevertheless raised the question: how does one design an effective, repeatable and scalable online course?

Much to our surprise (as well as a little disappointment), we found little guidance on this key question when we started designing the online KBAI course in December 2013. The Georgia Tech OMSCS program officially started only in January 2014, and, thus, when we started designing our course in December 2013, we had little knowledge of the demographics, backgrounds, and goals of students in the program and little prospect for learning from our

¹ An earlier version of this report was presented to the Panel on “Education in Cognitive Systems” at the Third Annual Conference on Advances in Cognitive Systems, Atlanta, Georgia, USA, May 2015

² (<http://www.lessonsfortheOMSCS/CS7637/program>) course of Cognitive Systems are available from Udacity

² All related lessons for the OMSCS CS 7637 online course of Cognitive Systems are available from Udacity

colleagues. Other than the materials from the extant inperson KBAI class, we were mostly on our own in designing the new online KBAI course.

Right at the start we made several strategic decisions. First, we elected to characterize the cognitive systems school of AI as *human-level, human-centered, and human-like AI*. This characterization builds on Goel & Davies (2011), and is compatible with other characterizations of cognitive systems such as Langley 2011 and Langley 2012. Langley (2012), for example, characterizes the vision of the cognitive systems school of AI as “understanding and reproducing, in computational systems, the full range of intelligent behavior observed in humans.” Second, we decided not to directly transfer the course materials from the legacy inperson KBAI class to the new online KBAI course. Although this would have been the simplest course of action, we decided to view the task of designing the online course as an opportunity to reflect on the learning goals, strategies, outcomes, and assessments of the course. Third, we decided not to follow the most common method for making online courses: replay of videotapes of inperson classes. Although again this would have been the simplest course of action, we thought that this method is both limited by the constraints of the old inperson medium and takes minimal advantage of the affordances of the new online medium. Fourth, we decided to view the design, development and delivery of the online KBAI course as an experiment in design-based research on online learning. Fifth, we decided to incorporate as many lessons from the learning sciences into the design of the online KBAI course as possible within the limits of resources available to us. Thus, we adopted learning strategies such as learning by example, learning by doing, project-based learning, collaborative learning, and more. Sixth, we decided to use as much interactive educational technology in the online course as possible, again within the limits of available resources. Thus, we developed and embedded ~150 interactive “microexercises” and ~100 AI agents acting as “nanotutors” into the interactive exercises of the online course. Seventh, as a quasi-experiment, we ran a typical inperson class with 75 students in parallel, with the same course syllabus, schedule and structure and the same assignments, projects and examinations, as the online course. The parallel inperson class provides a mechanism for summative evaluation of learning in the online class.

Based on the feedback from students in the online class and comparison of their performance with that of the inperson class's students, the OMSCS course appears to have been effective at learning about cognitive systems. Indeed, we offered other sections in Spring 2015 and Summer 2015, and plan to offer additional sections in Fall 2015 and Spring 2016. As we expected, some elements of the online course worked better than others. In this paper, we present the design, development, and delivery of the online KBAI class. We also discuss evaluation of the course based on student feedback and student performance.

2. History of the Class

Goel has been teaching an inperson semester-long course on Knowledge-Based AI at Georgia Tech in the fall of each year for more than a decade. The class typically consists of two sections, graduate and undergraduate, which meet together and follow exactly the same syllabus, readings, projects, assessments, etc. The size of the graduate and undergraduate sections over the last few years has varied from 15 to 25 and 25 to 40, respectively. The readings have come from several textbooks, including Patrick Winston's (1993) *Artificial Intelligence*, Mark Stefik's (1995) *Knowledge Systems*, Nils Nilsson's (1998) *Artificial Intelligence: A New Synthesis*, and Stuart Russell & Peter Norvig's (2009) *Artificial Intelligence: A Modern Approach*. (The website for the

Fall 2013 version of the inperson KBAI course before the development of the online course can be found at http://www.cc.gatech.edu/classes/AY2014/cs7637_fall/).

The teaching in the inperson KBAI course adopts a *design stance* towards learning about AI (Goel 1994). Thus, the class work includes four fairly intensive design and programming projects that build on one another. The projects explore selected KBAI concepts as well as the overall goals and methodologies of the cognitive systems paradigm in considerable detail. The assessments include in-class mid-term and final examinations. While the course does not teach AI programming, it provides access to AI programming resources such as the reimplementation of several classic AI systems in Python (Connelly & Goel 2013) originally described in Norvig's (1992) *Paradigms of AI Programming*. Joyner took the KBAI course in Fall 2010, was a teaching assistant (TA) for the course in 2012, and served as the dedicated course manager for the online course during the more than one year of its design, development, and (initial) delivery.

In recent years, the class projects in the inperson KBAI course have focused on visual analogy problems inspired by the Raven's Progressive Matrices (RPM) test of intelligence (Raven, Raven & Court 1998). Positive performance on RPM is known to have a high correlation with other intelligence tests. Thus, although visual in their nature, the RPM tests measure general human intelligence (Hunt 1974). As a result, they are often used as the psychometric measure of choice in educational and clinical settings. The RPM test has attracted much interest in cognitive systems research (Bringsjord & Schimanski 2003; Carpenter, Just & Shell 1990; Lovett et al. 2009) including in our research laboratory (Kunda, McGreggor & Goel 2013; Kunda et al. 2013; McGreggor, Kunda & Goel 2014; McGreggor & Goel 2014). In the projects in the KBAI class, students design, program, and test AI agents on visual analogy problems inspired by the Raven's test. We found that the class projects stimulated student engagement while providing an authentic opportunity to explore cutting-edge research (Goel, Kunda, Joyner, & Vattam 2013).

3. Georgia Tech Online MS in CS Program

The Georgia Tech Online Masters of Science in Computer Science program (OMSCS for short) is a new program of study launched in January 2014 (www.omscs.gatech.edu/). The online courses are supported and delivered by Udacity (www.udacity.com/georgia-tech), and the program is supported by a grant from AT&T. The goals of the OMSCS program are to offer the same educational programs and courses online that we offer on campus with the same depth, substance and rigor. Admission to the program is very selective and competitive. Once the students in the OMSCS program complete the same course requirements as the on-campus students, they will receive a Masters in Computer Science with no 'online' designation on the degree. However, while the on-campus MS in CS degree can cost several tens of thousands of dollars, the OMSCS program costs only several thousand dollars; thus, the OMSCS is an order of magnitude less costly to the students than the on-campus MS in CS program. The OMSCS program currently has more 3000 students, which is an order of magnitude more than the number of students in the on-campus MS in CS program.

4. Design of an Experiment on Online Learning

We view the online KBAI course as an experiment in design-based research on online learning, not as a final product. In design-based research (Brown 1992; Collins 1992), an educational intervention is introduced, extensive formative assessment is conducted, the intervention is redesigned in an iterative cycle of design, evaluate, redesign. Toward these ends, we conducted extensive formative assessment throughout the course. We obtained IRB approval and collected

student demographic data at the start of the class. We conduct anonymized student surveys at the one-quarter, half, and end points of the semester-long class. Georgia Tech also conducts its own summative course survey at the end.

As we mentioned earlier, in a quasi-experiment we ran a regular inperson class with 75 students in parallel with the online course, with the same course syllabus, structure, assignments, projects, examinations, and graders, but with the typical lectures and discussions and no online videos. In a quasi-experimental design (Shadish, Cook & Campbell 2002), there is a control group (the inperson class in our study) and an experimental group (the online class), but subjects in the two groups are not assigned randomly. Further, there can be confounding variables that are not controlled for in a quasi-experimental study, such as demographic differences. Nevertheless, the parallel inperson class provides a baseline for evaluation of learning in the online class. An interesting byproduct of the this quasi-experimental study was that in designing the online KBAI class we also extensively redesigned the syllabus, schedule, structure, assignments, projects, and examinations of the inperson class. This extensive redesign process was motivated by the incentive to dedicate extra time and energy to the development of materials given their intended reuse semester after semester; this redesign improved the inperson class as well.

5. Learning Goals, Outcomes, Assessments and Strategies

Our design for the KBAI class, both online and inperson, follows a four-tiered learning hierarchy. At the apex of the hierarchy are the *learning goals* of the class: the content we want students to have learned at the end of the course. At the second level are the *learning outcomes*: the demonstrable skills we expect students to possess to show mastery of the learning goals. At the third level are the *learning assessments*: tools we plan to use to evaluate the degree to which students have successfully demonstrated the learning outcomes. At the base of the hierarchy are the *learning strategies*: specific strategies to help students accomplish the learning goals.

5.1 Learning Goals

We have four main learning goals for the class.

G1 - Methods: Students will learn the core methods of KBAI. These methods include schemes for structured knowledge representation such as frames and scripts; methods for memory organization such as discrimination trees; methods for reasoning such as constraint propagation and case-based reasoning; methods for learning such as incremental concept learning and explanation-based learning; cognitive architectures such as production systems; and methods for meta-reasoning like strategy selection.

G2 - Tasks: Students will learn the common tasks addressed by KBAI such as classification, understanding, planning, explanation, diagnosis, and design.

G3 - Systems: Students will learn ways AI agents can use these methods to address these tasks.

G4 - Cognition: Students will learn the relationship between KBAI and cognitive science, using theories of human cognition to inspire design of human-centered, human-level, and/or human-like AI, and using designs of AI agents for insights about cognition.

5.2 Learning Outcomes

Learning outcomes refer to observable products and results of learning. We expect three main learning outcomes based on the above learning goals:

O1 - Build Systems: The primary learning outcome is that students should be able to design, implement, evaluate, and describe KBAI agents. The design and description of KBAI agents relates to the first learning goal G1: in order to design such an agent, knowledge of the core methods of KBAI is necessary. Implementation and evaluation, in turn, address the third learning goal G3, to carry out those designs in an actual agent.

O2 - Address Complex Problems: Students should also be able to use these strategies to address practical problems; this learning outcome addresses the second learning goal G2, where students will be able to articulate the relationship between KBAI agents and big real-world problems.

O3 - Reflect on Cognition: In addition, students should be able to use the design of KBAI agents to reflect on human cognition (and vice versa); this addresses the fourth learning goal G4.

5.3 Learning Assessments

In order to evaluate the learning outcomes (and, transitively, the learning goals), we use five types of assessment.

A1 - Projects: Four design and programming projects. The four projects are related, with each project building on preceding projects. Together these projects will address a big problem (such as taking an intelligence test). These projects will assess students' ability to actually design and implement a KBAI agent using the methods discussed in the class.

A2 - Assignments: Eight short assignments. In these written assignments, students will conceptually describe how a particular method might be used to complete the project. Students can choose the eight course topics to use (from the ~24 methods in all).

A3 - Tests: Two take-home tests, a midterm and a final, which examine students ability to reason through the application of the course's topics to a greater variety of problems than is covered in the project.

A4 - Exercises: A large number (~150) microexercises throughout the course. Although these are not incorporated into their grades, these provide us with a look at how students are interacting with and mastering the class material.

A5 - Interactions: Students' interactions with one another, the TAs and the professor. Participation here is not graded explicitly, but it is set as an explicit expectation at the beginning of the course, and interactions here will be used to evaluate students' progress.

5.4 Learning Strategies

In order to achieve the above learning goals and learning outcomes, we will use ten broad pedagogical motivations: learning by example, learning by doing, authenticity, project-based learning, personalized learning, collaborative learning, peer-to-peer learning, learning by teaching, communities of practice, and learning by reflection.

S1 - Learning by Example: Each of the ~25 lessons begins with an example of a real-world task for which we want to build an AI agent. This example is then used throughout the explanation of the method in that lesson to tie the method back to a particular practical problem. Further, each important concept in each lesson is explained with an example. The students see >150 examples of KBAI methods and concepts in the class.

S2 - Learning by Doing: Each lesson includes several micro-exercises, one for each main concept in the class, for a total of ~150 microexercises over ~25 lessons. As students address each of the exercises, they are given targeted feedback directly to the nature of their answer. To give this feedback, we have constructed nanotutors for most exercises. This feedback provides a

route from the student's current completion of the exercise to the right answer. Further, as part of each lesson, students actively solve a problem of the same sort they previously saw in the main example (S1 - Learning by Example) that builds on exercises in the lesson. Thus, in each lesson, students receive at least two problem-solving examples, one at the start of the lesson that we use for situating the lesson and the other from the problem they solve towards the end of the lesson.

S3 - Authenticity of Learning: Whenever possible, we take examples from the real world. Even when this is not possible, we relate the examples to the real world. Further, whenever possible, we relate each topic to current research.

S4 - Project-Based Learning: During the course, each student completes a semester-long project broken into four phases. The big project addresses a real, big and complex problem: taking an intelligence test (a kind of mini Turing test; A1). In service of this project, students are asked to conceptually relate several class lessons to the project in the form of written design assignments (A2). Thus, each topic covered in the class is tied directly to the overall project for the class. Further, students are able to run their projects against a set of sample problems to discern exactly how well their project currently operates. Students receive feedback in real time on its success and can revise the project accordingly

S5 - Personalized Learning: Personalization is incorporated throughout the course. First, on every exercise, students are given individualized, targeted feedback (S2). Similarly, on the projects, students are able to run their projects and receive feedback in real time on its success and can revise the project accordingly. Third, students are provided with a high-level mental map for the course, allowing them to navigate between topics to pursue their interests rather than always follow our scripted order of the course. Fourth, students are given leeway in choosing which assignments to complete for the course; they must complete 8 short assignments, but they can choose from around 24 different prompts.

S6 - Collaborative Learning: We form small “study groups” of all the students in the course. While the tests, the projects and the assignments in the course require individual work, we encourage the study groups to work together on all aspects of the course (including discussions about the projects and the exercises).

S7 - Peer-to-Peer Learning: After each test, project and assignment, we publicly post the best tests/projects/assignments along with our critiques. Students are requested to read through the outstanding work; they are expected to raise their own work to the same level of excellence.

S8 - Learning by Teaching: We empower the students and provide opportunities to students to act as teachers to one another. We ask the students to provide feedback to other students in their assignments (though we do not count this feedback as part of the grade).

S9 - Learning by Reflection: At the conclusion of each lesson, we ask each student to reflect on what they learned in the class. Each design project requires the writing of a design report that explains and critiques, and reflects on the student’s work on the project.

S10 - Community of Practice: We use an online discussion forum dedicated to the class to help develop a community of practice. We encourage all students to introduce themselves on the forum, and support peer-to-peer information sharing, question answering, as well as discussions and debates. The TAs and the professor not only monitor the forum and publicly answer questions, but they also seed discussions. We also hold regular office hours via Google Hangout.

6. Design of the KBAI Course

Design of the online CS7637 KBAI course began in December 2013. At present the course is comprised of 26 lessons. The lessons vary in length based on the topic (one of the advantages of

Current State

B A C
A on Table
B on Table
C on D
D on Table
 $\Delta = 2$

Available Operators:

Move(Object, Location)
e.g.:
Move(C, Table)
moves C onto the table
Move(C, B)
moves C onto B

Goal State

A B
B on C
C on D
D on Table

Try again!

One of your moves (move #2) didn't work. Block B has block A on top of it, so it can't be moved.

Discussions [See All](#)

No discussions for this part of the lesson yet!

[Start a Discussion](#)
[Report an Issue](#)

[Rewatch Instructions](#) [Submit Answer](#) [Continue to Answer](#)

Figure 1. An example of a microexercise and the output of the nanotutor on the right.

preparing the class in this medium), but average to approximately one hour per lesson when including the time students spend completing the interactive microexercises in each lesson. The videos of all 26 lessons are now available publicly and freely through Udacity at <https://www.udacity.com/course/knowledge-based-ai-cognitive-systems--ud409>. A (free) Udacity account is required to access these materials.

6.1 Course Syllabus

The KBAI course contains lessons on the following topics: (1) Introduction to the course, (2) Introduction to KBAI, (3) Semantic Networks, (4) Generate & Test, (5) Means-Ends Analysis and Problem Reduction, (6) Production Systems, (7) Frames, (8) Learning by Storing Cases, (9) Case-Based Reasoning, (10) Incremental Concept Learning, (11) Classification, (12) Logic, (13) Planning, (14) Understanding, (15) Commonsense Reasoning, (16) Scripts, (17) Explanation-Based Learning, (18) Analogical Reasoning, (19) Generalization and Version Spaces, (20) Constraint Propagation, (21) Configuration, (22) Diagnosis, (23) Learning by Correcting Mistakes, (24) Meta-Reasoning, (25) Advanced Topics, and (26) Course Wrap-Up.

6.2 Exercises and Tutors

Embedded in the 26 video lessons are ~150 interactive microexercises, averaging to approximately six exercises per lesson. This leads to an interactive microexercise approximately every eight minutes in a lesson. The exercises in the KBAI class go beyond what is typical in most MOOCs. First, the exercises are typically far more open-ended and iterative than simple multiple-choice or short fill-in-the-blank questions. Figure 1 illustrates one microexercise. The input for the exercise is comprised of free-response text that is instructed to follow a certain format. Other exercises in the course combine several multiple-choice questions, multiple free-response text boxes, and other more complicated structures. In this way, the activities that

students are offered in the midst of these lessons are far more authentic and creative than answering simple multiple-choice questions.

Secondly, a set of nanotutors augments the microexercises. The majority of the interactive exercises in the course are equipped with a nanotutor, building on our earlier work on intelligent tutoring systems (Joyner & Goel 2015). These nanotutors give students targeted, individualized, just-in-time feedback on students' responses to the current exercise. The tutors operate first by examining whether the input to the problem even makes sense. If not, the nanotutor supplies feedback on the type of input it will understand, guiding students along to the closed input set that it can process. Then, once it understands the input, it examines whether that input is valid; in the exercise above, it would check if all the moves are legal. If the input is valid according to the rules of the exercise, it moves on to checking correctness; in the exercise above, there exists valid feedback that does not answer the actual question of the exercise. Finally, for some microexercises, the nanotutor also checks to see if the answer is the *best* answer. In the exercise above, the nanotutor might comment that while the goal was achieved, it could have been achieved in fewer moves.

6.3 Readings

As in the inperson class in previous years, the recommended readings came from several textbooks, Stefik (1995), Russell & Norvig (2009), and Winston (1993). In particular, the course covered about three fourths of the Winston book (most all, except for the chapters on search, genetic algorithms, neural networks, and vision), about a third of the Stefik book (especially the chapters on classification, configuration and diagnosis), selected chapters from the Russell & Norvig book (such as planning). In addition to the above-recommended readings, we included several optional readings on selected topics in cognitive systems such as Lehman, Laird & Rosenbloom (2006) on the SOAR cognitive architecture.

6.4 Projects

As in the inperson class in previous years, the projects in the KBAI class are built around the Raven's Progressive Matrices (RPM) test of intelligence. Due to copyright and other issues, we are unable to use the actual RPM as part of the class projects, but instead we have developed a set of RPM-inspired problems that leverage the same transformations and reasoning strategies seen on the actual RPM. Figure 2 illustrates a 2x1 problem from our problem set. On the top is a 2x1 matrix with one entry missing. On the bottom are six choices. The task is to write an AI agent that can autonomously select one of the six choices on the bottom for insertion into the missing entry on the top and completion of the pattern in the 2x1 matrix. Similarly, Figure 3 illustrates a 3x3 problem from our problem set. On the left is a 3x3 matrix with one entry missing. On the right are six choices. The task again is to write an AI agent that can autonomously select one of the six choices on the right to insert into the missing entry on the left and thereby complete the pattern in the 3x3 matrix. For the Fall 2014 section of the class, we used 123 of such problems: 27 2x1, 48 2x2, and 48 3x3. Although 2x1 problems are not present on the actual RPM, we use them as a soft introduction to the type of reasoning that is needed on the test. In Fall 2014, students completed four projects. In projects 1, 2, and 3, students addressed the 2x1, 2x2, and 3x3 problems respectively; each project was also run against the problems from the previous project(s). In these three projects, students designed KBAI agents that operated on symbolic, verbal descriptions of the RPM problem designed by Joyner and given as part of the input.

2x1 Basic Problem 16

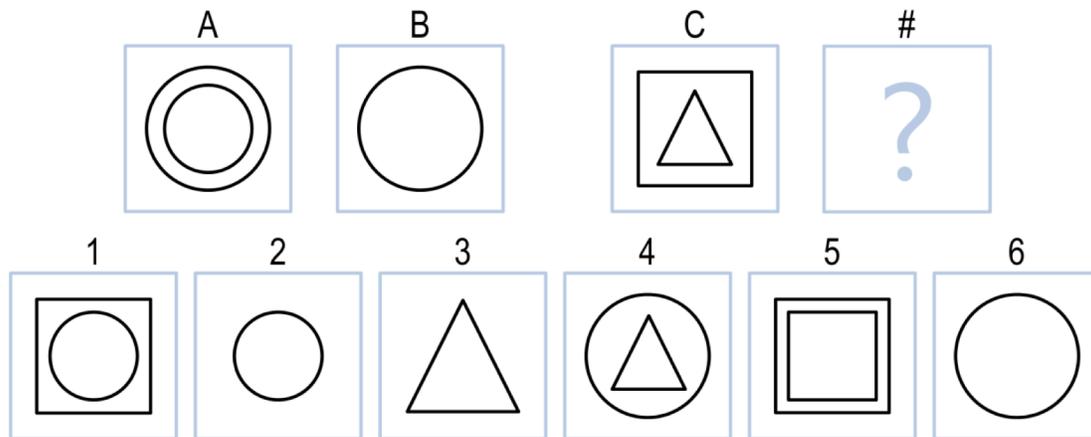


Figure 2. A 2x1 visual analogy problem inspired by the Raven's Progressive Matrices test.
(The correct answer is 3.)

In project 4, the input to the AI agents was the image files representing each frame of a problem. On each of the four projects, students' agents were run against three kinds of problems: Basic, Test, and Extra. Basic problems were provided to students during the design of their agents; part of students' grades on the projects was tied to how many Basic problems their agents solved correctly. Test problems were also used to calculate students' grades, but they were not provided to students in advance; in this way, Test problems helped check agents for overfitting to the provided problem set rather than general problem solving. Extra problems were provided to students during the design of their agents but were not counted for a grade; these were provided simply as a 'challenge' exercise for students who wanted to address some extra, harder problems. Joyner et al. (2015) provide more information on the nature of these projects.

6.5 Assignments

In Fall 2014, students also completed eight written assignments. Each assignment had the same general prompt: choose any of the topics covered in the class and discuss how that topic might be used to address RPM problems. 24 total topics are covered in the class, meaning that each student would choose 8 of the 24 topics to use at some point during the semester. Early in the semester, these assignments served to help students brainstorm and gather feedback on their approaches to designing their agents; later in the semester, these assignments served to help students think about these techniques could be used to address bigger, broader problems than the handful of RPM provided during the projects.

6.6 Examinations

There were 2 examinations: 1 mid-term examination and 1 final examination. Both examinations were of the "take home" type. Thus, the students had access to all kinds of information resources at their disposal. All questions on the mid-term and final examinations were open-ended. For the mid-term examination, the questions were based on a science fiction story. The questions on the final examinations originated from research projects in our laboratory.

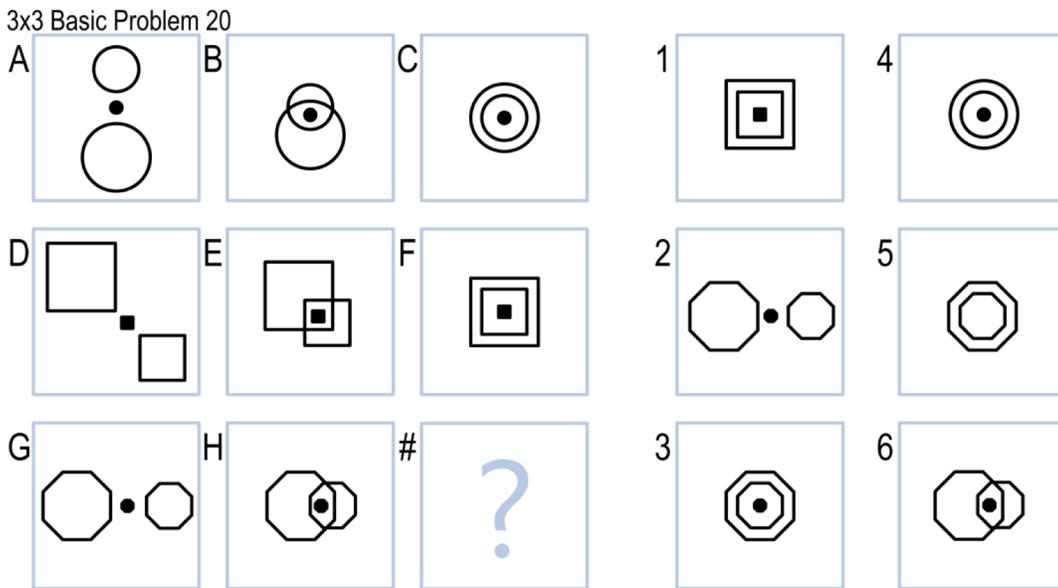


Figure 3. A 3x3 visual analogy problem inspired by the Raven's Progressive Matrices test. (The correct answer is 3.)

7. Development and Delivery of the Online Course

It is important to distinguish two different elements of the KBAI course: development and delivery. ‘Development’ refers to the process by which the course materials were assembled, recorded, edited, and reviewed. The result of the development process is a kind of “video textbook” of sorts, a collection of high-quality video lessons together with the readings, projects, assignments, grading rubrics, and other reusable materials. ‘Delivery’ refers to the act of actually teaching the course in a particular semester, which involves several facets that cannot be reused from semester to semester, such as office hours, virtual interactions on the discussion forums, examinations, and the actual grading.

7.1 Development of the Course

Development of the KBAI course began in February of 2014 with an intense two-day boot camp. During this boot camp, the course developers (and authors of this paper) visited the Udacity headquarters in Mountain View to learn the Udacity course development process and recording procedures. During this boot camp, we developed the learning goals, outcomes, and strategies for the entire course, as well as the entire structure of the course, identifying the 26 distinct lessons for production and recording.

After completing the boot camp, we began a two-month (March to May 2015) process of scripting the 26 lessons. Just as we had done for the course as a whole, for each lesson we articulated a set of learning goals, outcomes, and strategies, as well as a set of assessments and a lesson plan. Each lesson was constructed around a series of interactive microexercises and opened with a video of the professor describing the learning goals and outline for the lesson, then moved

on to a visually rich explanation of the topic. At the conclusion, course developer Joyner briefly recaps what was learned during the lesson and connects the lesson to other parts of the class.

At the conclusion of the scripting process, we spent two months (May to July 2015) recording the lessons and turning the scripts into polished, final videos. After recording all the filmed material for the class, we assembled the descriptions and rubrics for the course's four projects, eight assignments, and two exams (July to August 2015).

7.2 Delivery of the Course

The KBAI course launched in middle of August 2014 and ended in the middle of December 2014 for a total of 16 weeks of learning. We used a number of tools during the course: videos lessons were delivered via Udacity, assignments and announcements were given and received through Georgia Tech's Learning Management System T-Square; office hours were handled through Google Hangouts; and the discussion forum was hosted on A (www.piazza.com). We also an interactive tool developed locally by our colleague Joe Gonzales, a graduate student at Georgia Tech, for managing peer-to-peer feedback on the assignments. In addition, we wrote scripts for running and grading the students' AI agents on the 123 problems in the four projects.

The teaching team for the online KBAI and the parallel inperson class consisted of 9 TAs in addition to the two course developers (Goel and Joyner). Between the two courses, all assessments were identical: students completed the same assignments, projects, and examinations. Grading was performed blind as well; graders were not aware of which papers or projects came from online students and which came from inperson students.

7.2.1 Student Demographics

At the start of the Fall 2014 course, we administered a survey to students to examine the overall demographics of students in the online course in comparison to the inperson. (The results presented here are only from students who consented to participate in the IRB approved research study, though we do not anticipate systematic differences from the class as a whole because of the fairly large sample sizes. The student demographics in the online class were strikingly different compared to the demographics of the inperson class. OMSCS students were on average older, with nearly half (47%) falling between 25 and 34 years old. 82% of inperson students, by contrast, were under the age of 24. A total of 86% of OMSCS students were above age 25. OMSCS students were significantly more educated as well: 87% of OMSCS students had previously completed a Bachelor's degree, 11% already completed a Master's degree, and 2% already had a Doctoral degree. By contrast, 50% of the inperson course was working on their Bachelor's degree, while 43% already had a Bachelor's and 6% already had a Masters. In addition, OMSCS students had on average twice as much programming experience as the inperson students (10 years compared to 5 years). The majority of OMSCS students also reported that they also have full-time jobs and families, while the majority of inperson students are full-time students without families.

Interestingly, while the overwhelming majority of inperson graduate students in the College of Computing at Georgia Tech are international students, the majority of the OMSCS students were American students. 89% of students in the OMSCS KBAI class resided in the United States while 11% resided abroad. Additionally, while 24% of the inperson class was female, only 9% of the online class was female.

7.2.2 *Expectations and Retention*

When asked at the start of the term, how much time they anticipated spending on the class, OMSCS students replied an average of 10.9 hours per week, with a high standard deviation (5.6). For comparison, inperson students expected to spend only 7.9 hours per week with a standard deviation of 4.0. The general Georgia Tech heuristic is that students will spend three total hours for each credit hour taken; CS7637 is a three-credit hour class, and thus the actual time expectation prescribed by this heuristic would be 9 hours of time outside the classroom.

The online class started with 196 students. At the end of the semester, 170 students were enrolled in the class and received a grade for it, a completion rate of about 87%. This is a very high rate compared to other OMSCS classes and much higher than most MOOCs. The rate also mimics the common rate for inperson offerings; in Fall 2014, for example, the inperson offering of CS7637 had a completion rate of 88%.

7.2.3 *Discussions in the Online Classroom*

The Piazza discussion forum acted as the virtual classroom for the online KBAI class. The function that it plays is so important that students deeply appreciated it, but at the same time, it could become overwhelming for some students. The 170 students who concluded the class generated ~6500 contributions during the semester, while the inperson section generated only ~925 contributions from 70 students.

8. Evaluation of the KBAI Course

As we mentioned earlier, we view the online KBAI course as just one iteration in design-based research on online learning in cognitive systems. Thus, evaluation of learning in the KBAI course is an important question for further iterations of the course. In evaluating the online course, two approaches can be taken. First, we may look at student responses to the several surveys offered during the course; although student evaluations are not a complete metric for evaluating course success, they do offer valuable insights into students' perceptions of the course. To complete the picture, however, we may then look at student performance in the course, especially in comparison to the inperson course. The inperson course has been offered for many years and has led to strong learning in the past, and thus, the inperson section provides a useful "control" against which to examine the performance of the online students.

8.1 Student Feedback on the KBAI course

Questions on student perceptions of the course take two forms: open-ended survey questions and responses to Likert-scale questions asking for general thoughts on the quality of the course. Four surveys were administered throughout the semester asking for student feedback on the course: a quarter-course survey after 5 weeks, a mid-course survey after 9 weeks, an end-of-course survey after 15 weeks, and the final institute-sponsored survey after the class ended. Altogether, the student feedback collected for the class totals over a hundred pages, and is far too thorough to report conclusively here. In lieu of the fine-grained analysis of individual parts and tools used in the course, we instead present three high-level views of the course as a whole.

8.1.1 Internal Systematic Evaluations

During the surveys, we asked students several questions to examine the pace, rigor, and quality of the CS7637 course. In all these areas, we found the results we would desire and expect: students rated the course as appropriate, with a slight skew toward students finding the course a little too hard or fast. This matches the reputation of Georgia Tech as a tough university and the desired level of challenge we wanted to create.

With regard to pace, during the mid-course survey, 76% of the class responded that the pace of the course was about right, while 18% rated the course as too fast. With regard to rigor, students similarly rated the course as "about right" overall, with 73% of students rating the rigor as appropriate in the mid-course survey while 21% chose some level of rating the course as too difficult. Students also generally agreed with positive statements regarding the course's quality. In the mid-course survey, students were given four statements and asked to rate their level of agreement or disagreement. 97% of OMSCS students agreed that "The lectures are informative and easy to understand." 90% agreed that "The exercises provided during the lectures keep me engaged." 80% agreed that "The feedback I receive from the exercises enhances my understanding of the material." 99% agree that "Overall, the video lessons are valuable in helping me learn."

Students were also given a standard survey by the university itself to complete at the end of the course. This survey is notable because anonymity is more clearly protected; while we assure students we do not know who submits which survey responses, they are forced to take us at our word; in this Institute survey, however, their responses went directly to the Institute rather than to us. Here, too student ratings of the course were similarly positive. Students reported spending an appropriate amount of time on the course, with 55% reporting spending 9 to 12 hours a week on the course and 39% reporting spending more. Several Likert-scale questions were administered as well, and their interpolated medians on a scale of 1 to 5 were received. With regard to the amount learned, an interpolated median of 4.7 was observed, with 61% of the class selecting 5 and 30% selecting 4. Interpolated means of 4.6, 4.4, and 4.8 were observed respectively on questions regarding whether assignments facilitated learning, measured knowledge, and met course objectives; on each item, over 50% of respondents selected a 5. Most importantly, the course received an interpolated median of 4.8 with regards to its overall effectiveness, with 73% of students selecting 'Strongly Agree' when asked if the course was effective.

8.1.2 Compared to Other Courses

During the quarter-, mid-, and end-of-course surveys, we asked students to compare the KBAI course to three other categories of courses: other OMSCS courses, other online courses outside the OMSCS, and all other courses they have taken. The statistics cited here are from the mid-course survey, although the responses do not differ significantly between the different surveys.

With regard to the videos and lessons themselves (separate from any other element of the class), 69% of students who had taken other OMSCS courses rated the course's lessons and videos as better than other OMSCS courses; 83% of students who had taken other online courses rated KBAI's materials as superior; and, most interestingly, 79% rated them as better than the lectures for other classes they have taken as a whole, online or inperson. This is notable for three reasons. First, the development of the KBAI class received the same allocation of resources as other OMSCS courses, and yet a notable majority of students rated the course materials as superior to other OMSCS courses. This suggests that the focus on learning goals, strategies, assessments, and outcomes within the videos led to a significant difference in students'

perceptions of the materials. Second, this online classroom was rated by a significant majority of students in the class as better than other classes they have taken, including inperson classes; this is contrary to the somewhat common notion that online classes trade the quality of the class for increased accessibility and affordability. Instead, here we see that despite being online, the class was still rated superior. Third and perhaps most remarkably, 79% of students prefer the online KBAI course over their traditional classes but only 69% prefer it over other OMSCS classes; this suggests that the OMSCS program itself is perceived positively relative to traditional course experiences.

The above question focused on a comparison of the course materials, separate from the course as a cohesive whole. We also asked students to rate the course as a whole compared to the same three groups of other courses; in this comparison, students are comparing the assessments, projects, forum interactions, and the rest of the class rather than just the videos. On the mid-course survey, 76% of the students who have taken other OMSCS courses rated KBAI as better; 91% of the students who have taken other online courses rated KBAI as better; and 77% of the students rated KBAI as better than other classes as a whole. Here, we see a few additional interesting impacts. First, the third advantage seen previously disappears a bit; this suggests that the polished and produced videos are regarded as better than traditional lectures, but that the course experience as a whole is rated consistently between traditional classes and OMSCS classes. Second, there exists a jump in perception, suggesting that there are positive impressions of CS7637 separate from the videos themselves. Third, despite the equal resources, CS7637 is again rated significantly higher than other OMSCS courses, suggesting that the emphasis on a close tie between learning goals, outcomes, strategies, and assessments is improving the student experience.

8.1.3 Free-Response Feedback

Finally, we also asked students to provide free-response answers to two questions on the course as a whole: what elements they liked and wished to see used in other courses, and what elements they would suggest changing for future semesters. Students put forth several different elements that they would like to see used in other future classes. With regard to the design of the course videos and lessons, many students commented on the strong top-down structure of the course, wherein each lesson is mapped to the big picture explicitly throughout the course. Relatedly, students also remarked on the design and quality of the lessons themselves. With regard to the design of the course's assessments, students were very positive about the flexibility and project-based nature of the course. Many students expressed positivity that the project left room for them to explore in their own desired directions rather than being forced to meet a firm set of standards and metrics. The most common piece of positive feedback, however, on the level of communication with the course developers and the way in which the teaching team created the kind of course routine that would otherwise arise naturally in an inperson course.

Overall, students were overwhelmingly positive about the course. One student commented, "Please have other OMSCS courses follow the teaching methodology used in this course." Another replied, "Overall, one of the best courses I have ever taken either in person or online." And, perhaps most significantly, a third wrote, "This course impressed on me so much that I have changed my specialization from Software and DB to Interactive Intelligence."

Students also noted areas for potential improvement. Interestingly, the most common critique was that students wanted more depth in the material; many students themselves actually requested more readings, more coding assignments, and other additional required material. Beyond this

feedback, the majority of suggestions were largely hygienic, such as resolving the difficulty with navigating a very busy forum and the fragmented nature of interacting with so many different course tools.

8.2 Student Performance in the KBAI Course

Students in the Fall 2014 offering of the KBAI course completed eight written assignments, four projects, and two exams. All assignments were graded blindly; graders were not aware which students came from the online class and which came from the inperson class, and each grader received assignments to grade from both sections. The table below compares the performance of the inperson students with the performance of the online students. This comparison includes only those students who completed the entire course.

Table 1: Average grades on each assignment between the inperson and OMSCS sections of CS7637 in Fall 2014.

Item	Max	OMSCS (Mean)	OMSCS (SD)	Inperson (Mean)	Inperson (SD)	OMSCS - Inperson
Assignment 1	4	3.90	0.33	3.52	1.10	0.38
Assignment 2	4	3.94	0.26	3.70	0.89	0.24
Assignment 3	4	3.95	0.37	3.52	1.19	0.42
Assignment 4	4	3.92	0.30	3.83	0.82	0.09
Assignment 5	4	3.89	0.50	3.75	0.75	0.14
Assignment 6	4	3.86	0.48	3.62	0.98	0.24
Assignment 7	4	3.91	0.46	3.77	0.73	0.14
Assignment 8	4	3.97	0.16	3.90	0.51	0.08
Project 1	100	94.47	2.49	92.61	11.75	1.86
Project 2	100	92.74	5.08	89.64	16.18	3.10
Project 3	100	93.10	5.23	92.17	12.20	0.92
Project 4	100	92.0	6.20	88.5	16.5	3.53
Midterm	100	70.2	7.20	70.0	5.70	0.20
Final Exam	75	93.76	11.15	93.48	11.92	0.29
Final Grade	100	92.32	5.38	91.31	7.12	1.01

OMSCS students outperformed inperson students on every assessment in the class and in the class as a whole. On seven of the fourteen assessments (highlighted in grey above), this difference was statistically significant. Thus, in terms of duplicating the learning seen in the inperson KBAI class in the past, the OMSCS offering of CS7637 was successful: students in the OMSCS section performed as well as or better than students in the inperson class.

9. Design Guidelines

Our experience with the OMSCS KBAI class has led us devise several guidelines for developing online classes including the following ten:

1. **Establish learning goals, outcomes, strategies, and assessments first:** Actually this guideline has been known for a long time and is equally applicable to online and inperson learning. However, it is surprising (and also a little disturbing) how often educators start teaching without first establishing and articulating specific learning goals, outcomes,

strategies, and assessments. Once the learning goals, outcomes, strategies, and assessments have been articulated, they can be shared, critiqued, revised, and become the basis of later reflection. Section 5 enumerates our learning goals, outcomes, strategies, and assessments for the KBAI class in both the online and inperson versions.

2. **Allocate adequate time for design, development, and delivery:** Design, development and delivery of an online course is labor intensive and time consuming. We estimate that it took >1000 hours of our time (~200-250 hours of Goel's and ~750-800 of Joyner's) to design and develop the OMSCS KBAI course starting from the materials for the inperson class we had taught for several years. We further estimate that it took ~120-130 hours per week on average (~10 hours of Goel's, ~20 hours of Joyner's, and ~90-100 hours of TAs') to deliver the class over the 16 weeks of the Fall 2014 semester. It is important to note that delivery of an online class is as important as its design and development (though often the former receives much less attention than the latter).
3. **Deliberately recreate natural features of the inperson class:** The inperson classes typically have a structure and a rhythm, and it is important to recreate them in online courses. As an example, in an inperson class, the teacher may often share where the class stands in relation to the learning goals, highlight salient points learned recently, and give reminders of upcoming deliverables and deadlines. It is important to not only do the same for an online course, but also establish rhythm. In the OMSCS KBAI class, we found that it helped to post "start of week announcements" that sketched the upcoming lessons, deliverables and deadlines as well as "end of week announcements" that highlighted salient lessons and points.
4. **Leverage the advantages of digital media for online learning:** Digital media have unique affordances for online learning. For example, digital media allow for asynchronous learning in the form of watching the video lessons, completing the microexercises, discussing the class materials, etc. Discussions in online classes may unfold over days or even weeks, and attract scores or even hundreds of messages. As mentioned earlier, we used Piazza as the forum for online discussions, and actively encouraged discussions, which resulted in ~6500 contributions. Some of the conversations were quite deep and continued for weeks.
5. **Design project-based learning carefully:** As we mentioned above, and as we describe in detail in Joyner et al. (2015), the design projects in the OMSCS CS 7637 class were a success in accomplishing the outcomes desired of the class. The "learning by doing" projects motivated and engaged most students in the class. The student performance on the projects was better than we had expected based on our previous experiences with using similar projects in inperson classes (Goel, Kunda, Joyner, & Vattam 2013). On reflection, we think the projects were so successful because they were (a) personal (the students could relate to them in terms of their own cognition), (b) incremental (students could measure the progress they were making in terms of numbers of correctly addressed problems), and (c) challenging (few students could write AI agents that could address all the problems on the projects) but achievable (almost all students could write AI agents that could address at least a few problems).

6. **Understand the audience:** This applies to teaching in general, including both online and inperson classes. It is important to note the student demographics and learning goals of online students can be quite different from that of their inperson counterparts. While most students in the inperson class are in the 20 to 25 years range and have had only limited non-college experience, most students in the OMSCS program are much older, and have had more experience outside college. The OMSCS students typically are also more motivated and engaged. This means that it is not only appropriate but also productive to challenge their abilities.
7. **Break the isolation experienced by many online students:** Students in inperson classes get to meet the instructor and other students in the classroom. Students in inperson classes can also meet the instructor, the teaching assistants and other students outside the classroom. Online students typically do not have the same privileges, and thus some online students can sometimes feel isolated. It is important to help online students overcome this feeling of isolation. We used several strategies for this, including asking all students to post introductory messages to the Piazza forum in the first week of classes, encouraging them to post additional messages during the course and answering the messages promptly and positively, and holding regular and frequent office hours via Google Hangout,
8. **Solicit feedback and be ready to iterate:** Any class, whether online or inperson, is an experiment in teaching and learning. In case of the OMSCS KBAI class, we deliberately viewed the class as an experiment as described earlier. Thus, we sought formative assessment early and often: for example, we conducted (anonymized) quarterly survey of students. Further, we used the student feedback to revise the course both during the semester and between semesters.
9. **Leverage peer feedback and autograding wisely:** Peer feedback involves students giving each other scores and feedback on their assignments, while autograding involves an AI agent giving scores and feedback. Both of these can be powerful methods for scaling a class because they reduce the need to add more course staff as the class grows. However, these tools can also reduce the quality of the class. Peer-written feedback will rarely be as accurate and effective as expert-written feedback, while the process of systematizing projects to the point of being autogradeable may compromise the learning goals. Instead, the results of peer review should *inform* and *complement* the grading process rather than replacing it. Graders can be equipped with the results of peer review, or the peer review can simply provide additional feedback. Similarly, autograding results should be made available to students during the project design so that iterative improvement is possible rather than reserving the benefits for the graders and teaching team.
10. **Use the online class to enhance the inperson class:** A largely unexpected benefit of preparing the OMSCS 7637 class was that it led to significant improvements in the teaching and learning of the inperson CS 7637/4635 course. While we had been incrementally revising the inperson class over the years, preparation of the online class made us reflect on the learning goals, outcomes, assessments and strategies for the course. The revisions to the learning goals, outcomes, assessments and strategies

naturally led to revisions of the contents of the course. As a result, we found improved teaching and learning in the inperson class in Fall 2014.

10. Conclusions

The cognitive systems paradigm of AI is both very old and very young. In one sense, the paradigm goes back to the earliest days of AI in the 1950s, with the twin goals of using our understanding of human mind to inspire the design of intelligent systems and using our understanding of intelligent systems for insights into the design of human mind. However, the number and proportion of AI courses in the cognitive systems paradigm at major universities has diminished significantly over the last generation and at present there are few online AI courses in the paradigm. Thus, there are only modest opportunities for new generations of students to learn about the paradigm. In another sense, the cognitive systems paradigm is newly resurgent. This evidenced by new funding, new centers of research, new conferences and journals, and growing numbers of students taking courses in the school of thought.

In this paper we presented the design, development, and delivery of an online course on cognitive systems. We also presented data from a quasi-experiment that compares the performance of the students in the online course with the performance of students in a similar inperson class. The data appears to indicate both that the students in the online course performed at least as well as the students in the inperson class, and the feedback from the online students was quite positive.

Based on this positive experience, we have taken several additional steps to promote teaching cognitive systems online. First, we taught the course again in Spring 2015 and Summer 2015, and are committed to teaching it in Fall 2015 and Spring 2016. Second, we have agreed to increase the size of the individual offerings of the course; while the Fall 2014 section was capped at 200 students and the Spring 2015 was capped at 300, the Summer 2015 section has a cap of 400. This will likely bring the course to ~1500 students over a 2-year period from Fall 2014 through Spring 2016. We will continue to push the limits of how many students we can successfully teach in one offering of the course. Third, we have opened up the course video material to the world: all KBAI lesson materials are now freely available to anyone (www.udacity.com/course/ud409). Courses at other colleges are welcome to use the materials in whole or part, and anyone in the world can also access these materials; these materials include all video lectures for the course, as well as the interactive microexercises and set of nanotutors. Fourth, in Spring 2015, we partnered with Georgia Tech Professional Education (GTPE) to offer a more open section of KBAI: whereas students must be accepted to the OMSCS program to take the KBAI section described here, anyone could join the GTPE section – and about 20 students did - complete the projects and assignments, receive feedback from TAs, and earn a verified certificate of completion. Fifth, in Fall 2015 we will again offer both the inperson and online sections of the course in parallel and repeat the quasi-experiment described in this article.

This is an exciting time to be teaching and learning about AI and cognitive systems. We can only hope that the availability of this KBAI course inspires more students to learn about cognitive systems, inspires other schools to increase their investment in cognitive systems education and research, and inspires other professors and colleges to develop their own online courses on artificial intelligence and cognitive systems.

Acknowledgements

We thank several colleagues for their support in designing and developing the KBAI: the Georgia Tech OMSCS team, especially David White, Zvi Galil, Charles Isbell, and Jillian Morn; and the Udacity Georgia Tech team, especially Aaron Gross, Jason Barros and Sebastian Thrun. We are grateful to other members of the teaching team who helped deliver the KBAI class in Fall 2014 and analyze the results during Spring 2015: Lianghao Chen, Amish Goyal, Xuan Jiang, Rishikesh Kulkarni, Sridevi Koushik, Rochelle Lobo, Shailesh Lohia, Nilesh More, Sriya Sarathy, and especially Joe Gonzales, the developer of the Peer Feedback tool used extensively in both the online and the inperson sections of the course.

References

- Bringsjord, S., & Schimanski, B. (2003). What is artificial intelligence? Psychometric AI as an answer. In *International Joint Conferences on Artificial Intelligence*. 887-893.
- Brown, A. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions in classroom settings. *Journal of The Learning Sciences*, 2(2), 141–178.
- Carpenter, P., Just, M., & Shell, P. (1990). What one intelligence test measures: a theoretical account of the processing in the Raven Progressive Matrices Test. *Psychological Review*, 97(3), 404.
- Collins, A. (1992). Toward a design science of education. In E. Scanlon & T. O’Shea (Eds.), *New Directions in Educational Technology*. 15–22. New York: Springer-Verlag.
- Connelly, D., & Goel, A. (2013) Paradigms of AI Programming in Python. In Proc. Fourth Symposium on Educational Advances in AI (EAAI-2013), Bellevue, Washington, July 2013.
- Goel, A. (1994) Teaching Introductory Artificial Intelligence: A Design Stance. In *Proc. 1994 AAAI Fall Symposium on Improving Introductory Instruction of Artificial Intelligence*, New Orleans, November 1994.
- Goel, A., & Davies, J. (2011) Artificial Intelligence. In *Handbook of Intelligence*, R. Sternberg & S. Kauffman (editors), Chapter 23, pp. 468-484, Cambridge University Press, 3rd Edition, 2011.
- Goel, A., Kunda, M., Joyner, D., & Vattam, S. (2013). Learning about Representational Modality: Design and Programming Projects for Knowledge-Based AI. In *Procs. Fourth AAAI Symposium on Educational Advances in Artificial Intelligence*.
- Hunt, E. (1974). Quote the raven? Nevermore! In L. W. Gregg (Ed.), *Knowledge and Cognition*. 129-158. Hillsdale, NJ: Erlbaum.
- Joyner, D., Bedwell, D., Graham, C., Lemmon, W., Martinez, O, & Goel, A. (2015) Using Human Computation to Acquire Novel Methods for Addressing Visual Analogy Problems on Intelligence Tests. In *Procs. Sixth International Conference on Computational Creativity*. Park City, Utah.
- Joyner, D. & Goel, A. (2015) Improving Inquiry-Driven Modeling in Science Education through Interaction with Intelligent Tutoring Agents. In Proc. 20th ACM Conference on Intelligent User Interfaces, Atlanta, Georgia, pp. 5-16.
- Kunda, M., McGreggor, K., & Goel, A. (2013) A Computational Model for Solving Problems from the Raven’s Progressive Matrices Intelligence test using Iconic Visual Representations. *Cognitive Systems Research*, 22:47-66.

- Kunda, M., Soulieres, I., Rozga, A., & Goel, A. (2013). Methods for Classifying Errors on the Raven's Standard Progressive Matrices Test. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*. 2796-2801. Berlin, Germany.
- Langley, P. (2011) Artificial Intelligence and Cognitive Systems, *AISB Quarterly*.
- Langley P. (2012) The cognitive systems paradigm. *Advances in Cognitive Systems* 1:3–13.
- Lehman, J., Laird, J., & Rosenbloom, P. (2006). A Gentle Introduction to Soar, An Architecture for Human Cognition: 2006 Update.
<http://ai.eecs.umich.edu/soar/sitemaker/docs/misc/GentleIntroduction-2006.pdf>
- Lovett, A., Tomai, E., Forbus, K. & Usher, J. (2009). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science* 33(7), 1192-1231.
- McGreggor, K., & Goel, A. (2014) Fractal and Ravens. *Artificial Intelligence* 215: 1-23.
- Nilsson, N. (1998) *Principles of Artificial Intelligence: A New Synthesis*. Morgan Kauffman.
- Norvig, P. (1992) *Paradigms of AI Programming*. Morgan Kauffman.
- Raven, J., Raven, J. C., & Court, J. H. (1998). *Manual for Raven's Progressive Matrices and Vocabulary Scales*. San Antonio, TX: Harcourt Assessment.
- Russell, S., & Norvig, P. (2009) *Artificial Intelligence: A Modern Approach*. 3rd edition. Prentice Hall.
- Shadish, W., Cook, T., & Campbell, D. (2002) *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin.
- Stefik, M. (1995) *Knowledge Systems*. Morgan Kauffman.
- Winston, P. (1993) *Artificial Intelligence*. Third Edition, MIT Press.