

you're not teaching, you have to be able to communicate your work.

Professor Sternberg was a phenomenal advisor. I couldn't have imagined a better experience. His advising and teaching, and the incredible care he took, along with other faculty, showed me how to work with a graduate student. I learned you really don't gain anything by rushing through material. You put whatever time you need into teaching. When I started teaching, I decided: This was how I learned—and this is how I'm going to teach. It's hard to start out as a faculty member—there are many competing pressures. When I started at Brown in 1986, I was told by a faculty member at Brown, "Don't spend too much time on your teaching." And I thought, "Don't tell me that, because taking time is the way it works."

ENGenious: What are your thoughts on engineering education?

Blume: We know students learn best

in a hands-on, interactive environment. We take a project-based approach, and give students calculation-based design experience early on. Lectures are a part of the process, but most of the real learning takes place outside the classroom, when they're building their projects. If students get enough personal attention, everything falls into place. I think the experience of connection to a professor, to the departments, to engineering as a whole, encourages students to stay and graduate. The goals are to help them feel welcome and supported, and give them opportunities.

ENGenious: What about students who are great at teamwork and project work but fare less well on written tests?




Blume: I used to write somewhat clever tests. The problems required a strong test-taking ability. I would think, "Oh, these are clever questions," and pat myself on the back. Students who were good test-takers would do fine, but I realized the tests

were sinking others. Yet, these failing test-takers were made for engineering! They were very good at projects and organizing, always did the work on time, always attended office hours. It was heartbreaking—I wanted them to succeed! So I decided to stop being Miss Clever-Exam-Writer-Woman and rewrote the tests. The new versions included opportunities for both quantitative and qualitative input, and for partial credit.

ENGenious: Did rewriting the tests make a difference?

Blume: Yes, but it's a delicate balance. We don't want to push people through and set them up for failure later on. We work to gradually bring people up to the level necessary for a career—not to get all *As*, but to give them the personal attention and the tools they need to succeed. We try to recognize differences in learning styles and levels of understanding. Everybody gets to the same point so they can do the design projects, but more advanced students spend less time on the fundamentals. In the end, everyone has what they need to do the projects—and they do them together. We take care of everyone.


ENGenious: Everything you say has time and care in it.

Blume: As James K. (Jim) Knowles, William J. Keenan Jr. Professor of Applied Mechanics, said many years ago when I thanked him for all of the time he put into us students, "That's why we're here."   

Janet Blume is Associate Dean of the Faculty and Associate Professor of Engineering at Brown University.

Visit research.brown.edu/research/profile.php?id=1106970186.

Learning from Data: How to Deliver a Quality Online Course to Serious Learners



Yaser S. Abu-Mostafa is Professor of Electrical Engineering and Computer Science at Caltech. His main fields of expertise are machine learning and computational finance. He is a recipient of the Richard P. Feynman Prize for Excellence in Teaching, and he has won multiple Caltech student teaching awards throughout his career. In 2005, the Hertz Foundation established a perpetual graduate fellowship named the Abu-Mostafa Fellowship in his honor. ENGenious interviewed him to learn more about his research and his approach to teaching.

ENGenious: Why did you decide to create "Learning from Data: Introductory Machine Learning Course," Caltech's first-ever live broadcast of an entire course?

Abu-Mostafa: From a Caltech perspective, it's a good way to provide public service. And it's an opportunity for people to understand what Caltech is about—to see inside a Caltech class. This is a real class, delivered as I always teach it, with real students attending. In my opinion, there are many people in the world who could take a Caltech course and do adequately well, perhaps just not at the level we expect of our doctoral students. That doesn't mean they shouldn't have access. To enhance accessibility, I adjusted the presentation material to fit the medium, but the con-

tent and delivery are the same. And many people completed the course. When we reach out in this global way, Caltech becomes less remote; we are approachable, yet we offer the highest-quality learning experience.

ENGenious: What do you mean when you say your course is a "real" Caltech course?

Abu-Mostafa: Before creating the course, I surveyed other online university, non-profit, and commercial courses. The measure of success for some seems to be the number of followers, and the desire to get more followers often led to lowering the bar for the course content. But I wanted to deliver the real thing for disciplined students with a



Professor Janet Blume testing exercise equipment built by students in project course.



Yaser Abu-Mostafa

serious approach to science. So I kept the online course at the exact level of my Caltech course. It's not a video game. And I've had positive feedback about this approach. Some have even donated to Caltech as a result. One alumnus who graduated decades ago but had not previously given to Caltech sent a check after viewing my course. I believe the high quality of the course is key.

ENGenious: Why did you choose the subject of machine learning?

Abu-Mostafa: Machine learning is my research area. It has theory, mathematics, and algorithms, and it also carries a wide range of applications in multiple domains. For instance, retailers want to anticipate clients' tastes and present choices they like. I recently consulted with a women's fashion company and ended up making recommendations to women I never met about fashion items I have never seen. My recommendations were preferred by customers over those of professional stylists! Do I need to know fashion to do this? No. The key is to

extract the correct information, which is based on the right data. The ability to impact such a wide variety of applications keeps me intrigued with the field, and makes the appeal of machine learning courses quite broad. Almost 150 Caltech students from 15 different options took the course this year. Machine learning also has profound theoretical questions that need answers and algorithms and new techniques under development.

ENGenious: What is machine learning?

Abu-Mostafa: Put simply, machine learning is a branch of computer science that enables computers to learn from experience. It makes computers "smarter" than humans for a broad range of tasks. The most critical components of any machine-learning system are the data. Machine-learning algorithms can take existing data, search for patterns, and make predictions based on those patterns. Whether we know it or not, we encounter this process in many ways: Web searches result in more useful

links, Internet shopping is tailored to our preferences, medical lab results are more accurate—even dating services are more likely to find you a potential partner.

Various machine learning paradigms exist, and each develops its own attributes. Supervised learning is one such paradigm, and the most common. For example, supervised learning is used in medical diagnosis. Researchers can "supervise" a machine's learning process to identify cancerous cells by "training" the computer with image data that includes cancerous or noncancerous cells. The algorithm will learn to apply certain cell attributes—shape, size, and color, perhaps—to identify malignant cells.

Another paradigm is called reinforcement learning ("trial and error"). For example, a roboticist can design an algorithm that experiments with different kinds of limb movements that mimic those of a human. The algorithm will learn which movements, such as a particular gait or grasping technique, are most efficient—and which are not. As the learning process develops, both we and the machines

learn correct actions for different situations: the best movements or actions are reinforced, and less reliable movements or actions are avoided.

The mathematical theory of machine learning primarily focuses on the problem of "overfitting" the data. We look for genuine connections that fit the data while avoiding patterns that cannot be trusted. Another interesting challenge is the temptation to throw too much computing power at a problem. How can more power hurt? If the algorithm is too aggressive—that is, if it is using too sophisticated a model to fit a limited data sample—it could mislead itself by detecting coincidental patterns in a sample that does not reflect a true association.

An important point to remember is that machine learning works only for problems that have enough data. Machine learning does not create information; rather, it gets the information from the data.

ENGenious: How are interested people accessing your online course?

Abu-Mostafa: The course is offered through iTunes U, YouTube, and of course the Caltech server, in many formats and multiple bandwidths. On iTunes U alone, there are more than 60,000 course subscribers. People register, submit homework, and are graded automatically, or they can grade themselves using solution keys. We reach a broad audience of different backgrounds, including a large international following. Furthermore, there's no language barrier, because YouTube features automatic translation capability.

Many people who are proficient in machine learning have watched the course and are intrigued by my approach to the topic. This carries far-reaching professional dividends at the intellectual level. A winner of the

National Medal of Technology took the course. A Caltech trustee took the course. There are postdoctoral groups who have taken the course together. If it's 100, 10,000, or 2 billion people, that's fine; my main mission has been achieved: delivering a quality course to serious learners.

ENGenious: What surprised you?

Abu-Mostafa: Other than how much time it took to prepare the slides and how tricky it was to design meaningful multiple-choice homeworks, the impact of the course on people who already know machine learning was surprising. I have some non-mainstream views in machine learning, and I completely polished my arguments and offered them through this course, which is a permanent record—not just for students, but also for my peers. When your peers buy into new ideas, new research follows, and this was an unexpected professional reward.

Also, a live online course had not yet been done at Caltech, so the stakes were high. I very much appreciated the Caltech community's strong support for this effort. They had unmitigated confidence that this would come out right. The Division of Engineering and Applied Science, the Information Science and Technology initiative, and the provost's office provided the funds, Information Management System and Services (IMSS) and the Academic Media Technology office provided the technical support, and many Caltech units, including the Alumni Association, took care of publicity. I received strong encouragement from everyone, and you need encouragement to go through such an intense experience.

ENGenious: What did you learn?

Abu-Mostafa: Robert Heinlein said, "When one teaches, two learn." The diversity of the online audience intro-

duced me to a deeper understanding of how people view and apply the material. But if I hadn't done this, I never would have learned the difference between real-time feedback in a classroom setting and the delayed feedback you get with videotaping. My subjective conclusions on how I did after class weren't always correct. When I viewed the videos, I learned how to adjust my style to accommodate the medium.

At the educational level, I learned that delivering a quality online course is incredibly time consuming. For example, I thought a white board wouldn't fit the medium. But the speed one normally writes on a board is about the same pace people can follow and understand the details—you don't lose your students. So I produced almost 3,000 incremental viewgraphs for the video to match a board-writing pace.

ENGenious: How will the next session be different?

Abu-Mostafa: I think this course is very much the way I want it to be. I've taught the course many times and have also written a book. I am happy with the way it came out, and I will continue to offer it online based on the recorded lectures as long as the material remains viable. It takes a huge time commitment and effort to create a new online course of the right quality.

ENGenious: What inspires you?

Abu-Mostafa: Doing the right thing. I know it sounds clichéd, but it's not necessarily the easiest thing to do. In this case, the outcome offsets all the difficulty. ■ ■ ■

Yaser Abu-Mostafa is Professor of Electrical Engineering and Computer Science.

Visit work.caltech.edu/telecourse.html.